UNIVERSITY OF TORINO

DIPARTIMENTO DI MATEMATICA GIUSEPPE PEANO

Ph.D. in Modeling and Data Science

Cotutelle with the University of Pau and the Adour Region,
Laboratory of Mathematics and its Applications of Pau

Final dissertation, 19th December 2024

Statistical and Probabilistic Approaches to Hydrological Data Analysis: Rainfall Patterns, Copula-like Models and First Passage Time Approximations

University of Turin
Supervisor: Elvira Di Nardo
Co-supervisor: Stefano Ferraris

Candidate: Tommaso Martini

University of Pau and the Adour Region
Supervisor: Ivan Kojadinovic

ACADEMIC YEAR 2024/2025

Examining Committee

The following served on the Examining Committee for this thesis.

- Francisco de Asís Torres Ruiz, Ordinary Professor, Departamento de Estadística e Investigación Operativa, Universidad de Granada (referee and jury member)
- Fabrizio Durante, Ordinary Professor, Dipartimento di Matematica e Fisica "Ennio De Giorgi", Università del Salento (referee and jury member)
- Carlo De Michele, Ordinary Professor, Dipartimento di Ingegneria Civile e Ambientale, Politecnico di Milano (examiner and jury member)
- Elisa Perrone, Assistant Professor, Department of Mathematics and Computer Science, Eindhoven University of Technology (examiner and jury member)
- Elvira Di Nardo, Associate Professor, Dipartimento di Matematica "Giuseppe Peano", Università degli Studi di Torino (supervisor and jury member)
- Ivan Kojadinovic, Professor, Laboratoire de Mathématiques et de leurs Applications de Pau (LMAP), Université de Pau et des Pays de l'Adour (supervisor and jury member)

Summary

Analysis of rainfall data and subsequent modelling of the many variables concerning rainfall is fundamental to many areas such as agricultural, ecological and engineering disciplines and, due to the complexity of the underlying hydrological system, it relies heavily on historical records. Daily rainfall series are arguably the most used. In this context, we initially investigate the modelling of daily rainfall interarrival times through a family of discrete probability distributions known as the Hurwitz-Lerch-Zeta family, along with two other distributions which are deeply related to the latter and have never been considered with this aim. Building up on the relationships between the interarrival times and certain other temporal variables, fundamental in describing the alternation between periods of continuous rainfall and periods of drought, we delineate a methodology particularly well-suited for statistical applications. The latter procedure and the fitting performance of the aforementioned distributions is shown on a dataset composed of a variety of rainfall regimes.

Additionally, the multivariate modelling of rainfall variables has never been more important, as a perceivable shift in the inter-relationships between these variables could reflect climate changes in a region. In this context, copulas are well known and valued for their flexibility. However, they lose their charm when dealing with discrete random vectors. In this case, the uniqueness of the copula is compromised, leading to inconsistencies which basically break down the theoretical underpinnings of the inferential procedures commonly used in the continuous case. Recently, Gery Geenens made a compelling case for a new approach, grounding its beliefs in historical ideas regarding contingency tables. The theoretical insights he gives, coupled with a computational tool known as iterative proportional fitting procedure, open up the path to our development of novel (semi-parametric or parametric) models for finitely supported bivariate discrete random vectors. With this aim, we shall prove a sklar-like decomposition of a bivariate discrete probability mass function

between its margins and a *copula probability mass function*. Statistical models hinging upon this representation are built and related inferential procedures are studied both theoretically and empirically.

Of the same significance as modelling the behaviour of rainfall is its impact on water bodies and land surfaces. For instance, understanding the time it takes for rainfall to cause river levels to exceed a flood stage is of paramount importance for flood prediction and management. More in general, it is often crucial to determine the time at which certain hydrological thresholds are crossed by some hydrological quantity. When the latter's value in time is modelled by a stochastic process, this problem can be restated in terms of the first passage time. In this context, a practical computation of the first passage time probability density and distribution function is a delicate issue. Within this framework, we propose an approximation method based on a series expansion. Theoretical results are accompanied by discussions on the computational aspects. Extensive numerical experiments are carried out for the geometric Brownian motion and the Cox-Ingersoll-Ross process, showing the usefulness of the proposed method.

Resumè

L'analyse des données de précipitations et la modélisation des nombreuses variables associées sont fondamentales dans des domaines tels que l'agriculture, l'écologie et l'ingénierie. En raison de la complexité du système hydrologique sous-jacent, ces analyses reposent fortement sur des archives historiques, avec les séries de précipitations quotidiennes étant parmi les plus utilisées. Dans ce contexte, nous étudions initialement la modélisation des temps d'inter-arrivée des précipitations quotidiennes à travers la famille de distributions de probabilité discrètes connue sous le nom de Hurwitz-Lerch-Zeta, ainsi que deux autres distributions étroitement liées qui n'ont jamais été considérées à cette fin. En nous appuyant sur les relations entre les temps d'inter-arrivée et d'autres variables temporelles essentielles pour décrire l'alternance entre les périodes de pluie continue et de sécheresse, nous élaborons une méthodologie particulièrement adaptée aux applications statistiques. Cette procédure et la performance d'ajustement des distributions mentionnées sont démontrées sur un ensemble de données couvrant divers régimes de précipitations.

De plus, la modélisation multivariée des variables de précipitations est devenue cruciale, car un changement perceptible dans leurs interrelations pourrait refléter des modifications climatiques régionales. Les copules sont bien connues pour leur flexibilité, mais elles perdent de leur efficacité avec les vecteurs aléatoires discrets, compromettant leur unicité et entraînant des incohérences théoriques. Récemment, Gery Geenens a proposé une nouvelle approche basée sur des idées historiques liées aux tableaux de contingence. Ses perspectives théoriques, combinées à la procédure d'ajustement proportionnel itératif, nous ont conduits à développer de nouveaux modèles (semi-paramétriques ou paramétriques) pour des vecteurs aléatoires discrets bivariés à support fini. À cette fin, nous démontrons une décomposition de type Sklar d'une fonction de masse de probabilité discrète bivariée entre ses marges et une fonction de masse de probabilité copule. Des modèles statistiques basés sur cette représentation sont construits, et les procédures

d'inférence associées sont étudiées théoriquement et empiriquement.

D'une importance égale à la modélisation du comportement des précipitations est leur impact sur les masses d'eau et les surfaces terrestres. Comprendre, par exemple, le temps nécessaire pour que les précipitations fassent dépasser les niveaux des rivières au-delà d'un seuil de crue est essentiel pour la prévision et la gestion des inondations. Plus généralement, déterminer le moment où certains seuils hydrologiques sont franchis par une quantité hydrologique donnée est souvent crucial. Lorsque cette valeur est modélisée par un processus stochastique, le problème se reformule en termes de temps de premier passage. Dans ce contexte, le calcul pratique de la densité de probabilité et de la fonction de distribution du temps de premier passage est délicat. Nous proposons une méthode d'approximation basée sur un développement en série. Les résultats théoriques sont accompagnés de discussions sur les aspects computationnels. Des expériences numériques étendues sont menées pour le mouvement brownien géométrique et le processus de Cox-Ingersoll-Ross, démontrant l'utilité de la méthode proposée.

Acknowledgements

I can only start by thanking my supervisor, Elvira Di Nardo, or simply "prof," as I have never stopped calling her since the first day of my PhD. She has always been by my side, allowing me to grow both academically and personally. Our great adventure actually began with my master's thesis: her seriousness and dedication in guiding me encouraged me to continue our happy collaboration. This evolved as we delved into the then-unknown world of rain modelling, while never straying far from our initial topic of first passage times. Speaking of which, I take this opportunity to also thank Giuseppe D'Onofrio. He was not only an excellent collaborator but also a friendly source of advice and help in navigating the academic world. Thanks to his sociability, I also made new, interesting, and fun acquaintances (especially some PhD students from Campania).

Similar words can be said about my co-supervisor, Stefano Ferraris, whom I have often affectionately referred to, when talking with friends, as "the rain man." It is thanks to him and his knowledge that I entered this new world during my PhD. Additionally, he introduced us to the very talented, friendly, and unmistakably Sicilian Giorgio Baiamonte, Carmelo Cammalleri, and Carmelo Agnese, with whom we formed fruitful collaborations.

Most of this work was carried out within the welcoming walls of Palazzo Campana, so welcoming that I decided to stay there until the end of my PhD. Therefore, I also thank the Department of Mathematics of Turin, but more precisely room 16, where I spent many days: studying, writing, and drinking coffee with my dear colleagues. After a bitter start in this regard, due to the post-Covid period, things changed after a certain Sandro walked through the door of the old office where I used to stay alone. I don't even remember if he introduced himself or not, but I do remember that his first words were: "Shall we grab a coffee?" That was when I realised it would be the start of a good friendship. I also thank him immensely for introducing me, albeit quite late, to my primary lunch source: the "mensa" (university cafeteria),

which I will defend to the death. Then, Chiara joined us, with whom we shared the adventure of the summer school in Ferrara, the old Marco, Lidia, Davide, Margaux, Salvatore (who made the office a liveable place), and more recently Daniel, with whom I know I can always go out to eat or drink something. I also thank my colleagues in the computer science department, with whom I equally shared good times. The cohesion among colleagues was further fostered by the PhD program director, Laura Sacerdote, whom I also thank for her patience in always answering bureaucratic questions and managing the PhD program despite her countless commitments. I remember once accidentally seeing her Google calendar: it scared me quite a bit.

In the initial project of my PhD, probably inappropriately, I included the term "copulas." A statistical object initially unknown to me but one that immediately caught my attention: interesting from a mathematical standpoint and also essential in applications. Their presence in my project led to the chain of events that resulted in a joint supervision with the University of Pau in France. In search of someone who could help me untangle the issue of copulas in the discrete case, we turned to Ivan Kojadinovic, who promptly became my French supervisor. My collaboration with Ivan was excellent. He guided me in this new field, allowing me to produce more theoretical work. He often sat next to me to help me understand what was working and what was not in my work. He always emphasised the need to fight against our natural laziness, which might lead us to finish work prematurely. There was no shortage of laughs, either. During the co-supervision, I also spent a total of nine months in Pau, some of which were spent having lunch at the Pau "mensa" (which I brought along...) with Ivan. He also taught me some indoor rock climbing, which generated the most laughter, caused by the sight of me trying to tie the necessary knots.

During my time abroad, I was also fortunate to feel very comfortable within the university environment in Pau. Although I was practically alone in the French office, a strong and supportive network of PhD students in Pau allowed me to meet many new people in a short time. I'll be honest: I had never shaken hands with so many people from various countries around the world, which greatly enriched my cultural experience. Among these, I must absolutely mention Cherif and Jinping. The first, whom I met by chance when I realised we were both newcomers to that small city, accompanied me through most of the moments during my time abroad. Similarly, the second, whom I met through other Chinese friends, proved to be a trustworthy friend and a gateway to a new culture, full of remarkable kindness and respect. I

am confident our friendships will endure.

Speaking of friendships, I also thank the friends I met in Turin through various adventures, including Lorenzo, Davide, Elena, and my now "Londoner" friend Riccardo, who hosted me several times for refreshing breaks from work. Additionally, I cannot forget Stefania and Pietro, who introduced me to rock climbing, the only thing that can frustrate me more than research. Lastly, I thank my housemates, who were getting tired of always seeing the same graph on my computer screen when they walked into my room.

I also thank the friendships that have accompanied me for a lifetime: my friends from Acqui Terme, my hometown. Making a list would be unfair, and I am sure we will continue to share each other's milestones.

I also deeply thank my family. They have always supported and helped me in every possible way. My father even tried to understand what I was doing at times, giving me invigorating pats on the back, while my mother and sister were well aware of why I spent hours staring at a screen or scribbling.

Lastly, I must thank Benedetta. Even though we were often physically far apart, for me, we were never really distant: I always felt her support and constant affection, which were fundamental for my journey. These feelings were heightened by the fact that she is following the same path, which has probably allowed us to understand each other even more, as we will continue to do.

Ringraziamenti

Non posso che iniziare ringraziando la mia supervisor Elvira Di Nardo, o semplicemente "prof", come non ho mai smesso di chiamarla dal primo giorno del dottorato. E' sempre stata dalla mia parte e mi ha permesso di crescere sia accademicamente che umanamente. L'origine della nostra grande avventura risale in realtà alla mia tesi magistrale: la sua serietà e dedizione nel seguirmi mi hanno spinto a continuare la nostra felice collaborazione, che si è evoluta addentrandoci in questo mondo allora sconosciuto della modellizzazione della pioggia, ma allo stesso tempo non si è mai distaccata dall'iniziale topic relativo ai tempi di primo passaggio. Nel nominare quest'ultimi, colgo l'occasione per ringraziare anche Giuseppe D'Onofrio. Non è stato solo un ottimo collaboratore, ma anche un'amichevole fonte di consigli e aiuti per muoversi all'interno del mondo accademico. Inoltre, grazie alla sua socialità ho anche potuto fare delle nuove conoscenze interessanti e divertenti (specialmente alcuni dottorandi campani). Parole simili a quelle usate per Elvira si possono spendere per il mio co-supervisor Stefano Ferraris, a cui mi sono spesso amichevolmente riferito, parlandone con amici, come "l'uomo delle piogge". Infatti è grazie a lui e alla sua conoscenza che sono entrato per il periodo del dottorato in questo, per me, nuovo mondo. In aggiunta, ci ha fatto entrare in contatto con i bravissimi, simpaticissimi e sicilianissimi Giorgio Baiamonte, Carmelo Cammalleri e Carmelo Agnese, con cui sono nate delle fruttuose collaborazioni.

La maggior parte di questi lavori si sono svolti all'interno delle accoglienti mura di Palazzo Campana, così accoglienti che ho deciso di rimanerci fino al termine del dottorato. Dunque, ringrazio anche il dipartimento di matematica di Torino, ma più precisamente quella stanzetta 16 dove ho passato parecchie giornate: a studiare, scrivere, e prendere caffè con i miei cari colleghi. Dopo un inizio amaro da questo punto di vista, dovuto ad un periodo post covid, le cose sono cambiate dopo che un certo Sandro ha varcato la porta del vecchio ufficio dove alloggiavo da solo. Non mi

ricordo neanche se si sia presentato o no, ma solo che le sue prime parole furono: "ci prendiamo un caffè?" e così capii che sarebbe stato l'inizio di una buona amicizia. Inoltre, lo ringrazio immensamente per avermi fatto scoprire, ammetto veramente tardi, la fonte primaria dei miei pranzo: la mensa, che difenderò a spada tratta fino alla morte. Dopodichè, si aggiunsero Chiara, con cui abbiamo condiviso l'avventura della summer school Ferrara, il vecchio Marco, Lidia, Davide, Margaux, Salvatore (colui che ha reso l'ufficio un luogo vivibile) e più recentemente Daniel, con cui so sempre di poter andare a mangiare o bere qualcosa. Ringrazio anche i colleghi a informatica, con i quali ho ugualmente passato dei bei momenti. La coesione tra colleghi è stata inoltre fomentata dalla direttrice del corso di dottorato Laura Sacerdote, che ringrazio anche per la pazienza nel rispondere sempre alle domande di carattere burocratico e nel gestire il dottorato, nonostante gli infiniti impegni. Mi ricordo di aver per caso visto una volta il suo google calendar: mi sono spaventato parecchio.

Nel progetto iniziale del dottorato, probabilmente a sproposito, avevo inserito il nome *copule*. Un oggetto statistico a me inizialmente sconosciuto, ma che aveva attirato subito la mia attenzione: era interessante dal punto di vista matematico, ma anche fondamentale nelle applicazioni. Dalla loro presenza nel mio progetto ha avuto origine la catena di eventi che ha portato a finalizzare una cotutela con l'università di Pau in Francia. In cerca di qualcuno che potesse aiutarmi a districare un po' il problema delle copule nel caso discreto, abbiamo deciso di rivolgerci a Ivan Kojadinovic, che è di fatto poi immediatamente diventato il mio supervisor francese. La collaborazione con Ivan è stata ottima. Mi ha guidato in questa nuova materia e grazie a lui sono riuscito a produrre dei lavori più teorici. Spesso si è seduto accanto a me per farmi capire cosa andasse bene e cosa no nei prodotti del mio lavoro. Ha sempre sottolineato come sia necessario combattere la nostra naturale pigrizia, che ci porterebbe a terminare troppo precocemente dei lavori. Non sono mancate anche le risate. Infatti, per la cotutela ho passato un totale di 9 mesi a Pau, parte dei quali anche passati a pranzare nella mensa (me la sono portata dietro...) di Pau con Ivan. Mi ha anche insegnato un po' di arrampicata sportiva in palestra, ed è qui dove si sono generate le più forti risate: erano causate dalla visione del sottoscritto mentre provava a fare i necessari nodi.

Durante il periodo all'estero ho anche avuto la fortuna di trovarmi molto bene nell'ambiente universitario di Pau. Nonostante fossi praticamente solo nell'ufficio francese, l'esistenza di una solida e solidale rete di dottorandi a

Pau mi ha permesso di conoscere moltissime persone nuove in poco tempo. Sarò onesto: non avevo mai stretto le mani di così tante persone da svariati paesi del mondo, un fatto che ha contribuito a rendere questa esperienza una grandissima aggiunta al mio bagaglio culturale. Tra questi, voglio assolutamente nominare Cherif e Jinping. Il primo, conosciuto per caso quando mi sono accorto che eravamo entrambi dei nuovi arrivati in quella piccola città, mi ha accompagnato in più o meno in tutti i momenti vissuti nel periodo all'estero. Allo stesso modo il secondo, conosciuto attraverso altri amici cinesi, si è rivelato un amico fidato e una porta su una nuova cultura, fatta di una gentilezza e di un rispetto encomiabili. Sono sicuro che le nostre amicizie saranno conservate in futuro.

A proposito di amicizie, aggiungo anche ringraziamenti agli amici conosciuti a Torino per varie peripezie, fra cui Lorenzo, Davide, Elena e l'amico ormai "londinese" Riccardo che mi ha ospitato più volte per delle belle pause dal lavoro. Inoltre, non possono mancare Stefania e Pietro che mi hanno fatto scoprire l'arrampicata, l'unica cosa che può farmi frustrare di più della ricerca. Mancano ancora i miei coinquilini, che erano ormai stanchi di vedere sempre lo stesso grafico sul mio pc quando entravano in camera mia.

Ringrazio poi le amicizie che mi accompagnano da una vita: gli amici di Acqui Terme, la mia città natale. Fare una lista sarebbe ingiusto, e soprattutto, sono sicuro che continueremo a vivere assieme i rispettivi traguardi.

Ringrazio anche immensamente la mia famiglia. Mi hanno sempre supportato e aiutato in ogni modo a loro possibile. Mio padre ha anche cercato di capire cosa stessi facendo, qualche volta, dandomi delle rinvigorenti pacche sulle spalle, mentre mia madre e mia sorella erano ben consapevoli del motivo per cui passavo ore a fissare uno schermo o a scarabocchiare.

Infine, non mi rimane che ringraziare Benedetta. Anche se spesso fisicamente lontani, per me non ci siamo mai allontanati di molto: ho sempre sentito il suo supporto e il suo affetto costante e fondamentale per il mio percorso. Sensazioni acuite dal fatto che lei stia percorrendo la stessa strada, e questo ci ha probabilmente permesso di capirci ancora di più, come continueremo a fare.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

This thesis is dedicated to the exploration and development of some theoretically sound tools connected by the underlying unifying thread of hydrological data analysis and modelling. It is structured into three main parts. The first part focuses on the daily rainfall occurrence process, exploring the modelling of daily rainfall interarrival times and related temporal rainfall variables using a family of discrete probability distributions known as Hurwitz-Lerch-Zeta distributions, along with two additional distributions related to the latter, whose application in this context is new. The performance of these distributions, immersed in a broader statistical procedure concerning the aforementioned temporal variables, is shown on a novel and interesting dataset composed of 6 Italian and British stations spanning a variety of rainfall regimes. The second part uses the multivariate modelling of rainfall variables as a starting point to shift to a more theoretical tone, exploring new possibilities in the still relatively unexplored context of copula-like modelling of discrete random vectors. Finally, the third part, motivated by the problem of estimating the time at which some hydrological thresholds are reached, proposes a practical method for approximating the first passage time densities of some stochastic processes. The next sections will provide a general introduction to all the aforementioned topics.

1.1 Rainfall Modelling

1.1.1 Daily Rainfall Data

Due to the complexity of hydrological systems, their analysis and modelling relies heavily on historical records. Rainfall historical records are of various time scales, from hourly data to annual data. However, daily rainfall series are

arguably the most used information in environmental, climate, hydrological, and water resources studies (Serinaldi, 2009). Rain gauge networks measuring rainfall amount at a daily timescale operate in almost every country, and often provide the only available input for climate and hydrological analyses. A reliable and flexible single site model is the fundamental starting point of any more complex multi-site model taking into account the spatial correlations arising when observing a dense network of stations. Following Serinaldi (2009) and McMahon and Srikanthan (2001), single site daily rainfall models can be classified in the following way: (1) two-part models accounting separately for occurrence (wet/dry) and amount processes; (2) transition probability matrix models resorting to multistate Markov processes (3) resampling models based on nonparametric bootstrap and analogous techniques; and (4) time series models of the Autoregressive Moving Average (ARMA) type. Additionally, these models are crucial for developing *rainfall generators* (see Wilks, 1999b, for a review). Indeed, rainfall records are often too brief to conduct reliable and meaningful analyses, and, to generate longer alternative rainfall scenarios that are statistically consistent with observed data, stochastic models are commonly used to simulate synthetic series.

Given the ever-growing interest in modelling and analysing the alternation between period of continuous rainfall and periods of drought, the first part of this thesis is devoted to exploring and improving some recent advances in models of type (1), such as the ones contained in Agnese et al. (2014). The latter is a strong starting point for establishing the usefulness of a particular family of discrete probability distributions in describing the rainfall temporal variables involved in models of type (1). These temporal variables are introduced in the following section, together with the rainfall depth variable which is used to describe the amount of rain in the rainy periods.

1.1.2 Univariate Models of Rainfall Temporal Variables

Rainfall manifests one peculiar characteristic which is common to many other geophysical processes: intermittence (Wiscombe et al., 1994). Intermittence is found in variables which are related to the internal and external structure of rainfall. The most commonly seen for the external structure are the dry spells ds and wet spells ws, meaning the sequences of rainy days and non-rainy days. For the internal structure the usually considered variable is the rainfall depth h in the wet periods (Bernardara et al., 2007). Developing well-fitting models for these variables is crucial for constructing single-site models for rainfall occurrence and establishing marginal models that can

later be appropriately combined with copulas to extend the single-site model further. In the following section, we review the modelling approaches that form the foundation of the first part of this thesis.

Wet Spells, Dry spells and Chains

As already mentioned, at a local scale, a conventional method for addressing intermittency in rainfall records involves the statistical analysis and modelling of wet spells ws and dry spells ds, usually under the assumption of their independence. In his seminal investigation, Chatfield (1966) scrutinized a brief series of daily rainfall data from a single station in Kew (London) and examined the relationships between observed frequencies of increasing durations of wet spells. Chatfield (1966) identified a nearly constant probability that a wet day is succeeded by another rainy day. By the modelling point of view, this behaviour can be replicated by the memoryless property enjoyed in the discrete case solely by the geometric distribution. The latter has since been extensively employed to depict the distribution of wet spells in numerous studies (see, e.g. Kottegoda and Rosso, 1997; Racsko et al., 1991; Zolina et al., 2013). Additionally, Chatfield (1966) noted an inclination of the rainfall process towards persistence in the dry state, with the probability of consecutive dry days increasing with their preceding count. This prompted the adoption of the log-series distribution to model dry spells, owing to its escalating ratios of subsequent occurrences. While both geometric and log-series distributions have historically been utilized Green (1970) and remain prevalent for inferring the probability laws governing wet and dry spells (Chowdhury and Beecham, 2013; El Hafyani and El Himdi, 2022), their general applicability has been questioned by some scholars. For instance, Wilks (1999a) advocated for the use of a mixed geometric distribution to model wet spells in the United States, while Deni et al. (2010) demonstrated the efficacy of the compound geometric distribution in Peninsular Malaysia. Moreover, mixed distributions have been observed to perform adequately for both dry and wet spells in various contexts (Dobi-Wantuch et al., 2000; Deni and Jemain, 2009).

More recently, Agnese et al. (2014) suggested to model both ws and ds in a parsimonious way, by deriving their distribution from the one obtained by investigating the probabilistic law of the so-called interarrival times it, representing the series of times elapsed between two subsequent rainy days. The latter approach makes up an important part of this work and therefore it is of great interest to recall how the it have been treated in the literature,

especially concerning their role in describing the rainfall occurrence process.

The Interarrival Times

If we suppose that the interarrival times between rainy days are independent and identically distributed (i.i.d.), one natural way to model them is through the well known theoretical framework of renewal processes (Buishand, 1977). The fundamental assumption of a renewal process is that it probabilistically re-starts at each time of arrival, the so-called "renewal property". The simplest renewal process, the Bernoulli process, basically underlies the work Chatfield (1966), for example. Before expanding on possible extensions of the Bernoulli process which are more suitable to the discrete nature of temporal variables arising from daily rainfall measurements, it is worth to mention its very well known continuous counterpart: the Poisson process, which is widely used for its simple mathematical tractability. An underlying Poisson process implies that the interarrival times are independent and exponential in distribution. For instance, within the domain of eco-hydrology, Laio et al. (2001) advocated for the suitability of the exponential distribution in modelling daily rainfall data observed in specific locations in Texas. Conversely, Rodriguez-Iturbe et al. (1987) conducted an inference analysis on hourly rainfall data from Denver, Colorado, and concluded that direct reliance on Poisson process-based models offers limited utility. They proposed an alternative approach employing a rectangular pulse model to better encapsulate the underlying physics of the rainfall process. However, this alternative model exhibited inadequacies in replicating the distribution of dry period durations. Consequently, Rodriguez-Iturbe et al. (1988) proposed an extension to the model, with a primary emphasis on short-term prediction. Despite demonstrating favorable performance across various aggregation periods ranging from 1 hour to 24 hours, this extended model is characterized by a substantial number of parameters and poses challenges in data fitting using conventional techniques such as maximum likelihood and method of moments.

However, when modelling rainfall data at a daily scale, it is more natural to treat *it* as a discrete random variable. For example, Foufoula-Georgiou and Lettenmaier (1986) demonstrated the statistical advantages of treating the rainfall occurrence process as a discrete process, rather than a continuous one. That is why we now return to the simple Bernoulli process as a starting point. The latter's assumptions imply the that the interarrival times are independent and geometric in distribution. Geometrically distributed interarrival times

further imply that the rain probability is constant at any time, i.e., it is independent from the time elapsed from the last rainy day. Alternative methodologies capable of incorporating time-varying rain probabilities while preserving the renewal property exist. For instance, previous studies in hydrology have employed a straightforward Markov chain model for daily rainfall occurrence (Buishand, 1978; de Groen and Savenije, 2006; Gabriel and Neumann, 1962). In this model, the probability of rainfall on any given day is contingent upon the weather status (rain or no rain) of exclusively the preceding day. A more generalized approach involves relaxing the Markovian assumptions by introducing a discrete process with after-effects, wherein the rain probability at any time is influenced by the historical record since the last rainy day, albeit maintaining the renewal property. One such method entails the adoption of the logarithmic-series distribution, wherein the associated rain probability exhibits a monotonically decreasing trend with the time elapsed since the last rainy day (Gupta et al., 1997). Agnese et al. (2012) extended this concept by employing the polylogarithmic-series distribution to analyse it frequencies derived from daily rainfall data collected in Sicily and Piedmont. They noted that, while this distribution generally outperforms the logarithmic-series distribution, limitations persist, particularly during the 'warm season' (April to September). It was then natural to consider a family of discrete distributions which contains the aforementioned distributions as special cases. The three parameter family Hurwitz-Lerch Zeta (HLZ) satisfies the required property and shows to be a step forward with respect to other commonly seen distributions, as concluded in the recent works of Agnese et al. (2014) and Berro et al. (2019). Working on a dataset of Sicilian and Piedmont stations respectively, they show it is able to faithfully replicate statistical characteristics of interarrival times derived from rainfall data, such as very high standard deviation and skewness, and frequencies having a maximum at $it = 1$, followed by a monotone decrease towards a remarkable tail.

The HLZ family of discrete distributions is also considered the starting point of the second main idea developed in Agnese et al. (2014): jointly conveying the modelling of ws and ds by deriving their distributions from the one of it. Agnese et al. (2014) showed that both the ws and ds distributions can be easily derived from the it distribution, under the assumption that rainfall interarrival times are i.i.d. Indeed, geometrically distributed ws directly arise from the latter hypothesis on it, whereas the distribution of ds follows the same probabilistic law adopted for fitting the it probabilities,

albeit with a shifted support. However, an imposition of the geometric distribution for the ws could turn out to be restrictive, as already discussed in the previous part of this introduction. If the geometric distribution does not fit ws correctly, this implies that the rainfall probability varies within the rainfall event and two separated models for ws and ds may be needed for a reliable evaluation of both quantities. In other words, the modelling of ws and ds distributions separately may allow a relaxation of the i.i.d. hypothesis on it. At the same time, it must be noted that modelling first the it sample has the clear advantage of retrieving the probability distribution of both ws and ds from a single fitting, often reducing the number of parameters involved. One contribution of this thesis is then an extension of the approach of Agnese et al. (2014) mentioned above, which is developed in our work Baiamonte et al. (2024), where we conducted an in depth analysis of the theoretical and empirical differences between modelling separately ws and ds and modelling first it and deriving ws and ds as a byproduct. Empirical results will be presented when applying this methodology on a novel dataset of daily rainfall records composed of 6 Italian and British stations spanning a variety of rainfall regimes. This application will show that a comparison of the two procedures can shed light on some aspects of the data involved. The same approach is also employed for two additional time variables, deeply related to ws and ds: the wet chains wch, as previously introduced by Berro et al. (2019), and the dry chains dch, seemingly never thoroughly investigated in the literature before. These variables extend the concept of wet and dry spells to sequences characterised by an interruption of one no-rainy or one rainy day, respectively. They represent two quantities that may be of interest for practical hydrological applications.

Additionally, the scientific literature on the statistical inference of rainfall interarrival times is still rather sparse, hence another contribution of this work is providing further evidence of the suitability of the Lerch family to reproduce it frequencies in a wide range of rainfall regimes, possibly encouraging further applications of the methodology. Within the same context, to expand the range of options for fitting daily rainfall interarrival times data, we will also explore two additional distributions that are closely related to the Lerch family: the Poisson-stopped HLZ distribution, recently introduced in Ong et al. (2020), and the one-inflated HLZ distribution, which we construct following the approach in Gupta et al. (1995). Notably, neither of these distributions has been previously considered for modelling interarrival times. Hence, we shall report a description of the accurate

fittings of these distributions on interarrival times arising from the above mentioned dataset of rainfall measurements, accompanied by a comparison with their performance with the HLZ distribution, presenting some of the results contained in our work Agnese et al. (2022).

Rainfall Depths

Aa already stated, a main feature strictly related to the internal structure of rainfall and fundamental for modelling the rainfall process is the depth (or the intensity) h of the rainy days (Bernardara et al., 2007). In the literature, Yang et al. (2020); Porporato et al. (2006) rainfall depths are more often treated as continuous despite that sometimes these models fail to account for the time discreteness of the sample process (Foufoula-Georgiou and Lettenmaier, 1986). Moreover, daily rainfall depth measurements are almost always performed by automatically counting how many times a small bucket corresponding to 0.2 mm is filled. The latter causes an abundance of ties in the data, which led us to treat the rainfall depth h as a discrete random variable. According to this choice, another contribution of this thesis is showing the satisfactory results of fitting the three parameter HLZ family of discrete distributions and the PSHLZ distribution to h observations from the previously mentioned rainfall records, based on the findings of our article Agnese et al. (2022).

1.1.3 Bivariate Modelling of Hydrological Variables

The multivariate study of hydrological features is fundamental for aspects such as design and management purposes, where univariate frequency analysis proves to be unsatisfactory (Grimaldi and Serinaldi, 2007). Moreover, to be able to fully describe rain as a climate feature, multivariate analysis is often needed, as for example a rainfall hyetograph is completely characterized if peak intensity, volume event, duration, and peak time are modeled together. In the past, bivariate frequency modelling of rainfall features has usually been carried out following standard multivariate distributions commonly found on textbooks (see, for instance, Johnson et al., 2000), such as bivariate exponential and bivariate normal. However, these multivariate distributions provide margins of the same family. This clearly has strong limitations for rainfall data modelling, considering for example that many rainfall characteristics present a long tail and thus are not gaussian. The concept of copulas has then proved to be fundamental in this context. Indeed,

thanks to Sklar's theorem (Sklar, 1959), they let the joining of marginals from arbitrary families into a joint distribution having them as marginals and, conversely, one can extract a copula from a given multivariate distribution in order to explore the dependence structure given by the original multivariate distribution. Quite a few examples can be found in literature, such as Baets et al. (2008); Onof et al. (2011); Serinaldi (2009), and many others.

However, there has not been extensive work in the literature on the bivariate modelling of the interarrival times it and the subsequent rainfall depths h or of the well spells ws and the dry spells ds. Allowing the possibility of dependence between these variables would greatly increase the flexibility of a single site model for the daily rainfall process. When tackling these problems in the particular context of daily data, one would be facing two possible scenarios

1. one discrete (a daily temporal variable) and one continuous random variable (if we consider h as continuous): *mixed scenario*;

2. two discrete random variables (two temporal variables or one temporal variable and h as a discrete random variable): *discrete scenario*.

Both of them are not standard for copula modelling. Indeed, Sklar's Theorem does not guarantee uniqueness of the copula in the two above cases, leading to inconsistencies, especially in the context of statistical inference. This problem is non trivial and of mathematical, statistical and applied interest. The second main part of this thesis is devoted to the discrete scenario. Following the pioneering work Geenens (2020), we have developed a copula-like model and related inferential techniques for the bivariate discrete scenario, with a more general scope then the hydrological context just described, which nevertheless remains an application of undoubted importance. The next section is dedicated to introducing this line of research.

1.2 Copula-Like Models for Discrete Bivariate Random Vectors

The well known theorem of Sklar (1959) provides a fundamental basis for the extensive application of copulas, since it enables the separation of a random vector's dependence structure from its marginal components. Consequently, copulas have become an extremely useful tool for modelling dependence

between random vectors, valued for the flexibility that they are able to provide in this context: through them, a practitioner is able to separately model the margins and the dependence. Parametric copula families are naturally employed to achieve this and ad-hoc inferential procedures have been developed. However, this framework loses its charm when dealing with discrete random vectors. Indeed, Sklar's theorem affirms the uniqueness of the copula solely on the cartesian product of the ranges of the marginal distribution functions. Hence, the uniqueness of the copula is compromised in the case of discrete margins, leading to possible inconsistencies which complicates the inference of copulas for discrete random vectors, basically breaking down some of the theoretical underpinnings of the inferential procedures mentioned above (see Genest and Nešlehová, 2007, Section 2). More precisely, only a so-called subcopula is identifiable when the margins are discrete. A unique subcopula can then be extended in an infinite number of ways to a copula (see, for example, de Amo et al., 2017). Nevertheless, a practitioner could still decide to postulate a classical parametric copula model and proceed with the same inference tools which are used in the continuous case, as is advocated for a bivariate Bernoulli distribution in the nice essay on copulas for count data Genest and Nešlehová (2007), albeit with some required modifications and precautions. However, Faugeras (2017) argues that while the former technique of extending a subcopula is clearly subject to the arbitrary choice of the user, the latter parametric path could additionally lead to various issues of identification. To cite one, sampling fluctuations could lead to invalid estimated parameters (for example, an estimated $|\hat{\theta}_n| > 1$ in the case of the Farlie-Gumbel-Morgenstern family of copulas).

Furthermore, it is worth to mention here the literature involving the study of *discrete copulas* (see, e.g., Perrone et al., 2019, and the references therein), whose name would suggest a deep connection with our present context. Following (Durante and Sempi, 2015, Definition 3.1.7), discrete copulas are subcopulas with uniform grid domains. They have been investigated for their interesting geometrical properties: they admit a representation as convex polytopes, and this characterisation has been found to be extremely useful in building maximum entropy *checkerboard* copulas (Perrone et al., 2019, Section 3.1) which, interestingly enough, are mostly applied in climatology and hydrology (see AghaKouchak, 2014, for a review). Additionally, *empirical copulas*, which form the base of rank-based multivariate statistics analysis, are discrete copulas themselves, and this fact adds to the important role

that the latter may have in applications. However, as for instance can be seen in Durante and Perrone (2020), both the aforementioned applications of discrete copulas have found their use in the continuous case as far as statistical aspects may be concerned.

To sum up, even though a small number of strategies to deal with copula modelling in the case of discrete margins has been proposed in the literature, it still remains fragmentary and seemingly inconclusive. Recently, Gery Geenens in its seminal paper Geenens (2020) made a compelling case for an approach which takes its roots in a more fundamental view of what the role of a copula conceptually is: a *margin-free* representative of the dependence between the elements of a random vector. Equivalently, we could say that a copula embodies the information necessary to recover a joint distribution when only the margins are given. He argues that this basic idea can be suitably adapted to the discrete case and he develops his arguments in the context of finitely supported discrete bivariate random vectors, grounding its beliefs in historical ideas regarding the statistical analysis of contingency tables. A contingency table, informally speaking, is a matrix representation of a probability mass function (pmf) associated to a bivariate discrete random vector. This simple connection made it clear that investigating such a literature would provide interesting insights to build upon. Geenens (2020) starts by recalling the works of Yule (1912) and Goodman and Kruskal (1954), where it is advocated that adjusting frequencies of rows and columns of contingency tables to (discrete) uniform would allow an easier comparison of the association structures between tables. Following a suggestion dating back to Mosteller (1968), which believed that a contingency table should be decomposable into its two margins and a "nucleus" of association, he then introduces equivalence classes of dependence. This concept translates the choice of a representative of the dependence of a bivariate random vector into the selection of a unique representer of the appropriate equivalence class of dependence. Then, under the (arbitrary) choice of uniform margins as the most "margin-less" property to uniquely identify this representer, the usual concept of copula can fit this view. In the context of continuous random vectors, the probability integral transform is well known as the transformation which lets one "extract" the unique underlying copula from a given continuous multivariate distribution. However, it does not work in the discrete scenario. When restricting the focus on finitely supported discrete bivariate random vectors, another, albeit somewhat hidden, contribution of Geenens (2020) is then to propose the concept of I-projection (in the sense

of Csiszár, 1975) on a Fréchet class of pmfs with uniform margins (that is, the set of pmfs having uniform margins) as a way to associate a unique *copula pmf* u to a given (bivariate) pmf p. Informally, given a probability distribution p of interest and a set $\mathcal{M}$ of probability distributions, computing an I-projection consists in finding an element of $\mathcal{M}$, if it exists, that is the "closest" to p in the sense of the Kullback–Leibler (or information) divergence. When conducted on a Fréchet class, such projections turn out to preserve the *odds ratio matrix* which encodes the dependence in the considered discrete setting; see Agresti (2013, Section 2.4), Kateri (2014, p 43), Rudas (2018, p 123) or Geenens (2023, Section 4). In practice, the I-projection of a pmf on a Fréchet class can be carried out using the iterative proportional fitting procedure (IPFP), also known as Sinkhorn's algorithm or matrix scaling in the literature. The IPFP takes the form of an algorithm (actually of several equivalent algorithms) whose aim is to adjust the elements of an input matrix so that it satisfies specified row and column sums. The former conceptual and theoretical insights, coupled with the latter practical tool, opened up the path to the development of the novel (semi-parametric or parametric) models for finitely supported bivariate discrete random vectors which forms one of the main contribution of this thesis and has been investigated in our article Kojadinovic and Martini (2024).

As the first fundamental building block in this regard, the work of Geenens (2020) prompted us to explore the possibility of constructing a Sklar-like decomposition for a finitely supported bivariate pmf p. Thanks to the underlying concept of I-projections, we shall prove such a property. Afterwards, statistical models are built hinging upon this novel copula-like decomposition which, similarly as what happens in the framework of copula modelling for continuous random vectors (see, e.g., Hofert et al., 2018, Chapter 4 and the references therein), enjoy the flexibility of separately specifying the margins and the dependence, ideally governed by a copula pmf u. The latter characteristic is what separates our proposal from more traditional methods for modelling p, such as log-linear models (see, e.g., Agresti, 2013; Kateri, 2014; Rudas, 2018) and association models (see, e.g., Goodman, 1985; Kateri, 2014). Naturally, we then propose and investigate inferential procedures and goodness of fit techniques, providing a detailed study of their asymptotic properties. In this context, a result on the differentiability of I-projections on Fréchet classes shall be presented, which, on one hand, is fundamental in proving asymptotic properties of the statistical procedures we discuss, and, on the other hand, holds an independent importance, since it seems to have

never been treated before in the literature. Extensive numerical simulations will provide more insights on the provided theoretical results. It is crucial to emphasise that these analyses are conducted under the assumption of strict positivity of the initial unknown pmf p. Indeed, I-projections may not always exist, and to ensure the existence of the aforementioned copula-like decomposition of p, we will additionally assume p has rectangular support (that is, strictly positive). Interestingly, this assumption serves as the discrete counterpart to the commonly made assumption of strict positivity of the copula density within the interior of the unit square when modelling multivariate continuous distributions using copulas. It is worth noting that, as will be elaborated upon in our concluding remarks, the assumption of rectangular support for p could be substituted with alternative conditions, provided that certain practical challenges are managed. Fortunately, this assumption of rectangular support aligns well with many applications.

The differentiability result mentioned above is the starting point for another principal contribution of this thesis, based on the contents of our recently submitted work Geenens et al. (2024). A well-known class of alternatives to the Kullback–Leibler divergence (containing the latter) are the so-called ϕ-divergences (see, e.g., Ali and Silvey, 1966; Csiszár, 1967; Liese and Vajda, 1987; Csiszár and Shields, 2004, and the references therein) and ϕ-projections are merely the analogs of I-projections based on ϕ-divergences. As previously mentioned, a key step in providing asymptotic results for certain estimators of a copula pmf u is proving that I-projection on a Fréchet class of pmfs with fixed arbitrary (positive) margins are differentiable in a certain sense. We shall provide an extension of this result to the more general context of ϕ-divergences. Similar findings (though not differentiability results *per se*) can be found in Jiménez-Gamero et al. (2011, Section 2). In the case where $\mathcal{M}$ consists of probability vectors derived from a given parametric distribution and p is assumed to belong to $\mathcal{M}$, these findings include the well-known asymptotic properties of minimum ϕ-divergence estimators (see, e.g., Read and Cressie, 1988; Morales et al., 1995; Basu et al., 2011, and references therein). Since the results of Jiménez-Gamero et al. (2011) were derived without requiring p to belong to $\mathcal{M}$ (thereby accommodating possible model misspecification), they provide a foundation for developing a range of inference procedures related to goodness-of-fit testing and model selection, as detailed in Sections 3 and 4 of Jiménez-Gamero et al. (2011) (see also references therein).

The main objective of our developments in this context is to attempt to

unify and extend the previous results. Building on the approaches of Gietl and Reffel (2013, 2017), we consider finite measures on finite spaces, rather than restricting our attention solely to probability measures on finite spaces. With a slight strengthening of the main condition from Jiménez-Gamero et al. (2011) and a crucial additional assumption missing from that reference, we first present a general result on the continuous differentiability of ϕ-projections under some appropriate conditions. By additionally assuming the convexity of the set $\mathcal{M}$, we show that the previously mentioned conditions can be derived from simpler conditions that are straightforward to verify. From a statistical inference perspective, such results enable the immediate determination of the consistency and the asymptotic distribution of ϕ-projection estimators. Simulation examples shall provide an example of the practical application of the theoretical results provided.

Finally, we make note that when aiming to study the asymptotics of ϕ-projections and/or minimum ϕ-divergence estimators beyond the finite discrete setup, another line of research involves exploiting the dual representation of ϕ-divergences (see, e.g., Keziou, 2003; Broniatowski and Keziou, 2006, 2009). This more complex approach is out of the scope of this thesis, as it appears unnecessary for the discrete finite setting under consideration.

1.3 First Passage Times in Hydrology

The previous sections have introduced examples of tools from the field of probability and statistics developed to help in understanding the behaviour of rainfall. However, the discipline of hydrology encompasses various aspects such as the distribution, movement, and quality of water, both above and below the Earth's surface. Therefore, of the same significance of the behaviour of rainfall is its impact on water bodies and land surfaces. Closely related to the latter is a slightly more general problem: modelling the time at which certain hydrological thresholds are crossed by some hydrological quantity. When the latter's value in time is modelled by a stochastic process, the problem mentioned above can be restated in terms of the well known first passage time (FPT) problem (see, for instance, Stechmann and Neelin, 2014). As an example, consider a river basin where the water level follows a stochastic process influenced by continuous rainfall and other hydrological inputs. The FPT would then represent the time required for the water level to rise to a flood stage for the first time and knowing the distribution of the FPT would serve as a valuable tool for modelling the likelihood and timing

of potential flood events. This could be achieved by computing the FPT probability density function (pdf) and/or cumulative distribution function (cdf). In general, providing a closed form of the latter two is known to be a delicate issue. Exploring a practical way for doing so for some one dimensional diffusion processes is the third main contribution of this thesis. The next section introduces this line of research.

1.3.1 FPT Time Density Approximation

One-dimensional diffusion processes play a key role in the description of fluctuating phenomena belonging to different fields of applications as physics, biology, neuroscience, finance and others (Karlin and Taylor, 1981; Øksendal, 1998). These models are described by stochastic differential equation of the following type

$$\mathrm{d}Y_t = M(Y_t)\mathrm{d}t + \Sigma(Y_t)\,\mathrm{d}W_t, \quad Y_0 = y_0,$$

where the drift coefficient M and the diffusion coefficient Σ are real functions such that the above equation admits a unique solution with continuous trajectories and satisfying the Markov property. Here, $\{W_t\}_{t\geq 0}$ denotes a standard Wiener process and $y_0 \in \mathbb{R}$ is the initial condition.

In particular, the class enjoying a linear drift, that is

$$\mathrm{d}Y_t = (-\tau Y_t + \mu)\,\mathrm{d}t + \Sigma(Y_t)\,\mathrm{d}W_t, \quad Y_0 = y_0, \tag{1.1}$$

where $\tau > 0$ and $\mu \in \mathbb{R}$, is widely used for its mathematical tractability and flexibility. The volatility $\Sigma(Y_t)$ determines the amplitude of the noise and, according to its dependence on Y_t, it characterizes families of stochastic processes which are solution of (1.1). If

$$\Sigma(Y_t) = \sqrt{aY_t^2 + bY_t + c}, \quad a,b,c \in \mathbb{R}$$

the solution of (1.1) is called Pearson diffusion process (Forman and Sørensen, 2008). The coefficients a,b and c are such that the square root is defined for all the values of the state space (y_1, y_2) of Y_t, with $-\infty \leq y_1 < y_0 < y_2 \leq +\infty$. A wide range of well-known processes belongs to this class $(\sigma > 0)$:

- Ornstein-Uhlenbeck process: $a = b = 0, c = \sigma^2$ and $\Sigma(Y_t) = \sigma$;
- Inhomogeneous geometric Brownian motion: $a = \sigma^2, b = c = 0$ and $\Sigma(Y_t) = \sigma Y_t$;
- Jacobi diffusion: $a = -\sigma^2, b = \sigma^2$ and $\Sigma(Y_t) = \sigma\sqrt{Y_t(1-Y_t)+c}$;

- Feller process (CIR model): $b=\sigma^2, a=0$ and $\Sigma(Y_t)=\sigma\sqrt{Y_t+c}$.

Throughout this thesis we will focus on the geometric Brownian motion (GBM) and on the Feller process. The former has been selected for its mathematical tractability, which leads to the possibility of computing a closed form of the FPT pdf. In the hydrological context it has been used, for example, to model river flows (Lefebvre, 2002). The latter, albeit not as simple as the GBM, has been considered for its variety of applications not only in a biological context (Ditlevsen and Lansky, 2006; Feller, 1951; Lansky et al., 1995) but also in survival analysis, in the modelling of nitrous oxide emission from soil and in other applications such as physics and computer science (see Ditlevsen and Lansky, 2006, and references therein). In the mathematical finance it is known under the name of Cox-Ingersoll-Ross model (CIR) (Cox et al., 1985). For simplicity, we shall use the latter acronym as we continue. In this thesis we consider the dynamics of Y_t until it crosses a threshold S for the first time, the so called (upcrossing) FPT, defined as

$$T := \inf\{t \geq 0 : Y_t \geq S \,|\, 0 < y_0 < S\}.$$

Many contributions in the literature (Giorno et al., 1986; Going-Jaeschke and Yor, 2003; Linetsky, 2004; Masoliver and Perelló, 2014) focus on computing the Laplace transform (LT) of the pdf $g(t) := g(t \,|\, y_0, S)$ of T, namely

$$\tilde{g}(z) = \int_0^\infty e^{-zt} g(t) \mathrm{d}t, \quad z > 0. \tag{1.2}$$

The literature emphasizes the computation of the LT of g because the pdf itself is typically not known analytically and neither can be derived through direct inversion of (1.2). Nevertheless, from $\tilde{g}$ we can compute the probability of crossing the threshold S, that is $\mathbb{P}(T \,|\, y_0) = \int_0^\infty g(t)\mathrm{d}t$, and the mean FPT, that is $\mathbb{E}[T]$, in the following way

$$\mathbb{P}(T \,|\, y_0) = \tilde{g}(z)\big|_{z=0} \quad \text{and} \quad \mathbb{E}[T] = -\left.\frac{\mathrm{d}\tilde{g}(z)}{\mathrm{d}z}\right|_{z=0}$$

Moments of T of any orders can be computed using higher derivatives of $\tilde{g}$, when they exist. As is widely recognized, the moments of T provide valuable insights into the statistical properties of g and of FPT events. A different strategy to get the moments of T is using the transition pdf of the process defined as $f(y,t \,|\, y_0, \tau) = \frac{\partial}{\partial y}\mathbb{P}(Y_t < y \,|\, Y_\tau = y_0)$ for $y \in (y_1, y_2)$ and $t \geq 0$. Indeed, if $\{Y_t\}_{t\geq 0}$ admits a stationary distribution $\phi(y) = \lim_{t\to\infty} f(y,t \,|\, y_0, 0)$

independent of y_0, the Siegert formula (Siegert, 1951) allows us to compute the moments of T as

$$\mathbb{E}[T^n] = n \int_{x_0}^{S} \frac{2\,\mathrm{d}z}{[\Sigma(Y_t)]^2 \phi(z)} \int_{-\infty}^{z} \phi(x) \mathbb{E}\left[T^{n-1}\right] \mathrm{d}x, \quad n \in \mathbb{N}^+.$$

While they can be employed for the GBM, both the depicted strategies are impractical to compute the moments of T for a CIR process. In this case, although a closed-form expression for $\tilde{g}$ exists, computing its higher derivatives is cumbersome. Consequently, research has often concentrated on evaluating only the mean and variance of T (Ditlevsen and Lansky, 2006; D'Onofrio et al., 2018), or, at most, the third moment (Giorno et al., 1988). Similar computational challenges arise when calculating moments of T using the Siegert formula, despite the fact that the stationary distribution for a CIR process is known to be a shifted gamma distribution. Simulations of the underlying stochastic process paths using Monte Carlo methods remain an effective tool for obtaining manageable estimates of g, which are particularly valuable for analysing asymptotic properties. Another strategy involves writing the FPT distribution as a Sturm-Liouville eigenfunction expansion series, initially presented for the CIR process in Linetsky (2004), using the classical approach of Kent (1980) and Kent (1982). Although this strategy provides an expression for the FPT density, in these references information on the moments of T could be obtained only numerically and the procedure refers exclusively to diffusion processes without natural boundaries. A discussion on the FPT of the CIR process in the presence of entrance, exit and reflecting boundary at the origin is given in Martin et al. (2011), solving the Sturm-Liouville boundary problem in the case $\tau = 0$.

Recently, an important step in the FPT problem for the CIR process came from Di Nardo and D'Onofrio (2021), where the authors managed to compute closed form formulae for the cumulants of T of any order for the CIR process regardless of the nature of the boundaries. They relied on the particular form of $\tilde{g}$ for the CIR process, which is the ratio of two power series whose algebra is simplified if we consider $\log \tilde{g}$. They took advantage of the formal power series algebra (Charalambides, 2002) to give first a closed form expression of the cumulants. Recall that if T has moment generating function $\mathbb{E}\left[e^{zT}\right] < \infty$ for all z in an open interval about 0, then its cumulants $\{c_k(T)\}_{k \geq 0}$ are such that

$$\sum_{k\geq 1} c_k(T)\frac{z^k}{k!} = \log \mathbb{E}\left[e^{zT}\right]$$

for all z in some (possibly smaller) open interval about 0. Cumulants have nice properties compared with moments such as the semi-invariance and the additivity (McCullagh, 1987).

Overdispersion and underdispersion as well as asymmetry and tailedness of the FPT pdf might be analysed through the first four cumulants. Examples on how to employ the first four cumulants in the estimation of the parameters of a model fitted to data is given in (Antunes et al., 2020; Seneta, 2004). The employment of cumulants in the FPT literature is not new (Ramos-Alarcón and Kontorovich, 2013). However, their application has been limited to few cases and not in the direction addressed for the first time in Di Nardo and D'Onofrio (2021). Indeed, from the cumulants, one can easily recover the moments. The knowledge of the latter is the key to developing a Laguerre series expansion of the FPT pdf g which, when truncated to an order $n>0$, provides an approximant $\hat{g}_n$ of g. This approach for evaluating the FPT pdf of the CIR process seems to have been investigated only recently in Di Nardo and D'Onofrio (2021). Another main contribution of this thesis is then expanding on the latter reference and it is based on our works Di Nardo et al. (2023) and Di Nardo et al. (2024). We review the theoretical underpinnings of the methodology and we use them to provide sufficient conditions for the existence of the required series expansion when specifically considering the FPT context. Moving to the more practical aspects, an efficient evaluation of $\hat{g}_n$ is proposed through an iterative algorithm and corresponding stopping criteria are provided, improving over the classical convergence-based stopping criteria, which are commonly seen in the literature along with simple graphical checks (see, e.g., Provost and Ha, 2016). Among other considerations on computational aspects of the method, we discuss the possibility of correcting $\hat{g}_n$ when it is negative, a delicate issue which is known in literature, but, at the best of our knowledge, still requires attention. Additionally, all the previous aspects are also examined for the FPT cdf, which was not previously considered in Di Nardo and D'Onofrio (2021).

The possibility of constructing a bonafide approximating pdf opens the path to the development of interesting applications, seemingly never investigated in this framework. For instance, a novel acceptance-rejection type algorithm which exploits the form of $\hat{g}_n$ is proposed. The distinctiveness of

our proposal lies in its innovative use of the functional form of the series representation of the unknown density. Although acceptance-rejection methods have been utilised in FPT contexts, as highlighted by Herrmann and Zucca (2019) and Mijatović et al. (2015), they have been built with remarkably different techniques and, additionally, they have not been applied specifically to the CIR process. Furthermore, this approach is particularly beneficial given the lack of exact simulation methods for CIR sample paths.

Another novel application consists in an approximated maximum likelihood estimation procedure which has the aim of estimating the parameters of the underlying stochastic process. In this case, we show how the proposed approximation can be used to adapt the well known maximum likelihood estimation to the case where a closed form of the FPT density is not available, but the FPT moments are known and a sample of FPTs is available.

Noteworthily, both the mentioned applications are particularly amenable to possible extensions in a more general setup, as shall be explained in the concluding remarks which end this thesis.

An in depth analysis of the efficacy of the method and of the applications is conducted through various numerical experiments for the GBM, where the knowledge of the true FPT pdf enables a direct comparison with the approximant $\hat{g}_n$, and for the CIR, where the absence of a closed form of g creates the perfect scenario to show the usefulness of the proposed method.

1.4 Structure of the Thesis

This thesis is organised as follows. The first part of Chapter 2 is based upon the theoretical sections of our work Agnese et al. (2022). It is dedicated to reviewing the HLZ distribution (HLZD) and expressing its properties in a more compact way, eventually providing some new insights. For example, a new result on its convolution is provided. In order to further enlarge the choices for fitting daily rainfall interarrival times data, we also provide details on the Poisson-stopped HLZ distribution (PSHLZD), whose properties are expressed by exploiting the known Bell Polynomials, and on the one inflated HLZ distribution (OIHLZD). Afterwards, certain other temporal variables closely related to the daily rainfall interarrival times, namely the wet and dry spells and the wet and dry chains, are introduced. It is then shown how their distribution can be recovered from the one of the interarrival times and viceversa. Building up on these relations, a recently proposed methodology for the modelling and empirical analysis of these temporal variables is then

detailed, expanding on the contents of the methodological sections of our article Baiamonte et al. (2024). More in detail, firstly a procedure called direct method (DM) is presented, where the distribution of wet spells and dry spells (as well as of the corresponding chains) is derived as a consequence of the assumption of i.i.d. interarrival times. Note that in this case, the wet spells will have a geometric distribution. Secondly, the latter assumption is relaxed by using an indirect method (IM) where wet spells and dry spells are modelled separately, hence including the possibility of accounting for a non-constant rain probability inside a rainfall cluster.

Chapter 3 presents the empirical results obtained by using the procedure previously detailed on daily rainfall data, arising from measurements at 6 Italian and British stations which span a variety of rainfall regimes and were never investigated before in the literature. It is based on providing the details of the empirical analyses presented in our works Agnese et al. (2022) and Baiamonte et al. (2024). In the first part of this chapter, we describe and compare the fits of the HLZD, the PSHLZD and the OIHLZD on the interarrival times and the rainfall depths obtained from the above mentioned dataset. In the second part, we apply and compare the DM and IM on the same data. The results when using the DM highlight how the geometric distribution does not always reasonably reproduce the ws frequencies, even when it are well fitted by the Lerch distribution. Improved performances are shown to be obtained with the IM, owing to the relaxation of the assumption on the independence and identical distribution of the interarrival times. A further improvement on the fittings is obtained when the datasets are separated into two periods, suggesting that the inferences may benefit for accounting for the local seasonality.

Motivated by the problem of describing the dependence between some of the variables encountered in the previous chapters, in Chapter 4 we explore the theoretical development of copula-like models for discrete random vectors found in our work Kojadinovic and Martini (2024). We begin by recalling the ideas contained in the essay on dependence between discrete random vectors Geenens (2020), laying out a framework where a copula pmf plays the role of a representative of the dependence between two finitely supported discrete random variables. We will highlight how the main ingredient of this construction is the concept of I-projection (in the sense of Csiszár, 1975) on a Fréchet class of pmfs with fixed arbitrary (positive) margins. By using the

latter tool, we are able prove a Sklar-like decomposition for a bivariate pmf p into its two margins and a unique copula pmf u. Under the additional assumption that p has a rectangular support, this decomposition is then exploited to build a statistical model for p and ultimately the inference annd goodness of fit testing of the underlying unique copula pmf u is tackled. The main tool in proving asymptotic results for parametric and nonparametric estimation procedures for u is the delta method. As a result that is both key for the application of the latter and of independent interest, we prove that I-projections (in the sense of Csiszár, 1975) on a Fréchet class of pmfs with fixed arbitrary (positive) margins are differentiable in a certain sense. Theoretical results are complemented by finite-sample experiments and a data example. Unfortunately, the latter does not concern rainfall data. However, such a possible application shall be discussed in the future work section at the end of the thesis.

Chapter 5 is based on our recently submitted work Geenens et al. (2024) and develops some of the ideas found in Chapter 4 in a different direction. Indeed, one of the main contributions of Chapter 5 is an extension of the above mentioned differentiability result to the more general context of ϕ-divergences for finite measures on finite spaces. To this aim, we define ϕ-divergences in the latter more general context, recall their main properties as stated in Gietl and Reffel (2017), provide conditions under which they are strongly convex in their first argument, and define ϕ-projections. Subsequently, we present conditions under which ϕ-projections are continuous and continuously differentiable, and demonstrate that these can be replaced by considerably simpler conditions when the set $\mathcal{M}$ onto which one wishes to ϕ-project is convex. When the target set for the ϕ-projection is convex, we demonstrate that the necessary assumptions can be derived from easier conditions which are particularly amenable to verification. Given this, we show that for many common choices of ϕ-divergences, ϕ-projections are automatically continuously differentiable when $\mathcal{M}$ is a subset defined by linear equalities, a context typically arising in applications. We conclude the chapter by showing how these findings can be easily used to derive the asymptotics of ϕ-projection estimators (i.e., minimum ϕ-divergence estimators), providing examples for projections onto parametric sets of probability vectors, sets of probability vectors with specified fixed moments, and Fréchet classes of bivariate probability arrays.

After briefly returning to the context of probabilistic modelling in hydrology, in Chapter 6 a method for approximating the first passage time probability density and/or distribution function of some stochastic processes is detailed, giving the necessary details for what has been proposed in our works Di Nardo et al. (2023) and Di Nardo et al. (2024), on which this chapter is based upon. This approximation is obtained by truncating a series expansion involving the generalised Laguerre polynomials and the gamma probability density, and it relies on the knowledge of the moments, or equivalently, of the cumulants of the FPT random variable. We begin by providing the necessary background on the underlying theory. After particularising the latter to the FPT context, we provide the theoretical results fundamental for applying the procedure outlined previously to a FPT density g and FPT cdf G, such as sufficient conditions for the existence of the proposed series expansion and the study of the order of convergence. Then, the issue of the actual computation of the approximants is tackled. An iterative algorithm with suitable stopping criteria and theoretically sound corrections for the possible negativity of the approximants are some of the key aspects that are then explored. Afterwards, the FPT problem for the GBM and the CIR process are recalled, focusing on the known results useful for carrying out the proposed approximation. Given these, extensive numerical examples covering different shapes of the FPT pdfs and cdfs are presented for both the GBM and the CIR process aiming to discuss the strengths and weaknesses of the proposed approach. At last, three applications are considered: in addition to the previously mentioned acceptance rejection like algorithm and approximate maximum likelihood estimation, it will be shown how the laguerre gamma series estimation can give rise to a known estimator when only a sample of FPTs is available and neither the density nor the moments of T are known.

Finally, after summarising our contributions, in Chapter 7 we provide possible extensions that could lead to future projects.

Before the beginning of each chapter, a disclaimer will indicate which of the five published articles and one preprint, co-authored by the author of this work, it is primarily based upon.

A brief note regarding notation shall be given to the reader before commencing to go through the manuscript. As it is clear by the introduction and the structure of the thesis just delineated, the topics treated in this

work are linked by an unifying thread but undoubtedly remain somehow heterogeneous and may refer to different branches of literature. This is the reason why each chapter may propose its own notational section with slight changes with respect to the preceding ones, with the aim of abiding to the customs of the corresponding literature practices.

All simulations, numerical experiments or computations in the thesis were carried out using the R statistical environment (R Core Team, 2024) or the Mathematica software (Wolfram Research, Inc., 2023).

Chapter 2

Modelling Rainfall Interarrival Times and other Temporal Variables

This chapter is based on Agnese et al. (2022)

Modelling the frequency of interarrival times and rainfall depths with the poisson hurwitz-lerch zeta distribution, *Fractal and Fractional* 6(9)

and on Baiamonte et al. (2024)

Applying different methods to model dry and wet spells at daily scale in a large range of rainfall regimes across Europe, *Advances in Statistical Climatology, Meteorology and Oceanography*, 10(1), 51–67.

Rainfall data are usually available in the form of cumulative amounts over disjoint equispaced time intervals, and, as already mentioned in the introduction, the most available data are daily measurements. In this context, one way to model the rainfall process is to define an event as a day with measurable precipitation and consider a suitable discrete-time stochastic process to describe the probabilistic structure of the occurrence of rainfall. A description of the process giving the amount of rain on the rainy days is then provided separately. Consider $\{h_n\}_{n\in\mathbb{N}}$ to be the series of (possibly zero) rainfall amounts and let $h^* > 0$ be a rainfall threshold. Then, the binary series $\{x_n\}_{n\in\mathbb{N}}$ indicating the rainfall events as prescribed by h^* is defined by $x_n = \mathbf{1}_{\{h_n \geq h^*\}}$ for all $n \in \mathbb{N}$. A classical and basic choice for describing the mechanism which generates the relevant events is that of a renewal

process (see, e.g., the monograph Cox, 1970). The key series in describing a renewal process is the sequence of interrarrival times $\{it_n\}_{n\in\mathbb{N}}$ corresponding to the series of times passed between each event in $\{x_n\}_{n\in\mathbb{N}}$. Indeed, the defining property of a renewal process is that it "probabilistically" restarts at each event and it is formalised by imposing that the the interrarrival times $\{it_n\}_{n\in\mathbb{N}}$ are independent and identically distributed (i.i.d.). Then, a reasonable description of the rainfall process can be given by the bivariate discrete-time process $(it_n, h_n^*)_{n\in\mathbb{N}}$, which consists of the interarrival times and the subsequent rainfall depths greater than h^*. We shall consider that it_n is independent from h_n^* for all $n \in \mathbb{N}$. As the modelling of the latter is somehow secondary in the scope of this thesis, we shall briefly discuss it here before focusing more on the former.

Rainfall Depths and the Rainfall Amount Process

Over the course of the first two chapters of this thesis, we shall make two simple assumptions regarding $\{h_n^*\}_{n\geq 0}$. Firstly, we will consider that h_n^* are i.i.d., even though relaxations of one or the other restrictions can be found in some works, but mostly for other smaller time scales. Indeed, as explained in the comprehensive review Wilks (1999b), most stochastic weather generators often assume that precipitation amounts on wet days are independent and follow the same distribution. Adjusting the model to consider different probability distributions for precipitation based on the day's sequence within a wet spell (e.g., higher rainfall on a wet day after another wet day than after a dry day) has been studied by, for instance, Katz (1977), Buishand (1977, 1978) and Chin and Miller (1980). However, this additional complexity generally does not significantly affect the outcomes. Similarly, the autocorrelation between successive nonzero precipitation amounts is often statistically different from zero but is usually small and of little practical relevance (Katz, 1977; Buishand, 1977, 1978; Foufoula-Georgiou and Lettenmaier, 1987). On the other hand, accounting for the serial correlation of nonzero precipitation amounts becomes crucial when the precipitation model operates on an hourly (or smaller) time scale rather than daily (Katz and Parlange, 1995).

Additionally, as already said in the introduction, we define the common rainfall depth h as a discrete random variable. It is well established in the literature (see, for instance, Yang et al., 2020; Porporato et al., 2006) that rainfall depths are typically treated as continuous variables. However, daily rainfall depth measurements are often recorded by automatically counting

the number of times a small bucket, corresponding to 0.2 mm, is filled. This measurement process frequently results in a high number of ties in the data, which motivated us to treat the variable h as discrete. Additionally, this approach allows us to demonstrate the considerable flexibility of the discrete distributions that will be examined later in this chapter.

Interarrival Times

When the common distribution of the interarrivals is assumed to be geometric we are in the case of the Bernoulli renewal process. However, as well known, the geometric distribution enjoys the memoryless property. Consequently, it forces the conditional probability of having an event at time k given that it has not happened before to be constant for any $k \in \mathbb{N}$. The latter value is also known as failure rate. It is clear that this may be restrictive, especially with the aim of modelling the rainfall occurrence, where a constant failure rate would translate into a rainy day probability which does not take into account the past. A correction can be obtained by adopting a distribution for which the associated rain probability is a function of the time elapsed from the last rainy day. For instance, in the case of the logarithmic-series distribution such a relationship is monotonically decreasing. Actually, both the geometric and the logarithmic-series distributions are part of a larger three parameter family of discrete distributions known as the Hurwitz-Lerch-Zeta (HLZ) family. The full three parameter HLZ distribution (HLZD) has been recently proposed in Agnese et al. (2014) and in Berro et al. (2019) as a model for interarrival times data arising from rainfall measurements at the daily scale. In other words, they assumed that the rainfall events are generated following the mechanism of a renewal process whose interarrivals are i.i.d. with common law the HLZ distribution. The latter has proven to provide satisfactory fittings as it is able to replicate both the numerous occurrences of the value equal to one, which represent the uninterrupted sequences of rainy days, and some large values scattered over time and responsible for drought phenomena.

As a tentative step forward, in Agnese et al. (2022) we proposed to model the common it rv using the Poisson-stopped HLZD (PSHLZD), which is strictly related to the HLZD. This discrete distribution can model an excess of zeroes (paralleling the excess of $it = 1$) and can present a remarkable tail (Liew et al., 2020). The PSHLZD has been used in the same reference for comparisons with the negative binomial distribution, a popular model for fitting over-dispersed count data. Indeed, the PSHLZD can be seen both as

a Poisson-stopped sum of HLZD's as well as a generalisation of a negative binomial distribution. The Poisson contribution allows us to model the superposition of i.i.d. HLZD's in the observed time series as rare event.

This chapter is structured as follows.

We start the following section by briefly recalling the Bell Polynomials and some of their properties. The combinatorial aspects of these polynomials let us rewrite and prove known properties of the forthcoming probability distributions in a compact and modern way, in addition to providing some new minor results. Then, after establishing some notation and stating well known definitions, the HLZ family of discrete probability distributions is introduced and reviewed. We recall some of its properties, including expressions for moments and cumulants. New results on its log-concavity and convolution are provided. The PSHLZD is introduced immediately after, where the Bell Polynomials will be more involved. Finally, the One Inflated HLZ (OIHLZ) is dealt with in detail. All the descriptions of the distributions are accompanied by a brief explanation of how the maximum likelihood estimation (MLE) of their parameters can be carried out.

Afterwards, we introduce and describe some additional temporal variables which are linked to interarrival times and provide an alternative and slightly more complex way to model the rainfall occurrence process. We provide the probabilistic relationships between the interarrival times and the new introduced variables. Given these connections, two procedures for describing the rainfall occurrence process, particularly adapt to statistical applications, are proposed.

2.1 Bell Polynomials in a Nutshell

We refer the reader to Charalambides (2002) for an extensive treatment of the Bell polynomials and their role in combinatorics. In the following we try to collect some of their properties which are used in this chapter and throughout the thesis. The partial exponential Bell polynomials are usually written as

$$B_{n,j}(z_1,\dots,z_{n-j+1}) = \sum \frac{n!}{(1!)^{r_1} r_1! (2!)^{r_2} r_2! \dots} z_1^{r_1} \cdots z_{n-j+1}^{r_{n-j+1}} \quad n \in \mathbb{N}, \quad j \leq n \tag{2.1}$$

where the summation is over all the solutions in non-negative integers $r_1,\dots,r_{n-j+1}$ of $r_1+2r_2+\cdots+(n-j+1)r_{n-j+1}=n$ and $r_1+r_2+\cdots+r_{n-j+1}=j$. A lighter expression is obtained using partitions of the integer n with length j. Recall that a partition of an integer n is a sequence $\pi=(\pi_1,\pi_2,\dots)$ of weakly decreasing positive integers, named parts of π, such that $\pi_1+\pi_2+\cdots=n$. A different notation is $\pi=(1^{r_1},2^{r_2},\dots)$, where $r_1,r_2,\dots$, named multiplicities of π, are the number of parts of π equal to 1,2,... respectively. The length of the partition is $l(\pi)=r_1+r_2+\cdots$ and the vector of multiplicities is $m(\pi)=(r_1,r_2,\dots)$. We write $\pi\vdash n$ to denote that π is a partition of n. Thus the partial exponential Bell polynomials (2.1) can be rewritten as (see, e.g., Di Nardo et al., 2008)

$$B_{n,j}(z_1,\dots,z_{n-j+1})=\sum_{\pi\vdash n,l(\pi)=j} d_\pi z_\pi \tag{2.2}$$

where the sum is over all the partitions $\pi\vdash n$ with length $l(\pi)=j$ and

$$z_\pi=z_1^{r_1}z_2^{r_2}\cdots \qquad d_\pi=\frac{i!}{(1!)^{r_1}r_1!(2!)^{r_2}r_2!\cdots}. \tag{2.3}$$

Using integer partitions, the explicit expression of the partial exponential polynomials can be recovered in `R` using the `kStatistics` package (Di Nardo and Guarino, 2022). A useful property used in the following is

$$B_{n,j}(ab\,z_1,\dots,ab^{n-j+1}\,z_{n-j+1})=a^j\,b^n\,B_{n,j}(z_1,\dots,z_{n-j+1}) \tag{2.4}$$

with a,b constants. Equation (2.4) follows from (2.2) since from (2.3) we have

$$(ab\,z_1)^{r_1}(ab^2\,z_2)^{r_2}\cdots=a^{r_1+r_2+\cdots}b^{r_1+2r_2+\cdots}z_\pi=a^j\,b^n\,z_\pi$$

taking into account that $l(\pi)=r_1+r_2+\cdots=j$ and $r_1+2r_2+\cdots=n$.

The n-th complete exponential Bell polynomial in the indeterminates $z_1,\dots,z_n$ is defined as

$$B_n(z_1,\dots,z_n)=\sum_{j=0}^{n}B_{n,j}(z_1,\dots,z_{n-j+1}) \tag{2.5}$$

with $\{B_{n,j}\}$ the partial exponential Bell polynomials as in (2.1). Note that n is the positive integer corresponding to the maximum degree of the monomials in (2.5). This polynomial sequence satisfies the following recurrence

$$B_{n+1}(z_1,\dots,z_{n+1})=\sum_{j=0}^{n}\binom{n}{j}z_{j+1}B_{n-j}(z_1,\dots,z_{n-j}) \tag{2.6}$$

with the initial value $B_0 = 1$. The generating function of $\{B_n\}$ is the formal power series composition

$$\exp[h_z(t) - z_0] = \sum_{n \geq 0} \frac{t^n}{n!} B_n(z_1, \ldots, z_n) \in \mathbb{R}[[t]] \tag{2.7}$$

where $\mathbb{R}[[t]]$ is the ring of formal power series in t and $h_z(t)$ is the generating function of $\{z_k\}_{k \geq 0}$, that is

$$h_z(t) = \sum_{k \geq 0} \frac{t^k}{k!} z_k.$$

A different expression of the n-th complete exponential Bell polynomial involves integer partitions as follows

$$B_n(z_1, \ldots, z_n) = \sum_{\pi \vdash n} d_\pi z_\pi \tag{2.8}$$

where the sum is over all the partitions $\pi \vdash n, d_\pi$ and z_π are given in (2.3). In particular we have

$$B_n(\lambda z_1, \ldots, \lambda z_n) = \sum_{\pi \vdash n} \lambda^{l(\pi)} d_\pi z_\pi \tag{2.9}$$

with λ a constant. Now, suppose to replace $\lambda^{l(\pi)}$ in (2.9) with a numerical sequence $\{a_{l(\pi)}\}$. Thanks to this device, the complete exponential Bell polynomials result as a special case of a wider class of polynomial families, the generalized partition polynomials

$$G_n(a_1, \ldots, a_n; z_1, \ldots, z_n) = \sum_{\pi \vdash n} d_\pi a_{l(\pi)} z_\pi \tag{2.10}$$

where the sum is again over all the partitions $\pi \vdash n$. A different expression of (2.10) involves the partial exponential Bell polynomials $\{B_{n,j}\}$ in (2.1)

$$G_n(a_1, \ldots, a_n; z_1, \ldots, z_n) = \sum_{j=1}^{n} a_j B_{n,j}(z_1, \ldots, z_{n-j+1}). \tag{2.11}$$

An example of a well known and useful polynomial family arising from (2.11) is the logarithmic, defined as

$$\mathcal{L}_n(z_1, \ldots, z_n) = \sum_{j=0}^{n} (-1)^{j-1} (j-1)! B_{n,j}(z_1, \ldots, z_{n-j+1}). \tag{2.12}$$

Apart from their combinatorial properties, the Bell polynomials and the logarithmic polynomials have a well known use in probability, which is connecting moments and cumulants of a random variable. Such relationships are summarised in Appendix 2.A.

2.2 Notation and Preliminaries on Discrete Random Variables

Before proceeding, let us very briefly clarify the notation used in the first part of this chapter, and, afterwards, recall some useful definitions and results. Let $(\Omega,\mathcal{F},\mathbb{P})$ be an underlying probability space. As usual, with discrete random variable (rv) we will refer to a measurable function

$$X:\Omega\to T,$$

where T is countably infinite or finite. We will usually consider $T\subset\mathbb{Z}$. The law or distribution of X is fully characterised by its probability mass function (pmf) p_X defined as

$$p_X(x)=\mathbb{P}(X=x),\quad x\in T.$$

Actually, we shall mostly equivalently refer to the sequence of probabilities

$$\{p_x\}_{x\in T},\quad x\in T,$$

where $p_x=p_X(x)$ for $x\in T$.

2.2.1 Probability Generating Function

The probability generating function (pgf) provides a compact representation of a discrete probability distribution and facilitates the computation of moments and the study of other properties of the random variable. The pgf G_X of X is defined as:

$$G_X(t)=E[t^X]=\sum_{x\in T}p_x t^x,$$

where t is a real number for which the series converges. Indeed, as we continue, if not specified otherwise t is assumed to be in a such a way. Note that G_X always exists for $-1\le t\le 1$ with $G_X(0)=p_0$ and $G_X(1)=1$. We report some of its properties.

1. By the definition of G_X, for every $x\in T$ we have that

$$p_x=\frac{G_X^{(x)}(0)}{x!},$$

   where the superscript (x) denotes differentiation x times.

2.

$$E[X]=G_X^{(1)}(1).$$

3. If Y is a discrete random variable independent from X with pgf G_Y, then the PGF of $Z = X + Y$ is the product of the individual pgfs

$$G_Z(t) = G_X \circ G_Y,$$

where $\circ$ denotes the usual composition of functions.

2.2.2 Log-Convexity and Log-Concavity

A discrete distribution $\{p_x\}_{x\in T}$ is said to be log-concave if

$$p_x^2 \geq p_{x-1}p_{x+1}, \quad x \in T. \tag{2.13}$$

Log-convex is defined with the inequality sign reversed, that is it must hold that

$$p_x^2 \leq p_{x-1}p_{x+1}, \quad x \in T. \tag{2.14}$$

Log-convexity and log-concavity are useful in deriving other properties of a distribution. An example is given in the following.

2.2.3 Unimodality and Strong Unimodality

A discrete distribution $\{p_x\}_{x\in T}$ is said to be unimodal if there exists an $\bar{x} \in T$ such that $p_x \geq p_{x-1}$ for all $x \leq \bar{x}$ and $p_x \leq p_{x-1}$ for all $x \geq \bar{x}$. A discrete distribution $\{p_x\}_{x\in T}$ is strongly unimodal if the convolution of $\{p_x\}_{x\in T}$ with any unimodal distribution is unimodal. Strong unimodality has the following characterization (Keilson and Gerber, 1971, Theorem 2).

Theorem 2.2.1. *Let a discrete rv X with support $T \subset \mathbb{Z}$. A necessary and sufficient condition for strong unimodality of X is log-concavity of $\{p_x\}_{x\in T}$.*

Given the latter, the next proposition is straightforward.

Proposition 2.2.2. *A strongly unimodal discrete rv X with support $T \subset \mathbb{Z}$ is unimodal.*

Proof. Condition (2.14) is equivalent to

$$\frac{p_{x+1}}{p_x} \leq \frac{p_x}{p_{x-1}}, \quad x \in T. \tag{2.15}$$

From (2.15) it follows that $\inf\{x^* \in T : \frac{p_{x^*+1}}{p_{x^*}} \leq 1\}$ is the mode. □

2.2.4 Parameter Estimation and Model Selection

Consider the measurable space $(\Omega, \mathcal{F})$. A parametric family $\mathcal{P}$ is a set of probability measures P_θ on $(\Omega, \mathcal{F})$ defined as

$$\mathcal{P} = \{P_\theta : \theta \in \Theta\},$$

where $\Theta \subseteq \mathbb{R}^d$ for some fixed positive integer d. An usual assumption is that of identifiability, meaning that the map $\theta \mapsto P_\theta$ is injective. In the present context of discrete random variables, P_θ could be given in the form of a sequence of probabilities, which, with some abuse of notation, we can denote by $p^\theta = \{p_x^\theta\}_{x \in T}$. Let $\mathbf{x} = (x_1, x_2, \ldots, x_n)$ be a random sample from a population with pmf $\{p_x\}_{x \in T}$. Assuming there exists $\theta_0 \in \Theta$ such that $p^{\theta_0} = \{p_x\}_{x \in T}$, parameter estimation is a crucial step in statistical inference to find an appropriate θ such that the probability measure P_θ of a parametric family fits well to the given sample $\mathbf{x}$. One of the most known approaches to this is the classical maximum likelihood estimation which is the main parameter estimation procedure employed in the first two chapters of this thesis.

2.2.5 Maximum Likelihood Estimation

For a fixed sample $\mathbf{x} = (x_1, x_2, \ldots, x_n)$, the likelihood function is defined by

$$L_n(\theta \,|\, \mathbf{x}) = \prod_{i=1}^{n} p_{x_i}^\theta. \tag{2.16}$$

The maximum likelihood estimator of the parameter θ based on a sample $\mathbf{x}$ (see, for instance, Casella and Berger, 2002, Chapter 7) is

$$\hat{\theta}(\mathbf{x}) = \underset{\theta \in \Theta}{\operatorname{argmax}}\, L(\theta \,|\, \mathbf{x}).$$

In practice, is it often simpler to work with the natural logarithm of the likelihood function, the log-likelihood

$$l_n(\theta \,|\, \mathbf{x}) = \sum_{i=1}^{n} \log(p_{x_i}^\theta). \tag{2.17}$$

Usually, the log-likelihood function is differentiable and first order optimality conditions can be employed to find the so called maximum likelihood equations, which, when solved, provide the desired estimator. However, it is often the case that for certain parametric families the obtained system of

equations may not have a closed-form solution. Numerical methods such as the Newton-Raphson method or the EM algorithm could be implemented to seek numerical solutions. These methods usually return a local maximum that depends on the choice of the starting guess. A stochastic optimisation algorithm that returns the global maxima is preferred when the equations have no analytic solution. For instance, the simulated annealing type of random search algorithm will be used in this thesis to obtain the ML estimates for various models (see, for instance, Robert and Casella, 2004, Chapter 5).

As we noted in the introduction, the family of discrete distribution we will discuss is a generalization of other well known probability distributions. The act of generalising a parametric probability distribution introduces more parameters into the model and the full parameter form of a parametric family will perform, in terms of maximised log-likelihood values, at least as well as any other special cases having a smaller number of parameters. However, the inclusion of additional parameters in the model is not always statistically justified (Wilks, 1938). To avoid over-fitting, several information-theoretic-based criteria have been developed for model selection. These include the Akaike information criterion (AIC), Bayesian information criterion (BIC) that make use of different functional forms of the maximum log-likelihood statistic to measure the divergences between the proposed models and the real model (see, e.g., Van Der Hoeven, 2005). Other hypothesis testing approaches such as the likelihood ratio test, Wald's test, and the score test are also available (see, for example, Cox and Hinkley, 1974, Section 9.3). In this thesis, the likelihood ratio test will be used to select the appropriate model from a nested family of distributions.

2.2.6 Likelihood Ratio Test

Consider a simpler nested model with reduced parameter space Θ_0 such that $\Theta_0 \subset \mathbb{R}^{d_0} \subset \Theta$, with d_0 a positive integer. The likelihood ratio test statistic for testing

$$H_0 : \theta \in \Theta_0 \text{ versus } H_1 : \theta \in \Theta_0^c$$

is

$$D_n = -2\log \frac{\sup_{\Theta_0} L_n(\theta \,|\, \mathbf{x})}{\sup_{\Theta} L_n(\theta \,|\, \mathbf{x})}.$$

Let $\hat{\theta}_0(\mathbf{x})$ be the ML estimator for θ in Θ_0, the test statistic D_n can be written as

$$D_n = -2(\ell_n(\hat{\theta}_0(\mathbf{x}) \,|\, \mathbf{x}) - \ell_n(\hat{\theta}(\mathbf{x}) \,|\, \mathbf{x})).$$

A well known theorem of Wilks (1938) states that as the sample size n goes to infinity

$$D_n \xrightarrow{d} \chi^2_{d-d_0},$$

under H_0. That is, D_n is approximately χ^2 distributed, with degrees of freedom equal to the difference between the number of free parameters of the alternative and the null models.

2.3 The Hurwitz-Lerch Zeta Distribution

We shall employ the standard notation used for instance in Gupta et al. (2008). Let $\mathcal{T}(z,r,u) = z\Phi(z,r+1,u+1)$, where

$$\Phi(z,r,u) = \sum_{n=0}^{\infty} \frac{z^n}{(n+u)^r} \tag{2.18}$$

is the Lerch Transcendent function, defined for $u \notin \{0,-1,-2,-3,\dots\}$ and r in the set of complex number $\mathbb{C}$ when $|z|<1$, or r in the half plane $Re(r)>1$ when $|z|=1$ (Bateman and Erdélyi, 1953).

Definition 2.3.1. *A discrete random variable* $Y \sim \text{HLZD}(a,\theta,s)$ *if*

$$p_y = \frac{\theta^y}{\mathcal{T}(\theta,s,a)(y+a)^{s+1}}, \quad y \in \mathbb{N}^+, \tag{2.19}$$

where $a > -1$ *and* $s \in \mathbb{R}$ *if* $\theta \in (0,1)$ *or* $s > 0$ *if* $\theta = 1$.

Expression (2.19) comprises a wide range of well known discrete distributions, some of which are reported in Table 2.1 below.

Table 2.1: Lerch family of probability distributions with the corresponding parameter domains.

ID	Probability distribution	θ	s	a
1	3-par Lerch	(0,1)	$\mathbb{R}$	$(-1,\infty)$
2	2-par polylogarithmic	(0,1)	$\mathbb{R}$	0
3	1-par logarithmic	(0,1)	0	0
4	1-par geometric	(0,1)	-1	1
5	2-par extended log	(0,1)	0	$(-1,\infty)$

The pgf of $Y \sim \text{HLZD}(a,\theta,s)$ is

$$G_Y(z) = \frac{\theta\Phi(z\theta,s+1,a+1)}{\Phi(\theta,s,a)}, \quad 0 < z\theta \leq 1, \tag{2.20}$$

with $G_Y(0) = 0$.

2.3.1 Moments and Cumulants

HLZD moments have a closed form expression involving the Lerch Transcendent function. Differently from Aksenov and Savageau (2005), we find this closed form expression using (2.19).

Proposition 2.3.2. *If* $Y \sim \text{HLZD}(a,\theta,s)$, *then*

$$\xi_k := \mathbb{E}[Y^k] = \sum_{j=0}^{k} (-a)^{k-j} \binom{k}{j} \frac{\Phi(\theta, s+1-j, a)}{\Phi(\theta, s+1, a+1)}, \quad k \in \mathbb{N}^+. \tag{2.21}$$

Proof. Fix $k \in \mathbb{N}^+$. Using the binomial expansion of $y^k = (y-a+a)^k$, we have

$$\begin{aligned}
\xi_k &= \sum_{y=1}^{\infty} y^k \mathbb{P}(Y=y) = \sum_{y=1}^{\infty} y^k \frac{\theta^{y-1}}{\Phi(\theta, s+1, a+1)(y+a)^{s+1}} \\
&= \sum_{y=1}^{\infty} \left(\sum_{j=0}^{k} \binom{k}{j} (y+a)^j (-a)^{k-j} \right) \frac{\theta^{y-1}}{\Phi(\theta, s+1, a+1)(y+a)^{s+1}} \\
&= \sum_{j=0}^{k} \binom{k}{j} (-a)^{k-j} \frac{1}{\Phi(\theta, s+1, a+1)} \sum_{y=1}^{\infty} \frac{\theta^{y-1}}{(y+a)^{s+1-j}}
\end{aligned}$$

from which (2.21) follows by taking into account (2.18). □

As a corollary, the mean and the variance are respectively:

$$\mathbb{E}[Y] = \frac{\mathcal{T}(\theta, s-1, a)}{\mathcal{T}(\theta, s, a)} - a \quad \text{Var}[Y] = \frac{\mathcal{T}(\theta, s-2, a)}{\mathcal{T}(\theta, s, a)} - \left(\frac{\mathcal{T}(\theta, s-1, a)}{\mathcal{T}(\theta, s, a)} \right)^2.$$

Remark 2.3.3. Suppose $\theta = 1$. Recall that we can have convergence of the series $\Phi(1, s, a)$ only if $s > 0$. By inspecting the r.h.s. of (2.21), we have that ξ_k will be infinite for any $k \geq k'$, where k' is the first integer such that $s + 1 - k' \leq 0$. □

More generally, for $k \in \mathbb{N}^+$ the k-th central moment can be recovered as

$$\xi'_k := \mathbb{E}[(Y - \xi_1)^k] = \sum_{j=0}^{k} \binom{k}{j} \xi_1^j \xi_{k-j}$$

and the factorial moments as

$$(\xi)_k = \mathbb{E}[Y(Y-1)\cdots(Y-k+1)] = \sum_{j=0}^{k} s(k,j) \xi_k \tag{2.22}$$

with $s(k,j)$ the Stirling numbers of the first kind (Charalambides, 2002). HLZD cumulants are such that

$$\kappa_n(Y) = \mathcal{L}_n(\xi_1,\ldots,\xi_n), \quad n \in \mathbb{N}^+,$$

where $\{\xi_j\}$ are the moments of $Y \sim \text{HLZD}(a,\theta,s)$, given in (2.21), and $\mathcal{L}_n$ is the n-th logarithmic polynomial (2.12). Let us recall that, if the moment generating function (mgf) $M_Y(t)$ of Y is well defined in a suitable neighborhood of 0, then the coefficients $\{\kappa_n(Y)\}_{n\geq 1}$ in the expansion

$$M_Y(t) = \exp\left(\sum_{n\geq 1} \frac{t^n}{n!}\kappa_n(Y)\right)$$

are the cumulants of Y. The first cumulant is the mean $\mathbb{E}[Y]$, the second cumulant is the variance $\text{Var}(Y)$, the skewness and the kurtosis of Y can be recovered using the third and the fourth cumulant of Y respectively. More details are given in Appendix 2.A.

2.3.2 Mode

The HLZ distribution is a particular case of a wider class of distributions called the Modified Power Series Distributions (MPSD) introduced by Gupta (1974).

Definition 2.3.4. A discrete random variable $Y \sim \text{MPSD}(a,g,f)$ if

$$p_y = \frac{a(y)g(\theta)^y}{f(\theta)}, \quad y \in T \subset \mathbb{N} \tag{2.23}$$

where $a(y)$, $g(\theta)$ and $f(\theta)$ are positive, bounded, and differentiable functions of y and θ respectively with $f(\theta) = \sum_{y\in T} a(y)g(\theta)^y$. □

Using this wider class of distributions, we will prove that $Y \sim \text{HLZD}(a,\theta,s)$ is unimodal for all $s \in \mathbb{R}$. The following result follows straightforwardly from Theorem 2.2.1 by substituting (2.23) into (2.14).

Proposition 2.3.5. *$Y \sim \text{MPSD}(a,g,f)$ is strongly unimodal if and only if $\{a(y)\}_{y\in T}$ is a logarithmically concave sequence.*

Given these facts, we can state the next result.

Proposition 2.3.6. *Suppose $Y \sim \text{HLZD}(a,\theta,s)$.*

(i) If $s \geq -1$, the sequence $\{q_y\}_{y\geq 1}$ is monotonically decreasing and the mode is $y = 1$.

(ii) If $s < -1$, Y is strongly unimodal and therefore also unimodal.

Proof. Similarly to what stated in Section 2.3 of Gupta et al. (2008), we have

$$\frac{p_y}{p_{y-1}} = \theta\left(1 - \frac{1}{a+y}\right)^{s+1}, \quad y \in \{2,3,\dots\}. \tag{2.24}$$

Since $\theta \in (0,1), a > -1$ and $s \geq -1$, the rhs of (2.24) is always between 0 and 1, thus *(i)* follows. From Theorem 2.2.1 and Proposition 2.2.2, for *(ii)* we have to prove that $\{p_y\}_{y\geq 1}$ is log-concave, that is it satisfies (2.14). Using Proposition 2.3.5, since $a(y) = \frac{1}{(y+a)^{s+1}}$, we simply need to show that $(1-(y+a)^{-2})^{s+1} \geq 1$ for all $y \in \mathbb{N}$. The latter holds if $s < -1$. □

2.3.3 Failure Rate

From Definition 2.3.1 we have that $Y \sim \text{HLZD}(a,\theta,s)$ has cumulative distribution function (cdf)

$$\mathbb{P}(Y \leq k) = 1 - \theta^k \frac{\mathcal{T}(\theta,s,a+k)}{\mathcal{T}(\theta,s,a)}, \quad k \in \mathbb{N}^+. \tag{2.25}$$

Following P. L. Gupta and Tripathi (1997), we define the failure rate r of a discrete rv X supported on $\mathbb{N}^+$ as

$$r_X(k) = \frac{\mathbb{P}(X = k)}{\mathbb{P}(X \geq k)}, \quad k \in \mathbb{N}^+. \tag{2.26}$$

From (ii) in Theorem 1 of Gupta et al. (2008) we have that if X is log-convex then r_X is a non-increasing function of k. Hence we can state the following simple result.

Proposition 2.3.7. *Let $Y \sim \text{HLZD}(a,\theta,s)$. Then it holds that*

$$r_Y(k) = \frac{\theta}{(a+k)^{s+1}\mathcal{T}(\theta,s,a+k-1)}, \quad k \in \mathbb{N}^+. \tag{2.27}$$

Moreover, if $s \geq -1$, r_Y is a non-increasing function of k.

Proof. Expression (2.27) immediately follows from (2.25) and (2.19). According to the discussion slightly above, we have to prove that $p_y^2 \leq p_{y-1}p_{y+1}$,

$y \in \mathbb{N}^+$ or, equivalently, that $\frac{p_y}{p_{y-1}} - \frac{p_{y+1}}{p_y} \leq 0$, $y \in \mathbb{N}^+$. As in the proof of Theorem 1 in Gupta et al. (2008) we have

$$\frac{p_y}{p_{y-1}} - \frac{p_{y+1}}{p_y} = \theta \left[\left(\frac{a+y-1}{a+y} \right)^{s+1} - \left(\frac{a+y}{a+y+1} \right)^{s+1} \right] = \theta \left[\frac{(b^2-1)^{s+1} - (b^2)^{s+1}}{[b(b+1)]^{s+1}} \right],$$

where $b = a+k$. The above centered display is less or equal then zero when $s \geq -1$. □

2.3.4 Convolution

The family of HLZ distributions is not closed under convolution. Nevertheless, as a subclass of MPS distributions, the HLZD convolution still returns a MPSD. Indeed, more in general we will prove that the family of MPS distributions is closed under convolution.

Theorem 2.3.8. *If $Y_1, \ldots, Y_j$ are independent rvs identically distributed to $Y \sim \text{HLZD}(a, \theta, s)$, then $Y_1 + \cdots + Y_j \sim \text{MPSD}(a_j, g, f_j)$ with $f_j(\theta) = [\mathcal{T}(\theta, s, a)]^j$,*

$$a_j(y) = \frac{j!}{y!} \sum_{\pi \vdash y, l(\pi) = j} d_\pi (a_\pi)^{s+1} \quad \text{with} \quad a_\pi = (a+1)^{-r_1}(a+2)^{-r_2} \cdots, \; y \in T_j, \tag{2.28}$$

d_π as given in (2.3) and $T_j = \{y_1 + \cdots + y_j \in \mathbb{N} : y_1, \ldots, y_j \in T\}$.

Proof. Observe that if $Y_1, \ldots, Y_j$ are rvs i.i.d. to $Y \sim \text{MPSD}(a, g, f)$, then $Y_1 + \cdots + Y_j \sim \text{MPSD}(a_j, g, f_j)$ with $f_j(\theta) = f^j(\theta)$ and

$$a_j(y) = \begin{cases} \frac{j!}{y!} B_{y,j}[a(1), \ldots, a(y-j+1)], & y \in T_j \\ 0, & y \notin T_j. \end{cases} \tag{2.29}$$

Indeed in (2.23), set $a(y) = 0$ if $y \notin T$ and consider the sequence $\{p_y\}_{y \geq 1}$ such that $p_y = 0$ if $y \notin T$. By using Lemma 1 in Eger (2016), we have

$$\mathbb{P}(Y_1 + \cdots + Y_j = y) = \frac{j!}{y!} B_{y,j}(p_1, \ldots, p_{y-j+1}) \tag{2.30}$$

where $\{B_{y,j}\}$ are the partial exponential Bell polynomials (2.1). From (2.30) with p_i replaced by $a(i)g(\theta)^i / f(\theta)$ for $i = 1, \ldots, y-j+1$ and using (2.4) we have

$$\mathbb{P}(Y_1 + \cdots + Y_j = y) = \frac{j!}{y!} \frac{g(\theta)^y}{f^j(\theta)} B_{y,j}[a(1), \ldots, a(y-j+1)]. \tag{2.31}$$

Thus $Y_1+\cdots+Y_j \sim \text{MPSD}(a_j, g, f_j)$ with $f_j(\theta)=f^j(\theta)$ and $a_j(y)$ given in (2.29). From (2.31) note that $a_j(y)=0$ if $1,2,\ldots,y-j+1 \notin T$. By replacing $g(\theta)=\theta, f(\theta)=\mathcal{T}(\theta,s,a)$ and $a(k)=(k+a)^{-(s+1)}$ for $k=1,\ldots,y-j+1$ in (2.31) we have

$$a_j(y)=\frac{j!}{y!}B_{y,j}\left[(a+1)^{-(s+1)},\ldots,(a+y-j+1)^{-(s+1)}\right], \quad y\in\mathbb{N}^+.$$

The result follows after some manipulations, rewriting the partial Bell exponential polynomials as in (2.2). □

2.3.5 Tail Behaviour

Following Section 4 of Liew and Ong (2012), we report a brief description of the tail behaviour of $Y \sim \text{HLZD}(a,\theta,s)$.

Proposition 2.3.9. *Let $Y \sim \text{HLZD}(a,\theta,s)$. Then Y has a long tail if and only $\theta=1$.*

Proof. From (2.19) it can be readily computed that

$$\lim_{y\to\infty}\frac{p_{y+1}}{p_y}=\theta,$$

after which the statement is immediate. □

The latter result indicates that θ is the parameter which governs the thickness of the tail.

Remark 2.3.10. This fact can also be deduced by the following argument. We have that the mgf of $Y \sim \text{HLZD}(a,\theta,s)$ is

$$G_Y(e^t)=\frac{\theta\Phi(e^t\theta,s+1,a+1)}{\Phi(\theta,s,a)}, \tag{2.32}$$

with G_Y the pgf of Y. Recall that in general the Lerch transcendent function $\Phi(z,r,u)$ converges for $u>0$ if z and r are any complex numbers with $|z|>1$ or $|z|=1$ and $Re(r)>0$ (Bateman and Erdélyi, 1953). Under the parametric assumptions in (2.19), when $\theta\in(0,1)$ there always exists a neighbourhood of the origin where $G_Y(e^t)$ exists finite. Hence, Y does not have a heavy tail and therefore neither a long tail. Since we must have $\theta\leq 1$ for the convergence of the series $\Phi(\theta,s,a)$, if $\theta=1$ we have that $G_Y(e^t)$ will be infinite for any $t>0$ from the fact that $e^t\theta>1$ for any $t>0$. To see this from yet another perspective, recall Remark 2.3.3. □

2.3.6 Maximum Likelihood Estimation

Consider a vector $\mathbf{y} = (y_1, \ldots, y_n)$ of independent observations of $Y \sim \text{HLZD}(a, \theta, s)$. The MLE of (θ, s, a) is

$$(\hat{\theta}, \hat{s}, \hat{a}) = \underset{(\theta,s,a)\in\Theta}{\operatorname{argmax}}\, \ell_n(\theta, s, a \,|\, \mathbf{y}), \tag{2.33}$$

with $\Theta = (0,1) \times (-\infty, +\infty) \times (-1, \infty)$, $\ell_n(\theta, s, a \,|\, \mathbf{y}) = \log L_n(\theta, s, a \,|\, \mathbf{y})$ the log-likelihood function and

$$L_n(\theta, s, a \,|\, \mathbf{y}) = \prod_{i=1}^{n} \mathbb{P}(Y = y_i).$$

The MLE of the HLZD parameters (θ, s, a) has been studied by Gupta in Gupta et al. (2008), where the Author derives the following formula

$$\ell_n(\theta, s, a \,|\, \mathbf{y}) = \log\theta \sum_{i=1}^{n} y_i - (s+1) \sum_{i=1}^{n} \log(a + y_i) - n \log[\Phi(\theta, s+1, a)]. \tag{2.34}$$

Additionally, he showed that the three likelihood equations arising from maximizing the log-likelihood correspond to the following equations of the generalised method of moments. In particular we have

$$\sum_{i=1}^{n} \frac{y_i}{n} = \mathbb{E}[Y], \tag{2.35}$$

$$\sum_{i=1}^{n} \frac{\log(a + y_i)}{n} = \mathbb{E}[\log(a + Y)], \tag{2.36}$$

$$\sum_{i=1}^{n} \frac{1}{n(a + y_i)} = \mathbb{E}\left[\frac{1}{a + Y}\right]. \tag{2.37}$$

For the case of the Log distribution (ID = 1 in Table 2.1), the single parameter θ can be estimated by solving (2.35) for $s = 0$ and $a = 0$, for the case of the Polylog distribution (ID = 2 in Table 2.1) the two parameters θ and s can be obtained by solving the system of equations (2.35) and (2.36) with $a = 0$ and for the case of the 2-parameters Log distribution (ID = 5 in Table 2.1) the two parameters θ and a can be obtained by solving the system of equations (2.35) and (2.37) with $s = 0$.

Closed form solutions of the above equations are not available and the moments $\mathbb{E}[\log(a + Y)]$ and $\mathbb{E}[1/(a + Y)]$ must also be numerically approximated. As noted in Section 2.2.5, the likelihood equations may be solved

by standard numerical methods to obtain the MLE. However, it is well known that this does not guarantee that global maxima of the likelihood have been achieved. In order to avoid this problem, a global optimisation method can be employed to solve (2.33). The global optimisation method takes advantage of the bounds of the parameters. For instance, the MLE of the parameters can be obtained through a global optimisation algorithm known as *Simulated Annealing*. Simulated annealing is a stochastic global optimisation technique applicable to a wide range of discrete and continuous variable problems. It makes use of Markov Chain Monte Carlo samplers, to provide a means to escape local optima by allowing moves which worsen the objective function, with the aim of finding a global optimum. Technical details can be found in Bélisle (1992), a variant of which is the algorithm implemented in the `Optim` function in the base `Stats` `R`-package. The latter has been used in this thesis.

2.4 The Poisson-Stopped Hurwitz-Lerch Zeta Distribution

Definition 2.4.1. Let $\lambda>0$, $\theta\in(0,1)$, $a>-1$ and $s\in\mathbb{R}$. A discrete random variable $X\sim\text{PSHLZD}(\lambda,a,\theta,s)$ if its pgf is

$$G_X(t)=\exp\left(\lambda\left[\frac{\theta\Phi(t\theta,s+1,a+1)}{\Phi(\theta,s,a)}-1\right]\right),\quad 0<\theta t\leq 1, \tag{2.38}$$

where Φ is the Lerch Transcendent as in (2.20). □

According to Definition 2.4.1, $X\sim\text{PSHLZD}(\lambda,a,\theta,s)$ takes non-negative integer values and belongs to the class of generalized rvs as defined in Charalambides (1977). Indeed given two independent rvs Z and Y, with pgf G_Z and G_Y respectively, the generalized rv X has pgf

$$G_X=G_Z\circ G_Y. \tag{2.39}$$

The composition (2.38) matches (2.42) when $Y\sim\text{HLZD}(a,\theta,s)$ and Z is a Poisson (PS) rv of parameter $\lambda>0$, independent of Y, since $G_Z(t)=\exp[\lambda(t-1)]$. As we continue, to avoid confusion, we will refer to the sequence of probabilities associated to $Y\sim\text{HLZD}(a,\theta,s)$ as $\{q_x\}_{x\in\mathbb{N}^+}$, unless specified otherwise.

In the following we analyse in detail the properties of the PSHLZD using the complete exponential Bell polynomials. Some of the properties given in Liew et al. (2020) will also be briefly recalled.

Proposition 2.4.2. *If* $X \sim \text{PSHLZD}(\lambda, a, \theta, s)$ *then*

$$p_x = \begin{cases} e^{-\lambda}, & x = 0, \\ \frac{e^{-\lambda}}{x!} B_x(\lambda q_1, \ldots, \lambda x! q_x), & x \in \mathbb{N}^+, \end{cases} \tag{2.40}$$

where B_x *is the complete exponential Bell polynomial* (2.5) *of degree* x.

Proof. Observe that $G_Y(0) = 0$ and $G_Y(t) = \sum_{x \geq 1} x! q_x t^x / x!$. The result follows from (2.7) with $z_0 = 0$ and $h_z(t) = G_Y(t)$, since from (2.38) we have

$$\exp(-\lambda) \exp[\lambda G_Y(t)] = \sum_{x \in \mathbb{N}} \frac{t^x}{x!} e^{-\lambda} B_x(\lambda 1! q_1, \ldots, \lambda x! q_x).$$

□

Corollary 2.4.3. *If* $X \sim \text{PSHLZD}(\lambda, a, \theta, s)$ *then* $p_0 = P(X = 0) = e^{-\lambda}$ *and*

$$p_x = \theta^x e^{-\lambda} \sum_{\pi \vdash x} \left(\frac{\lambda}{T(\theta, s, a)} \right)^{l(\pi)} \frac{(a_\pi)^{s+1}}{m_\pi!} \quad x = 1, 2, \ldots. \tag{2.41}$$

where the sum is over all the partitions $\pi \vdash x$, $m_\pi! = r_1! r_2! \cdots$ *and* a_π *is given in* (2.28).

Proof. From (2.40), using (2.8) and (2.3), we have

$$\begin{aligned} p_x &= \frac{e^{-\lambda}}{x!} \sum_{\pi \vdash x} \frac{x!}{(1!)^{r_1} r_1! (2!)^{r_2} r_2! \cdots} (\lambda 1! q_1)^{r_1} (\lambda 2! q_2)^{r_2} \\ &= e^{-\lambda} \sum_{\pi \vdash x} \frac{\lambda^{r_1 + r_2 + \cdots}}{r_1! r_2! \cdots} \frac{[(a+1)^{-(s+1)}]^{r_1} [(a+2)^{-(s+1)}]^{r_2} \cdots \theta^{r_1 + 2r_2 + \cdots}}{T(\theta, s, a)^{r_1 + r_2 + \cdots}} \end{aligned}$$

by which (2.41) follows observing that $r_1 + r_2 + \cdots = l(\pi)$ and $r_1 + 2r_2 + \cdots = x$. □

As a corollary of Proposition 2.4.2 and recursion (2.6), the sequence $\{p_x\}_{x \in \mathbb{N}}$ in (2.40) satisfies the following equations.

Corollary 2.4.4. *If* $X \sim \text{PSHLZD}(\lambda, a, \theta, s)$ *then*

$$p_{x+1} = \frac{\lambda}{x+1} \sum_{j=0}^{x} (j+1) q_{j+1} p_{x-j}, \quad x \in \{1, 2, \ldots\}$$

and $p_0 = e^{-\lambda}$.

Proof. The result follows using (2.6) since we have

$$\begin{aligned} p_{x+1} &= \frac{e^{-\lambda}}{(x+1)!} B_{x+1}(\lambda q_1, \ldots, \lambda(x+1)! q_{x+1}) \\ &= \frac{e^{-\lambda}}{(x+1)!} \sum_{j=0}^{x} \binom{x}{j} B_{x-j}(\lambda q_1, \ldots, \lambda(x-j)! q_{x-j})(j+1)! q_{j+1} \\ &= \frac{e^{-\lambda}}{(x+1)!} \sum_{j=0}^{x} \frac{x!}{(x-j)!j!} p_{x-j} \frac{(x-j)!}{e^{-\lambda}} (j+1)! q_{j+1}. \end{aligned}$$

□

If $s \geq 0$ and $-1 < a \leq 0$, the PSHLZD is unimodal. It is a consequence of (Liew et al., 2020, Property 1), where the authors, with the above parameters, prove the discrete self-decomposability of a PSHLZD, which implies unimodality (Steutel and van Harn, 1979).

2.4.1 Log-Concavity

Under suitable conditions, the PSHLZD is log-concave.

Proposition 2.4.5. *If $X \sim \text{PSHLZD}(\lambda, \theta, a, s)$ and $s \geq -1$, then X has a log-concave cdf, that is*

$$[\mathbb{P}(X \leq x)]^2 \geq \mathbb{P}(X \leq x-1)\mathbb{P}(X \leq x+1), \quad x \in \mathbb{N}.$$

Proof. According to (Badía et al. (2021), Theorem 1), a random sum $\sum_{i=1}^{Z} Y_i$ of i.i.d. rvs has a log-concave cdf if Z is strongly unimodal and the common distribution of $\{Y_i\}_{i \geq 1}$ has a decreasing pdf. Thus, the result follows as $X \sim \sum_{i=1}^{Z} Y_i$ with $Z \sim \text{PS}(\lambda)$, which has a log-concave pmf (strongly unimodal), and $Y \sim \text{HLZD}(\theta, a, s)$ with a decreasing pmf when $s \geq -1$ (see Proposition 2.3.6). □

Proposition 2.4.5 gives a sufficient condition to get cdf log-concavity. A different way is to consider the sequence of probabilities $\{p_x\}$. Indeed, if X has a log-concave pmf (2.14), then its cdf is also log-concave (Badía et al., 2021). In the more general setting of generalised rvs, we can state the following. Recall that X is a generalised random variable if given two independent rvs Z and Y, with pgf G_Z and G_Y respectively, it has pgf

$$G_X = G_Z \circ G_Y. \tag{2.42}$$

Proposition 2.4.6. *Suppose X is a generalised rv. Then it has a log-concave pdf if and only if the sequence*

$$p_x = \frac{1}{x!}\mathcal{B}_x(1!\tilde{q}_1,\ldots,x!\tilde{q}_x;1!q_1,\ldots,x!q_x), \quad x\in\mathbb{N}^+, \tag{2.43}$$

with $p_0=\mathbb{P}(X=0)=G_Z[G_Y(0)]$, $\tilde{q}_x=\mathbb{P}(Z=x)$ and $q_x=\mathbb{P}(Y=x)$, is log-concave.

Proof. Expression (2.43) follows from (2.3) in Charalambides (1977) using the general partition polynomials (2.8). □

When $Z\sim\mathrm{PS}(\lambda)$ a necessary and sufficient condition to recover strong unimodality is related to the magnitude of q_1 and q_2, as the following theorem shows.

Theorem 2.4.7. *If X is a generalized rv with Y strongly unimodal and $Z\sim PS(\lambda)$, then X is strongly unimodal if and only if $\lambda q_1\geq 2q_2$.*

Note that a similar result is proved in (Yu (2009), Theorem 4). We provide a different proof using the following lemma.

Lemma 2.4.8. *If $\{z_j\}_{j\geq 1}\in[0,\infty)$ is a log-concave sequence, then the sequence $\{\frac{1}{n!}B_n(z_1,\ldots,z_n)\}_{n\geq 1}$ is log-concave if and only if $z_1\geq 2z_2$, with $\{B_n\}_{n\geq 0}$ given in (2.5).*

Proof. If $\{z_j\}_{j\geq 0}$ with $z_0=1$ is a log-concave sequence of non-negative real numbers and the sequence $\{a(n)\}_{n\geq 0}$ is defined by

$$\sum_{n=0}^{\infty}\frac{a(n)}{n!}y^n=\exp\left(\sum_{j=1}^{\infty}\frac{z_j}{j!}y^j\right) \tag{2.44}$$

then the sequence $\{\frac{a(n)}{n!}\}_{n\geq}$ is log-concave (Bender and Canfield, 1996). Equation (2.44) parallels (2.7). Therefore, the sequence $\{\frac{1}{n!}B_n(z_1,\ldots,z_n)\}_{n\geq 1}$ results as log-concave if the sequence $\{z_j\}_{j\geq 0}$ is log-concave. Note that for $j\geq 2$ we have

$$\frac{z_j^2}{[(j-1)!]^2}\geq\frac{z_{j-1}}{(j-2)!}\frac{z_{j+1}}{j!},$$

which easily reduces to $j\,z_j^2\geq(j-1)\,z_{j-1}\,z_{j+1}$ always satisfied when $\{z_j\}_{j\geq 1}$ is log-concave. Now let $j=1$. We have $\{z_j\}_{j\geq 0}$ is log-concave if and only if $z_1\geq 2z_2$ and the result follows. □

Proof of Theorem 2.4.7. Following the same arguments of Proposition 2.4.2, for a generalized rv with $Z \sim PS(\lambda)$, (2.43) reduces to

$$\mathbb{P}(X=x) = \frac{e^{-\lambda}}{x!} B_x(\lambda 1! q_1, \ldots, \lambda x! q_x)$$

with $q_x = \mathbb{P}(Y=x)$ for $x \in \mathbb{N}^+$. The sequence $\{\frac{e^{-\lambda}}{x!} B_x(\lambda 1! q_1, \ldots, \lambda x! q_x)\}_{x\geq 1}$ is log-concave if and only if the sequence $\{\frac{1}{x!} B_x(\lambda 1! q_1, \ldots, \lambda x! q_x)\}_{x\geq 1}$ is log-concave. The result follows using Lemma 2.4.8. □

Corollary 2.4.9. *If $s < -1$, $X \sim$ PSHLZD(λ, θ, a, s) is strongly unimodal if and only if $\lambda q_1 \geq 2q_2$.*

2.4.2 Moments and Cumulants

Exploiting the properties of Bell polynomials, we can recover closed form expressions of PSHLZD moments and cumulants in terms of moments of $Y \sim$ HLZD(a, θ, s).

Proposition 2.4.10. *If $X \sim$ PSHLZD(λ, a, θ, s) then*

$$\mu_k := \mathbb{E}[X^k] = B_k(\lambda \xi_1, \ldots, \lambda \xi_k), \quad k \in \mathbb{N}^+, \tag{2.45}$$

with B_k the k-th complete exponential Bell polynomial (2.5) and $\xi_1, \ldots, \xi_k$ the first k moments of $Y \sim$ HLZD(a, θ, s) given in (2.21).

Proof. If M_X and M_Y are the mgfs of $X \sim$ PSHLZD(λ, a, θ, s) and $Y \sim$ HLZD(a, θ, s) respectively, then

$$M_X(t) = G_X(e^t) = e^{\lambda[G_Y(e^t)-1]} = e^{\lambda[M_Y(t)-1]} \tag{2.46}$$

from (2.42). Equation (2.45) follows as the rhs of (2.46) can be written as (2.7), with $h_z(t) = \lambda M_Y(t)$ and $z_0 = \lambda$. □

Remark 2.4.11. Taking into account (2.46), if $X \sim$ PSHLZD(λ, a, θ, s) then $X \sim Y_1 + \cdots + Y_Z$ with $Y \sim$ HLZD(a, θ, s) and $Z \sim$ PS(λ), that is X is a compound Poisson rv. Therefore the PSHLZD is an infinitely divisible distribution (Sato, 2013). □

Moments (2.45) can be explicited written using (2.9). A straightforward corollary of recursion (2.6) is the following.

Corollary 2.4.12. $\mu_{k+1} = \lambda \sum_{j=0}^{k} \binom{k}{j} \mu_{k-j} \xi_{j+1}, \quad k \in \mathbb{N}^+$.

If $\mu'_k := \mathbb{E}[(X-\mu_1)^k]$ denotes the k-th central moment of $X \sim \text{PSHLZD}(\lambda,a,\theta,s)$ then

$$\mu'_k = \sum_{k=0}^{n} \binom{n}{k} (-\lambda\xi_1)^{n-k} B_k(\lambda\xi_1,\ldots,\lambda\xi_k), \quad k \in \mathbb{N}^+.$$

Proposition 2.4.13. *If $X \sim \text{PSHLZD}(\lambda,a,\theta,s)$ then*

$$(\mu)_k := \mathbb{E}[X(X-1)\cdots(X-k+1)] = B_k(\lambda(\xi)_1,\ldots,\lambda(\xi)_k), \quad k \in \mathbb{N}^+ \tag{2.47}$$

where $(\xi)_1,\ldots,(\xi)_k$ are the first k factorial moments of $Y \sim \text{HLZD}(a,\theta,s)$ given in (2.22).

Proof. Let us recall that, if $Q_X(t)$ is the factorial mgf of $\{(\mu)_k\}$, then $Q_X(t) = G_X(t+1)$ with G_X the pgf of X. Therefore we have

$$Q_X(t) = G_X(t+1) = \exp\big(\lambda[G_Y(t+1)-1]\big) = \exp\big(\lambda[Q_Y(t)-1]\big), \tag{2.48}$$

with $Q_Y(t)$ the generating function of the factorial moments $\{(\xi)_k\}$. Equation (2.47) follows as the rhs of (2.48) can be written as (2.7), with $z_0 = \lambda$ and $h_z(t) = \lambda Q_Y(t)$. □

Proposition 2.4.14. *If $\kappa_n(X)$ is the n-th cumulant of $X \sim \text{PSHLZD}(\lambda,a,\theta,s)$ then $\kappa_n(X) = \lambda\xi_n$, for $n = 1,2,\ldots$ where ξ_n is the n-th moment of $Y \sim \text{HLZD}(a,\theta,s)$ given in (2.21).*

Proof. The result follows since

$$\log M_Y(t) = \log[e^{\lambda(M_X(t)-1)}] = \lambda[M_x(t)-1] = \sum_{n\geq 1} \frac{t^n}{n!}\lambda\mathbb{E}[X^k].$$

□

2.4.3 Maximum Likelihood Estimation

Suppose to have a sample $\mathbf{x} = (x_1,\ldots,x_n)$ of independent observations of $X \sim \text{PSHLZD}(\lambda,a,\theta,s)$. The MLE of (λ,θ,s,a) is

$$(\hat\lambda,\hat\theta,\hat s,\hat a) = \underset{(\lambda,\theta,s,a)\in\Theta}{\text{argmax}}\ \ell_n(\lambda,\theta,s,a\,|\,\mathbf{x}),$$

with $\Theta = (0,\infty)\times(0,1)\times(-\infty,+\infty)\times(-1,\infty)$, $\ell_n(\lambda,\theta,s,a,\mathbf{x}) = \log L_n(\lambda,\theta,s,a\,|\,\mathbf{x})$ the log-likelihood function and

$$L_n(\lambda,\theta,s,a,\mathbf{x}) = \prod_{i=1}^{n} \mathbb{P}(X = x_i).$$

The MLE of the PSHLZD parameters in this case must be directly tackled with the global optimisation method described in Section 2.3.6, since $\ell_n(\lambda,\theta,s,a)$ is not analytically tractable referring to (2.41).

2.5 The One Inflated Hurwitz-Lerch Zeta Distribution

Definition 2.5.1. A discrete random variable $Z \sim \text{OIHLZD}(p,a,\theta,s)$ if

$$\begin{aligned} \mathbb{P}(Z=1) &= p+(1-p)\mathbb{P}(Y=1), \\ \mathbb{P}(Z=x) &= (1-p)\mathbb{P}(Y=x), \quad x \in \{2,3,\ldots\}, \end{aligned} \tag{2.49}$$

with $p \in [0,1]$ and $Y \sim \text{HLZD}(a,\theta,s)$. □

This definition parallels the definition of the Zero Inflated Modified Power Series Distribution given by Gupta in Gupta et al. (1995). If $G_Z(t)$ denotes the pgf of $Z \sim \text{OIHLZD}(p,a,\theta,s)$ then

$$G_Z(t) = pt + (1-p)G_Y(t) \tag{2.50}$$

and the HLZD is retrieved by setting $p=0$.

2.5.1 Moments and Cumulants

From (2.50), the ordinary moments of Z can be easily recovered from the ordinary moments of $Y \sim \text{HLZD}(a,\theta,s)$.

Proposition 2.5.2. *If $Z \sim \text{OIHLZD}(p,a,\theta,s)$ then*

$$\nu_k := \mathbb{E}[Z^k] = p+(1-p)\xi_k \qquad k \in \mathbb{N}, \tag{2.51}$$

with ξ_k the k-th moment of $Y \sim \text{HLZD}(a,\theta,s)$.

Proof. The result follows from (2.50) since

$$M_Z(t) = G_Z(e^t) = pe^t + (1-p)G_Y(e^t) = pe^t + (1-p)M_Y(t).$$

□

For example, we have

$$\begin{aligned} \mathbb{E}[Z] &= p+(1-p)\mathbb{E}[Y], \\ \text{Var}[Z] &= (1-p)\left[\text{Var}[Y] + p(1+\mathbb{E}[Y]^2) - 2\mathbb{E}[Y]\right]. \end{aligned}$$

Proposition 2.5.3. *If $Z \sim \text{OIHLZD}(p,a,\theta,s)$ then*

$$(\nu)_k := \mathbb{E}[Z(Z-1)\cdots(Z-k+1)] = \begin{cases} p+(1-p)(\xi)_1, & k=1, \\ (1-p)(\xi)_k, & k \in \{2,3,\ldots\}. \end{cases} \tag{2.52}$$

with $(\xi)_k$ the k-th factorial moment of $Y \sim \text{HLZD}(a,\theta,s)$.

Proof. The result follows from (2.50) since

$$Q_Z(t) = G_Z(t+1) = p(t+1) + (1-p)G_Y(t+1) = p(t+1) + (1-p)Q_Y(t) \tag{2.53}$$

with $Q_Y(t)$ the factorial mgf of $Y \sim \text{HLZD}(a,\theta,s)$. Expression (2.52) follows by expanding the rhs of (2.53) since we get

$$Q_Z(t) = 1 + [p + (1-p)(\xi)_1]t + (1-p)\sum_{k\geq 2} \frac{t^k}{k!}(\xi)_k.$$

□

The OIHLZD cumulants are

$$\kappa_n(Z) = \mathcal{L}_n(\nu_1, \ldots, \nu_n), \quad n \in \mathbb{N}^+,$$

with $\{\nu_j\}$ moments of $Z \sim \text{OIHLZD}(p,a,\theta,s)$ given in (2.51), and $\mathcal{L}_n$ the n-th logarithmic polynomial (2.12).

2.5.2 Maximum Likelihood Estimation

To estimate the OIHLZD parameters using MLE, let us first rewrite (2.49) using (2.23), that is

$$\begin{aligned} \mathbb{P}(Z=1) &= 1-w, \\ \mathbb{P}(Z=x) &= (1-p)\frac{a(x)g(\theta)^x}{f(\theta)}, \quad x \in \{2,3,\ldots\} \end{aligned} \tag{2.54}$$

where we set $w = (1-p)[1-\mathbb{P}(Y=1)]$, $g(\theta) = \theta$, $a(x) = (a+x)^{-(s+1)}$ and $f(\theta) = T(\theta,a,s)$. Further rewrite (2.54) as

$$\begin{aligned} \mathbb{P}(Z=1) &= 1-w \\ \mathbb{P}(Z=x) &= w\mathbb{P}(W=x), \quad x \in \{2,3,\ldots\} \end{aligned}$$

where W has a One Truncated Hurwitz-Lerch Zeta Distribution (OTHLZD) (Conceição et al., 2017), that is

$$\mathbb{P}(W=x) := \frac{1}{1-\frac{a(1)g(\theta)}{f(\theta)}}\frac{a(x)g(\theta)^x}{f(\theta)}, \quad x \in \mathbb{N}^+. \tag{2.55}$$

Suppose $\mathbf{z} = (z_1, \ldots, z_n)$ is a vector of independent observations of $Z \sim \text{OIHLZD}(p,a,\theta,s)$ and $l(\theta,a,s,w,\mathbf{z}) = \log L_n(\theta,a,s,w \,|\, \mathbf{z})$ is the log-likelihood function with

$$L_n(\theta,a,s,w \,|\, \mathbf{z}) = \prod_{i=1}^{n} \mathbb{P}(Z=z_i).$$

If n_j is the number of times the integer j appears in the vector $\mathbf{z}$ for $j \in \mathbb{N}^+$ then the log-likelihood $\ell_n(\theta, a, s, w \,|\, \mathbf{z})$ can be written as

$$\ell_n(\theta, a, s, w \,|\, \mathbf{z}) = n_1 \log(1-w) + (n-n_1)\log(w) + \sum_{j=2}^{\infty} n_j \log\left(\frac{\mathbb{P}(Y=j)}{1-\mathbb{P}(Y=1)}\right).$$

Now set

$$\ell_{n,1}(w \,|\, \mathbf{z}) = n_1 \log(1-w) + (n-n_1)\log(w) \tag{2.56}$$

and

$$\ell_{n,2}(\theta, a, s \,|\, \mathbf{z}) = \sum_{j=2}^{\infty} n_j \left(\frac{\mathbb{P}(Y=j)}{1-\mathbb{P}(Y=1)}\right). \tag{2.57}$$

From (2.56) and (2.57), the parameters (θ, a, s, w) can be estimated separately, that is the estimation $\hat{w}$ can be recovered from $\ell_{n,1}(w \,|\, \mathbf{z})$ and the estimations $(\hat{\theta}, \hat{a}, \hat{s})$ from $\ell_{n,2}(\theta, a, s \,|\, \mathbf{z})$. The latter ones give the MLE of the parameters of $W \sim \text{OTHLZD}(\theta, s, a)$ in (2.55) using the vector $\mathbf{z}$ restricted to the observations which are greater than 1.

Remark 2.5.4. The estimation $\hat{p}$ of p can be recovered from $\hat{w}$ as

$$\hat{p} = 1 - \frac{\hat{w} T(\hat{\theta}, \hat{a}, \hat{s})}{T(\hat{\theta}, \hat{a}, \hat{s}) - a(1)\hat{\theta}}.$$

□

Fitting results of the three presented discrete distributions on a dataset of daily rainfall measurements will be provided in the next chapter. In the following sections we shall describe other temporal variables connected to the interarrival times.

2.6 Models for Other Temporal Variables

As already noted in the introduction, at a local scale, a conventional method for addressing intermittency in rainfall records involves the statistical analysis of uninterrupted sequences of rainy days, referred to as wet spells *ws*, and sequences of non-rainy days, termed dry spells *ds*, usually under the assumption of their independence. Of the same interest are two additional time variables, strongly associated to *ws* and *ds*, called wet chains *wch* and dry chains *dch*. More precisely, in the following with *wch* we will refer to sequences of rainy days, possibly interrupted by 1-day dry periods, in

(wet) days. Note that also a cluster of rainy days ws is counted in wch, we only add the possibility of having holes (1-day dry periods) inbetween the clusters. Similarly, with dch we will refer to sequences of dry days, possibly interrupted by 1-day rainy periods, in (dry) days. Again, also a cluster of dry days ds is counted in dch, with the added possibility of having 1-day rainy interruptions inbetween the clusters. From a hydrological perspective, the aforementioned wch could be seen as a single rainy sequence in a broad sense. This characterization stems from the understanding that the occurrence of a singular non-rainy day amidst a series of consecutive rainy days does not substantially impact the overall wet status of the sequence, as if the entirety of this period can be attributed to a single meteorological perturbance. Similarly, an analogous rationale can be extended to dch, wherein the interruption of a sequence of dry periods by a lone rainy day may not significantly alter the prevailing dry conditions, as long as we assume that the single rainy interruptions enjoy a limited rainfall depth. The latter hypothesis seems reasonable when considering that the rainfall intensity during such isolated rainy days tends to be minimal, typically lasting for only a few hours. It should be noted that the definition of chains can be easily extended to interruptions longer than a single day (i.e., a 2-day wch is a sequence of $it=1$, $it=2$, and $it=3$, ending with $it>3$). However, the chains can become less and less realistic as single rainy or dry sequences when increasing the lengths of the possible interruptions.

Before proceeding, let us now briefly delineate the connections of the upcoming discussion with the framework of renewal processes delineated in the beginning of this chapter. Recall that we have initially defined a renewal process by specifying a sequence $\{it_n\}_{n\geq 0}$ of i.i.d. interarrival times. Under this assumption, it is easy to prove that a sequence of rainy days ws has a geometric distribution and a sequence of dry days ds share the same distribution as it, albeit with a shifted support (see upcoming Propositions 2.6.1 and 2.6.2). However, while more complex, a natural way to be less restrictive is to select arbitrary distributions for ws and ds. In this scenario, the counterpart of the renewal process is its generalization known as *alternating renewal process* (see, for instance, the appendix of Buishand, 1977). It is used to model a system which alternates between two states, namely the wet state and the dry state in our application. Hence, it is characterized by the sequence of times spent in each state which we denote by $\{ws_n\}_{n\geq 0}$ and $\{ds_n\}_{n\geq 0}$. Again, the fundamental assumption is that $\{ws_n\}_{n\geq 0}$ (resp. $\{ds_n\}_{n\geq 0}$) are i.i.d. with common distribution ws

(resp. ds), and, additionally, the two sequences $\{ws_n\}_{n\geq 0}$ and $\{ds_n\}_{n\geq 0}$ are independent. Finally, as done in the case of the standard renewal process, the description of the rainfall process is finalized by additionally considering the process $\{h_n^*\}_{n\geq 0}$ describing the amount of rain on the rainy days. After this discussion, the approach we proposed in Agnese et al. (2014) and that we delineate in the forthcoming sections can be summarised as follows. With the aim of a parsimonious procedure, one can opt for the ordinary renewal process and, consequently, choose a model for it and then derive the distributions of the other temporal variables of interest, as shall become clear in the next section. If more flexibility is needed, one can consider the alternating renewal process and, as shall be shown later, derive the distribution of it.

We will start by exposing the probabilistic relationships between it, ws, wch and dch in the two different settings delineated above. Additionally, we shall also investigate the derivation of the probability distribution of the total volume of rainfall in a given wet spell, denoted by hcl. The role of the standing assumptions of each scenario will be emphasized, showing how they reasonably simplify the computations involved.

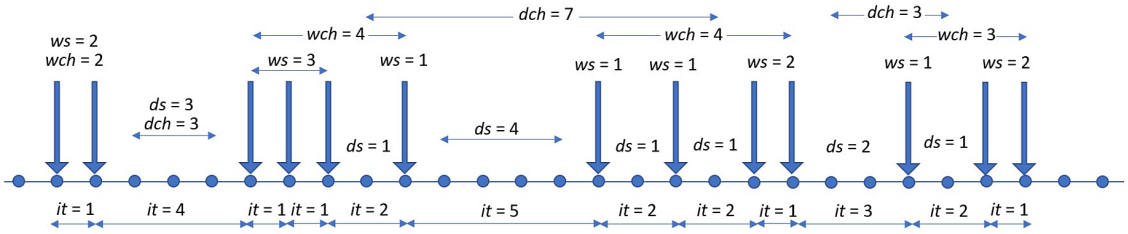

Figure 2.1: Example of a possible realisation of the binary rainfall occurrence process $\{x_n\}_{n\in\mathbb{N}}$ where observations of it, ws, ds, wch and dch are highlighted. The arrows denote the rainy days.

2.6.1 Probabilistic Relationships

As we continue, we will work by alternating between the two following assumptions.

Assumption 1 (Direct method)**.** *The rainfall occurrence process is a standard renewal process and we consider the associated sequence* $\{it_n\}_{n\geq 0}$ *of i.i.d. interarrival times.*

Assumption 2 (Indirect method)**.** *The rainfall occurrence process is an alternating renewal process. We consider the two associated sequences of holding times* $\{ws_n\}_{n\geq 0}$ *and* $\{ds_n\}_{n\geq 0}$*, which are i.i.d. with common distribution* ws *and* ds *respectively. Additionally, the two sequences* $\{ws_n\}_{n\geq 0}$ *and* $\{ds_n\}_{n\geq 0}$ *are independent.*

Let us start by focusing on the wet spells ws and the interarrival times it. It is natural to consider the derivation of ws from it under Assumption 1 first.

Proposition 2.6.1 (ws from it)**.** *Suppose Assumption 1 holds. Then we have that*

$$ws \sim Geom(\mathbb{P}(it>1)).$$

Proof. Consider the case $ws=m$, where m is a positive integer. In order to properly define the event of a wet spell of length m in terms of interarrival times, we have to consider an associated sequence of interarrival times random variables $it_0,\ldots,it_m$. Given the definition of it and ws, we have that

$$\{ws=m\}=\begin{cases}\{it_0>1\} & m=1,\\ \bigcap_{i=0}^{m-1}\{it_i=1\}\cap\{it_m>1\} & m>1.\end{cases} \tag{2.58}$$

Under the assumption $it_n \overset{\text{iid}}{\sim} it$ for all $n\in\mathbb{N}$ it follows that

$$\mathbb{P}(ws=m)=\begin{cases}\mathbb{P}(it>1) & m=1\\ [\mathbb{P}(it=1)]^{m-1}\mathbb{P}(it>1) & m>1.\end{cases} \tag{2.59}$$

□

For what regards ds, the result is straightforward by their definition.

Proposition 2.6.2 (ds from it)**.** *Under Assumption 1 we have that*

$$\mathbb{P}(ds=k)=\frac{\mathbb{P}(it=k+1)}{\mathbb{P}(it>1)},\quad k\in\mathbb{N}^+. \tag{2.60}$$

Proof. A little thought reveals that for any $n\in\mathbb{N}$ and $k\in\mathbb{N}^+$ we have

$$\{ds_n=k\}=\{it_n=k+1|it_n>1\}.$$

The result then follows from Assumption 1. □

Viceversa, one would like to obtain the distribution of it from the ones of ws and and ds under Assumption 2. The approach used for this purpose is different from the one in this section and the derivation is thus postponed to Section 2.6.2.

Let us now focus on deriving the probability distributions of the chains wch and dch from the one of it under Assumption 1, or, alternatively, from the ones of ws and ds under Assumption 2. We start by considering wch. Intuitively, under Assumption 1, we can derive the distribution of wch from the one of it by accounting for the probabilities that it is equal to 1 and/or 2, and weighting the probability of different combinations of sequences of rainy days: indeed, a chain of $m \in \mathbb{N}^+$ wet days could be interrupted by m' holes, where $0 \leq m' \leq m-1$, with different occurrence probabilities at varying m'. Let us be more detailed in the proof of the following proposition.

Proposition 2.6.3 (wch from it)**.** *Suppose Assumption 1 holds. Then we have that*

$$\mathbb{P}(wch = m) = \sum_{i=0}^{m-1} \binom{m-1}{i} \mathbb{P}(it=2)^i \mathbb{P}(it=1)^{m-1-i} \mathbb{P}(it>2), \quad m \in \mathbb{N}^+. \tag{2.61}$$

Proof. Consider the case $wch = m$, where $m > 1$. In order to properly define a wet chain event of length m in terms of interarrival times, we consider an associated $m+1$ sequence of interarrival times random variables $it_0, it_1, \ldots, it_m$. Let $I_m = \{0, \ldots, m-1\}$. Then we have that, from the definition of wch,

$$\{wch = m\} = \bigcup_{i=0}^{m-1} \bigcup_{\substack{(j_1,\ldots,j_i) \in C_i \\ j_{i+1},\ldots,j_m \in I_m \setminus \{j_1,\ldots,j_i\}}} \{it_{j_1}=2, \ldots, it_{j_i}=2, it_{j_{i+1}}=1, \ldots, it_{j_m}=1, it_m>2\} \tag{2.62}$$

where

$$C_i = \{(j_1,\ldots,j_i) : j_k \in \{0,\ldots,m-1\} \text{ for } k=1,\ldots,i \text{ and } j_k < j_{k+1} \text{ for } k=1,\ldots,i-1\}.$$

Consider now that $it_n \overset{\text{iid}}{\sim} it$ for all $n \in \mathbb{N}$. It follows that for every $(j_1,\ldots,j_i) \in C_i$ and $j_{i+1},\ldots,j_m \in \bar{I}_m$ with $i \in I_m$ we have

$$\begin{aligned}
&\mathbb{P}(\{it_{j_1}=2, \ldots, it_{j_i}=2, it_{j_{i+1}}=1, \ldots, it_{j_m}=1, it_m>2\}) \\
&= \prod_{k=1}^{i} \mathbb{P}(\{it_{j_k}=2\}) \prod_{k=i+1}^{m-1} \mathbb{P}(\{it_{j_k}=1\}) \mathbb{P}(\{it_{j_m}>2\}) \\
&= \mathbb{P}(it=2)^i \mathbb{P}(it=1)^{m-1-i} \mathbb{P}(it>2).
\end{aligned}$$

Hence, the result is obtained from the fact that $|\bigcup_{i=0}^{m-1} C_i| = \binom{m-1}{i}$ and from the additivity of $\mathbb{P}$. □

Remark 2.6.4. Let us show that the probability distribution derived for *wch* arises in a simple probabilistic experiment. Consider $p \geq 0$ and $q \geq 0$, playing the role respectively of $\mathbb{P}(it=2)$ and of $\mathbb{P}(it=1)$ in Proposition 2.6.3 and its proof. Consider a sequence of independent trials $E_1, E_2, \ldots$ each with three possible outcomes

$$\begin{cases} \text{Type 1 success} & \text{w.p. } p, \\ \text{Type 2 success} & \text{w.p. } q, \\ \text{Failure} & \text{w.p. } 1-p-q. \end{cases}$$

Then let us consider the random variable

$$X \coloneqq \text{ number of trials before a failure.}$$

We have that by combinatorial arguments

$$\mathbb{P}(X=k) = \sum_{i=0}^{k-1} \binom{k-1}{i} p^i q^{k-1-i}(1-p-q), \quad k \in \mathbb{N}^+.$$

□

It is also of interest to derive the probability distribution of *wch* from the ones of *ws* and *ds* under Assumption 2. The setting and derivation are slightly more convoluted, as can be seen in the proof of the next proposition. Recall that a partition of an integer n is a sequence $\pi = (\pi_1, \pi_2, \ldots)$ of weakly decreasing positive integers, named parts of π, such that $\pi_1 + \pi_2 + \cdots = n$. A different and useful notation is $\pi = (1^{r_1}, 2^{r_2}, \ldots)$, where $r_1, r_2, \ldots$, named multiplicities of π, are the number of parts of π equal to 1,2,... respectively. The length of the partition is $l(\pi) = r_1 + r_2 + \cdots$ and the vector of multiplicities is $m(\pi) = (r_1, r_2, \ldots)$. We write $\pi \vdash n$ to denote that π is a partition of n.

Proposition 2.6.5 (*wch* from *ws* and *ds*)**.** *Suppose Assumption 2 holds. Then we have that for* $m \in \mathbb{N}^+$

$$\mathbb{P}(wch=m) = \sum_{\pi \vdash m} \frac{l(\pi)!}{r_1! \cdots r_m!} \prod_{\pi' \in \pi} \mathbb{P}(ws=\pi')[\mathbb{P}(ds=1)]^{l(\pi)-1}\mathbb{P}(ds>1), \quad (2.63)$$

or, equivalently,

$$\mathbb{P}(wch=m)=\sum_{k=1}^{m} p_{ws}^{k*}(m)[\mathbb{P}(ds=1)]^{k-1}\mathbb{P}(ds>1), \tag{2.64}$$

where we denote the pmf of ws by p_{ws} and the k-fold convolution of p_{ws} by p_{ws}^{k*}.

Proof. Consider the case $wch=m$, with $m\in\mathbb{N}^+$. In order to define the event of a chain of length m in terms of wet spells and dry spells, we have to consider two associated sequences of wet spells random variables $ws_0,ws_1,\ldots,ws_{m^*}$ and dry spells random variables $ds_0,ds_1,\ldots,ds_{m'}$, with $m^*\leq 2m-1$ and $m'\leq 2m$. By the definition of a wet chain wch, we have that

$$\{wch=m\}=\bigcup_{\pi\vdash m}\bigcup_{I_\pi}\{ws_{j_1}=\pi_1,ds_{j_1}=1,ws_{j_2}=\pi_2,ds_{j_2}=1,\ldots,ws_{j_{l(\pi)}}=\pi_{l(\pi)},ds_{j_{l(\pi)}}>1\}, \tag{2.65}$$

where

$$I_\pi=\{(j_1,\ldots,j_{l(\pi)}):j_k\in\{1,2,\ldots,l(\pi)\}\text{ and }j_k\neq j_l\text{ for }k,l=1,\ldots,l(\pi)\text{ with }l\neq k\}.$$

Now suppose that $ws_n\overset{\text{iid}}{\sim}ws$ and independent from $ds_n\overset{\text{iid}}{\sim}ds$ for all $n\in\mathbb{N}$. It follows that for every $(j_1,\ldots,j_{l(\pi)})\in I_\pi$ with $\pi\vdash m$ we have

$$\begin{aligned}\mathbb{P}\Big(\{ws_{j_1}=\pi_1,ds_{j_1}=1,ws_{j_2}=\pi_2,ds_{j_2}=1,\ldots,ws_{j_{l(\pi)}}=\pi_{l(\pi)},ds_{j_{l(\pi)}}>1\}\Big)&\\ =\prod_{k=0}^{l(\pi)}\mathbb{P}(\{ws_{j_k}=\pi_k\})\prod_{k=0}^{l(\pi)-1}\mathbb{P}(\{ds_{j_k}=1\})\mathbb{P}\Big(\{ds_{j_{l(\pi)}}>1\}\Big)&\\ =\mathbb{P}(ws=\pi')[\mathbb{P}(ds=1)]^{l(\pi)-1}\mathbb{P}(ds>1).&\end{aligned}$$

Hence, (2.63) is obtained from the fact that $|I_\pi|=\frac{l(\pi)!}{r_1!\cdots r_m!}$ and from the additivity of $\mathbb{P}$. Now, suppose to rearrange (2.63) in order to specify the sum over partitions of m of a fixed length k, that is

$$\mathbb{P}(wch=m)=\sum_{k=1}^{m}\sum_{\substack{\pi\vdash m\\ l(\pi)=k}}\frac{k!}{r_1!\cdots r_m!}\prod_{\pi'\in\pi}\mathbb{P}(ws=\pi')[\mathbb{P}(ds=1)]^{k-1}\mathbb{P}(ds>1).$$

The term

$$\sum_{\substack{\pi\vdash m\\ l(\pi)=k}}\frac{k!}{r_1!\cdots r_m!}\prod_{\pi'\in\pi}\mathbb{P}(ws=\pi')$$

in the previously centered display is actually the probability that the sum of k i.i.d. wet spells of any duration is equal to m, that is $\mathbb{P}(\sum_{j=1}^{k} ws_j = m)$. The pmf of the sum $\sum_{j=1}^{k} ws_j$ for an arbitrary number k of i.i.d. ws is known and obtained as the the k-fold convolution of p_{ws}, that is p_{ws}^{k*}. Then, we can rewrite (2.63) with a more condensed notation as (2.64). □

The derivation of the distribution of the dry chains dch from the ones of ds and ws can be based upon (2.63). Indeed, a little thought reveals that it is sufficient to reverse the roles of ws and ds in the derivation of (2.63) to obtain the following proposition, stated without proof since the arguments are the same as the ones employed for the previous result.

Proposition 2.6.6 (dch from ws and ds)**.** *Suppose Assumption 2 holds. Then we have that for* $m \in \mathbb{N}^+$

$$\mathbb{P}(dch = m) = \sum_{\pi \vdash m} \frac{l(\pi)!}{r_1! \cdots r_m!} \prod_{\pi' \in \pi} \mathbb{P}(ds = \pi')[\mathbb{P}(ws = 1)]^{l(\pi)-1} \mathbb{P}(ws > 1), \quad (2.66)$$

or, equivalently,

$$\mathbb{P}(dch = m) = \sum_{k=1}^{m} p_{ds}^{k*}(m)[\mathbb{P}(ws = 1)]^{k-1} \mathbb{P}(ws > 1).$$

From (2.66) and the reasoning behind its computation we can infer the derivation of the distribution of dch from the one of it, when taking into account how to obtain ws and ds from it as shown in Propositions 2.6.1 and 2.6.2.

Proposition 2.6.7 (dch from it)**.** *Suppose Assumption 1 holds. Then we have that for* $m \in \mathbb{N}^+$

$$\mathbb{P}(dch = m) = \sum_{\pi \vdash m} \frac{l(\pi)!}{r_1! \cdots r_m!} \prod_{\pi' \in \pi} \mathbb{P}(it = \pi' + 1)[\mathbb{P}(it > 1)]^{l(\pi)-1} \mathbb{P}(it = 1). \quad (2.67)$$

Let us now derive the probability distribution of the total volume of rainfall in a given wet spell, that is hcl, from the ones of the interarrival times and rainfall depths.

Proposition 2.6.8 (hcl from it and h)**.** *Suppose Assumption 1 holds. Additionally, consider a sequence of rainfall depths* $h_n \overset{\text{iid}}{\sim} h$ *and independent from the sequence of interarrival times* $it_n \overset{\text{iid}}{\sim} it$ *for all* $n > 0$*. Then for* $k \in \mathbb{N}^+$ *we have that*

$$\mathbb{P}(hcl = k) = \mathbb{P}(h = k)\mathbb{P}(it > 1) + \sum_{m=2}^{k} p_h^{m^*}(k)\,[\mathbb{P}(it = 1)]^{m-1} \mathbb{P}(it > 1), \quad (2.68)$$

where p_h is the pmf of h and $p_h^{m^*}$ is the m-fold convolution of p_h with itself.

Proof. Initially fix a wet spell ws of length m. To properly define the event of the total volume of rainfall being equal to k, we consider an associate m sequence of rainfall depth random variables $h_1,\dots,h_m$. We have that for $k>m$

$$\{hcl=k,ws=m\}=\bigcup_{\substack{\pi\vdash k\\ l(\pi)=m}}\bigcup_{(j_1,\dots,j_m)\in C_m}\{h_{j_1}=\pi_1,\dots,h_{j_m}=\pi_m,ws=m\}, \tag{2.69}$$

where

$$C_m=\{(j_1,\dots,j_m):j_k\in\{0,\dots,m-1\}\text{ and }j_k\neq j_l\text{ for }k,l=1,\dots,m\}.$$

Under our assumptions, we have that the sequence of rainfall depths $h_n\overset{\text{iid}}{\sim}h$ is independent from the sequence $ws_n\overset{\text{iid}}{\sim}ws$ obtained from the sequence $it_n\overset{\text{iid}}{\sim}it$, $n\in\mathbb{N}$. Using combinatorial arguments analogous to the ones employed in the proofs of the previous propositions, we can write

$$\mathbb{P}(hcl=k,ws=m)=\sum_{\substack{\pi\vdash k\\ l(\pi)=m}}\frac{m!}{r_1!\cdots r_k!}\prod_{\pi'\in\pi}\mathbb{P}(h=\pi')\mathbb{P}(ws=m). \tag{2.70}$$

where $\sum_{\substack{\pi\vdash k\\ l(\pi)=m}}\frac{m!}{r_1!\cdots r_k!}\prod_{\pi'\in\pi}\mathbb{P}(h=\pi')$ amounts to computing the m-fold convolution of p_h with himself[1], that is $p_h^{m^*}$. Then, by simply marginalising, we can write for $k\in\mathbb{N}^+$ that

$$\begin{aligned}\mathbb{P}(hcl=k)&=\sum_{m=1}^{\infty}\mathbb{P}(hcl=k,ws=m) &(2.71)\\ &=\sum_{m=1}^{\infty}p_h^{m^*}(k)\mathbb{P}(ws=m) &(2.72)\\ &=\sum_{m=1}^{k}p_h^{m^*}(k)\mathbb{P}(ws=m) &(2.73)\\ &=\mathbb{P}(h=k)\mathbb{P}(it>1)+\sum_{m=2}^{k}p_h^{m^*}(k)\left[\mathbb{P}(it=1)\right]^{m-1}\mathbb{P}(it>1), &(2.74)\end{aligned}$$

where the third equality comes from the fact that if $m>k$, $\{\pi\vdash k|l(\pi)=m\}=\emptyset$[2] and the fourth equality comes from Proposition 2.6.1. □

[1] Indeed, under our hyphoteses, the random variable hcl is equal to the random sum $\sum_{i=1}^{ws}h_i$ where $h_i\overset{\text{iid}}{\sim}h$ and all independent from it and thus from ws.

[2] This is equivalent to saying that if $m>k$, $k\notin\text{Supp}(\sum_{i=1}^{ws}h_i|ws=m)$.

Remark 2.6.9. Note that since ws appears in (2.73) and the distribution of ds can be easily recovered from the one of it, the pmf of hcl can be also derived from the ones of ds, ws and h. □

2.6.2 Relationships Between the Variables' Counts in a Series of Observations

In the previous section we have postponed the derivation of the distribution of it from the ones of ds and ws under Assumption 2. The reason why is that an intuitive probablistic approach, as the one employed in the previous derivations, seems to be unattainable in this case. Therefore we have proceeded in a slightly different way, which stems from relationships regarding the counts of temporal variables arising from the same series of rainfall observations, that shall be described below and also hold independent interest. Let a series of rainfall data be defined as $\mathbf{h}=\{h_1,\ldots,h_n\}$, where h (mm) is the rainfall depth recorded at a fixed uniform unit τ of time (e.g., a day). Let any $k\in\{1,\ldots,n\}$. A day k is considered rainy if the rainfall depth $h_k\geq h^*$, where h^* is a fixed rainfall threshold. The sub-series of $\mathbf{h}$ of the rainy days can be defined as the event series $\mathbf{t}=\{t_1,\ldots,t_{n_r}\}$, where $n_r\leq n$ is an integer multiple of the time-scale τ. The sequence built with the times elapsed between each element of $\boldsymbol{t}$ (except the first one) and the immediately preceding one is defined as the interarrival time series $\mathbf{it}=\{it_2,\ldots,it_{n_r}\}$. Consider $k\in\{1,\ldots,n_r\}$. If $it_k=1$, the rainy day k immediately follows another rainy day $k-1$, while an isolated rainy day k is identified if it satisfies the condition

$$it_k>1 \text{ and } it_{k+1}>1. \tag{2.75}$$

Thus, any sequence of $m\in\mathbb{N}^+$ consecutive 1 values in $\mathbf{it}$ represents an uninterrupted sequence of rainy days and is an observation of a wet spell ws of duration (size) $m+1$ and we will refer with $\mathbf{ws}$ to the series of these observations. Additionally, the series of non-rainy days $\mathbf{ds}$, composed of observations of the dry spell ds, can be derived from $\mathbf{it}$ by using the relationship

$$ds_k=it_k-1 \text{ for any } it_k>1. \tag{2.76}$$

Hence, as we continue, note that with $\mathbf{it}$, $\mathbf{ws}$ and $\mathbf{ds}$ we refer to series of observations of it, ws and ds respectively obtained from the same $\mathbf{h}$. We denote by $N_{\mathbf{x}=k}$ (resp. $N_{\mathbf{x}>k}$) the cardinality of the set $\{x\in\mathbf{x}:x=k\}$ (resp. $\{x\in\mathbf{x}:x>k\}$), where $\mathbf{x}\in\{\mathbf{it},\mathbf{ws},\mathbf{ds}\}$. Moreover, let $N_{\mathbf{x}}$ be the length of the

series $\mathbf{x}$. It is easily derived that

$$N_{\boldsymbol{it}=1} = \sum_{k=\min\limits_{x\in\mathbf{ws}}}^{\max\limits_{x\in\mathbf{ws}}} N_{\mathbf{ws}=k}\,(k-1), \tag{2.77}$$

$$N_{\boldsymbol{it}>1} = N_{\mathbf{ws}}, \tag{2.78}$$

$$N_{\boldsymbol{it}=k} = N_{\mathbf{ds}=k}\,k \text{ whenever } k>1. \tag{2.79}$$

Let us also recall the concept of empirical frequency. For a series of observations $\mathbf{x}\in\{\mathbf{it},\mathbf{ws},\mathbf{ds}\}$, we define the empirical frequency distribution as the map $f_{\mathbf{x}}:\mathbb{N}^+\to[0,1]$ defined as

$$f_{\mathbf{x}}(k) = \frac{1}{N_{\mathbf{x}}}\sum_{i\in\mathbb{N}^+}\mathbf{1}_{\{x_i=k\}} = \frac{N_{\mathbf{x}=k}}{N_{\mathbf{x}}}, \quad k\in\mathbb{N}^+.$$

The domain of $f_{\mathbf{x}}$ has been defined as $\mathbb{N}^+$ since the three considered random variables have such a common support. Throughout the first two chapters of this thesis, the vector $\{f_{\mathbf{x}}(k)\}_{k\geq 1}$ shall usually be referred to as the empirical frequencies or observed frequencies of the series $\mathbf{x}$.

Proposition 2.6.10. *We have that*

$$f_{\boldsymbol{it}}(1) = \frac{\sum\limits_{k=\min\limits_{x\in\mathbf{ws}}}^{\max\limits_{x\in\mathbf{ws}}} f_{\mathbf{ws}}(k)\,(k-1)}{\sum\limits_{k=\min\limits_{x\in\mathbf{ws}}}^{\max\limits_{x\in\mathbf{ws}}} f_{\mathbf{ws}}(k)\,k}. \tag{2.80}$$

Proof. Exploiting the properties of the empirical frequencies and the relationships (2.77) and (2.78) we have that

$$\frac{\sum_{k=\min_{x\in\mathbf{ws}}}^{\max_{x\in\mathbf{ws}}} f_{\mathbf{ws}}(k)\,(k-1)}{\sum_{k=\min_{x\in\mathbf{ws}}}^{\max_{x\in\mathbf{ws}}} f_{\mathbf{ws}}(k)\quad k} = \frac{\sum_{k=\min_{x\in\mathbf{ws}}}^{\max_{x\in\mathbf{ws}}} f_{\mathbf{ws}}(k)\,(k-1)}{\sum_{k=\min_{x\in\mathbf{ws}}}^{\max_{x\in\mathbf{ws}}} f_{\mathbf{ws}}(k)\,(k-1)+\sum_{k=\min_{x\in\mathbf{ws}}}^{\max_{x\in\mathbf{ws}}} f_{\mathbf{ws}}(k)} =$$

$$= \frac{\sum_{k=\min_{x\in\mathbf{ws}}}^{\max_{x\in\mathbf{ws}}} f_{\mathbf{ws}}(k)\,(k-1)}{\sum_{k=\min_{x\in\mathbf{ws}}}^{\max_{x\in\mathbf{ws}}} f_{\mathbf{ws}}(k)\,(k-1)+1} = \frac{\sum_{k=\min_{x\in\mathbf{ws}}}^{\max_{x\in\mathbf{ws}}} f_{\mathbf{ws}}(k)\,N_{\mathbf{ws}}\,(k-1)}{\sum_{k=\min_{x\in\mathbf{ws}}}^{\max_{x\in\mathbf{ws}}} f_{\mathbf{ws}}(k)\,N_{\mathbf{ws}}\,(k-1)+N_{\mathbf{ws}}} =$$

$$= \frac{\sum_{k=\min_{x\in\mathbf{ws}}}^{\max_{x\in\mathbf{ws}}} N_{\mathbf{ws}=k}\,(k-1)}{\sum_{k=\min_{x\in\mathbf{ws}}}^{\max_{x\in\mathbf{ws}}} N_{\mathbf{ws}=k}\,(k-1)+N_{\mathbf{ws}}} = \frac{N_{\boldsymbol{it}=1}}{N_{\boldsymbol{it}=1}+N_{\boldsymbol{it}>1}} = \frac{N_{\boldsymbol{it}=1}}{N_{\boldsymbol{it}}} = f_{\boldsymbol{it}}(1).$$

□

Let p_{it} and p_{ws} be the pmfs of it and ws respectively.

Proposition 2.6.11. *Suppose that* $\boldsymbol{it}$ *and* $\boldsymbol{ws}$ *are series of i.i.d. observations of* it *and* ws*, obtained from the same series* $\boldsymbol{h}=\{h_1,h_2,\ldots,h_n\}$*, where* n *is the sample size. Then we have that*

$$p_{it}(1)=\frac{\mathbb{E}[ws]-1}{\mathbb{E}[ws]}. \tag{2.81}$$

Proof. Let $\bar{\mathbf{x}}_n=\sum_{i=1}^{n}\frac{x_i}{n}$ denote the sample mean of a series of observations $\mathbf{x}$. Thus (2.80) can be rewritten as

$$f_{\boldsymbol{it}}(1)=\frac{\bar{\mathbf{ws}}_n-1}{\bar{\mathbf{ws}}_n}. \tag{2.82}$$

Let the sample size n go to infinity. Then we have that $N_{\mathbf{it}}\to\infty$ and $N_{\mathbf{ws}}\to\infty$. Under the assumptions that **it** and **ws** are series of i.i.d. observations of it and ws, the Strong Law of Large Numbers (SLLN) implies that $\bar{\mathbf{ws}}_n$

converges almost surely to $\mathbb{E}[ws]$ as $n \to \infty$. In the same way, the SLLN implies that the empirical frequencies $f_{\mathbf{it}}(k)$ a.s. converge to the probabilities $p_{it}(k)$ as $n \to \infty$ for all $k \in \mathbb{N}$. The uniqueness of the limit then provides the desired result. □

Remark 2.6.12. Expression (2.80) can then be used for computing the probability of $it = 1$ when interested in deriving the distribution of it from the ones of ds and ws under Assumption 2. Then for $k > 1$ we use (2.76) and we get

$$\mathbb{P}(it = k) = \begin{cases} \frac{\mathbb{E}[ws]-1}{\mathbb{E}[ws]} & k = 1, \\ \mathbb{P}(ds = k-1)[1 - \mathbb{P}(it = 1)] & k > 1. \end{cases} \tag{2.83}$$

□

2.7 The Direct and Indirect Methods

Given the relationships presented above and recalling the discussion at the beginning of Section 2.6, one could proceed in the two following ways to provide a model for the temporal variables described above.

Direct method

Assume the rainfall occurrence process is governed by a renewal process and consider the associated sequence of interarrival times $it_n \overset{\text{iid}}{\sim} it$. Additionally, assume that $it_n \overset{\text{iid}}{\sim} it$ is independent from the sequence $h_n \overset{\text{iid}}{\sim} h$ describing the rainfall amount process. Knowing p_{it} and p_h, derive the pmfs of $\{ws, ds, wch, dch, hcl\}$;

Indirect method

Assume the rainfall occurrence process is governed by an alternating renewal process and consider the independent associated sequences of wet spells $ws_n \overset{\text{iid}}{\sim} ws$ and dry spells $ds_n \overset{\text{iid}}{\sim} ds$. Additionally, assume that $ws_n \overset{\text{iid}}{\sim} ws$ is independent from the sequence $h_n \overset{\text{iid}}{\sim} h$ describing the rainfall amount process. Knowing p_{ws}, p_{ds} and p_h obtain the pmfs of $\{it, wch, dch, hcl\}$.

In Fig. 2.2 we can see a summary of them.

The next chapter is devoted to applying the contents of what has been presented in these sections to data consisting of daily rainfall measurements.

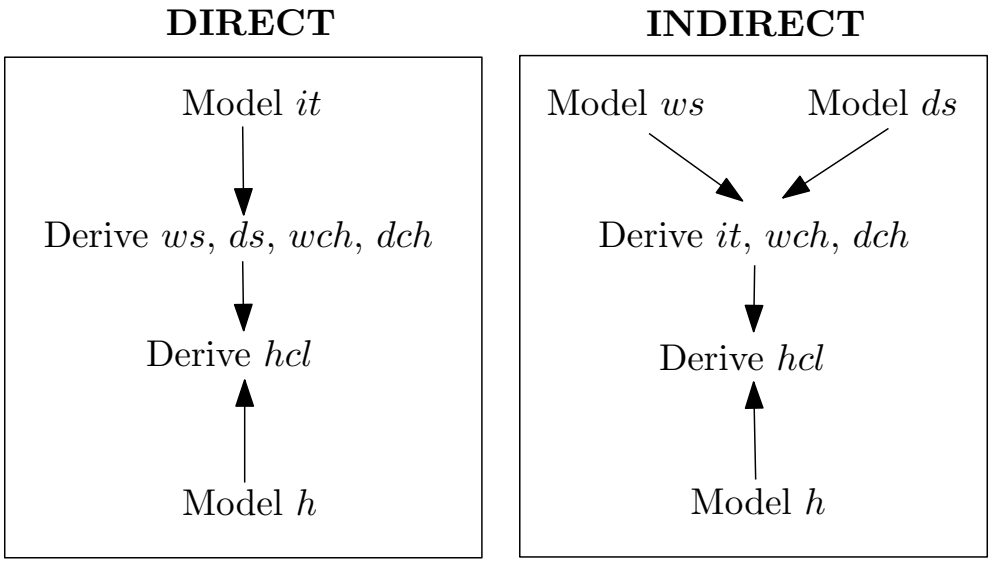

Figure 2.2: Direct and indirect methods.

More specifically, a first part shall present comparisons of fittings of the HLZD, PSHZLD and OIHLZD. A second part will illustrate a statistical application and comparison of the direct method and indirect method.

Appendix

2.A Cumulants, Moments and Bell Polynomials

Recall the definition of the moment generating function (mgf) or a random variable X, that is

$$M_X(z) = \mathbb{E}\left[e^{zX}\right].$$

If X is a random variable whose moment generating function $M_X(z) < \infty$ for all z in an open neighborhood of 0, its moments $\{m_n\}_{n\geq 0}$ can be recovered as $m_n = \left.\dfrac{d^n M_X}{dz^n}\right|_{z=0}$.

Definition. *If X is a random variable having moment generating function $M_X(z) < \infty$ for all z in an open neighborhood of 0, its cumulants $\{\kappa_n\}_{n\geq 0}$ are the coefficients of the series*

$$K_X(z) = \sum_{n=1}^{\infty} \kappa_n \frac{z^n}{n!} = \log M_X(z) \tag{2.84}$$

for all z in some possibly smaller neighborhood of 0. The series in (2.84) is called the cumulant generating function and the cumulants $\{\kappa_n\}_{n\geq 0}$ can be evaluated as

$$\kappa_n = \left.\frac{\mathrm{d}^n K(t)}{\mathrm{d}t^n}\right|_{t=0}.$$

The moments can be recovered in terms of cumulants and viceversa by firstly simply noting that

$$m_n = \left.\frac{d^n \exp(K_X(t))}{dt^n}\right|_{t=0} \quad \text{and} \quad \kappa_n = \left.\frac{d^n \log(M_X(t))}{dt^n}\right|_{t=0}$$

and then using *Faà di Bruno's formula* (see, e.g., Johnson, 2002) for higher derivatives of composite functions, which, in terms of the partial exponential Bell polynomials reads

$$\frac{d^n}{dx^n} f(g(x)) = \sum_{k=0}^{n} f^{(k)}(g(x)) B_{n,k}\left(g'(x), g''(x), \ldots, g^{(n-k+1)}(x)\right), \quad n \in \mathbb{N}+,$$

for functions f and g with a sufficient number of derivatives. Hence, a little thought reveals that for $n \in \mathbb{N}^+$ we have

$$m_n = B_n(\kappa_1, \ldots, \kappa_n),$$

$$\kappa_n = \mathcal{L}(m_1, \ldots, m_n).$$

Moreover, thanks to the properties of the polynomials involved, the following well-known recursive relation holds

$$m_n = \kappa_n + \sum_{i=1}^{n-1} \binom{n-1}{i-1} \kappa_i m_{n-i}. \tag{2.85}$$

Chapter 3

Analysis and Modelling of Temporal Variables at the Daily Scale in a Large Range of Rainfall Regimes Across Europe

This chapter is based on Agnese et al. (2022)

Modelling the frequency of interarrival times and rainfall depths with the poisson hurwitz-lerch zeta distribution, *Fractal and Fractional* 6(9)

and on Baiamonte et al. (2024)

Applying different methods to model dry and wet spells at daily scale in a large range of rainfall regimes across europe, *Advances in Statistical Climatology, Meteorology and Oceanography*, 10(1), 51–67.

In the first part of the previous chapter we have studied the HLZ family of discrete distributions, the PSHZLD and the OIHLZD. Next, we have delineated two approaches for modelling various rainfall temporal variables that we denoted as direct method (DM) and indirect method (IM). Following this partition, in this chapter we will first present a comparison of the fits of the three previously mentioned discrete distributions to interarrival times and rainfall depths using a real dataset. Secondly, using the same dataset, we will apply the DM and the IM, focusing on how a comparison of the two

procedures can shed light on some aspects of the data involved.

The available dataset is composed of various series of observations **h** corresponding to daily rainfall records. To discriminate between rainy and no-rainy days, the fixed rainfall threshold h^* is chosen equal to 1 mm, according to the conventional value established by the World Meteorological Organisation. Using this, from any **h** we can derive the series of rainfall amounts on the rainy days denoted by $\mathbf{h}^*$, which, for brevity, shall usually be referred to as the series of rainfall depths. As explained in Section 2.6, from **h** we can additionally obtain the series of observations **it**, **ws**, **ds**, **wch**, **dch** of rainfall temporal variables *it*, *ws*, *ds*, *wch*, *dch*, which shall also be analysed. The next section gives a preliminary description of the data.

3.1 The Data

The dataset at our hands comprises six time series gathered across various European latitudes, ranging from 38°N to 58°N, spanning locations from Trapani and Floresta in Sicily to Stornoway in Northern Scotland. Fig 3.1.1 illustrates the geographical positions and altitudes of these six stations. Each rain gauge station provides approximately 70 years of recorded data: Ceva (1950-2016), Floresta (1951-2015), Oxford (1950-2017), Stornoway (1950-2020), Torino (1950-2017), and Trapani (1950-2015), with only a minimal amount of missing data. Given the significance of rainfall patterns in determining the distribution of rainy and non-rainy days, the analyses were not only conducted for the entire year Y but also for two different additional subsets, S1 from April to September, and S2 from October to March.

The analyzed stations exhibit distinct rainfall regimes, as depicted in Fig. 3.1.2 by the average number of rainy days per month (panel a) and the fraction of the yearly-average rainy days per month (panel b), which is a standardization of the values in panel a. Trapani (TRA) and Floresta (FLO) typify the Mediterranean climate, characterized by pronounced seasonality in rainfall (see Fig. 3.1.2b) with most precipitation occurring in the S2 season. These stations differ significantly in their average annual rainfall: TRA receives a low 420 mm, while FLO receives a high 1133 mm, which is reflected in the disparate number of rainy days per month (Fig. 3.1.2a). Torino (TOR) and Ceva (CEV) exhibit a mid-latitude sublitoranean climate, with a high frequency of rain in the spring (Fig. 3.1.2b). CEV also shows a secondary peak in autumn, primarily due to the Tyrrhenian Sea's summer

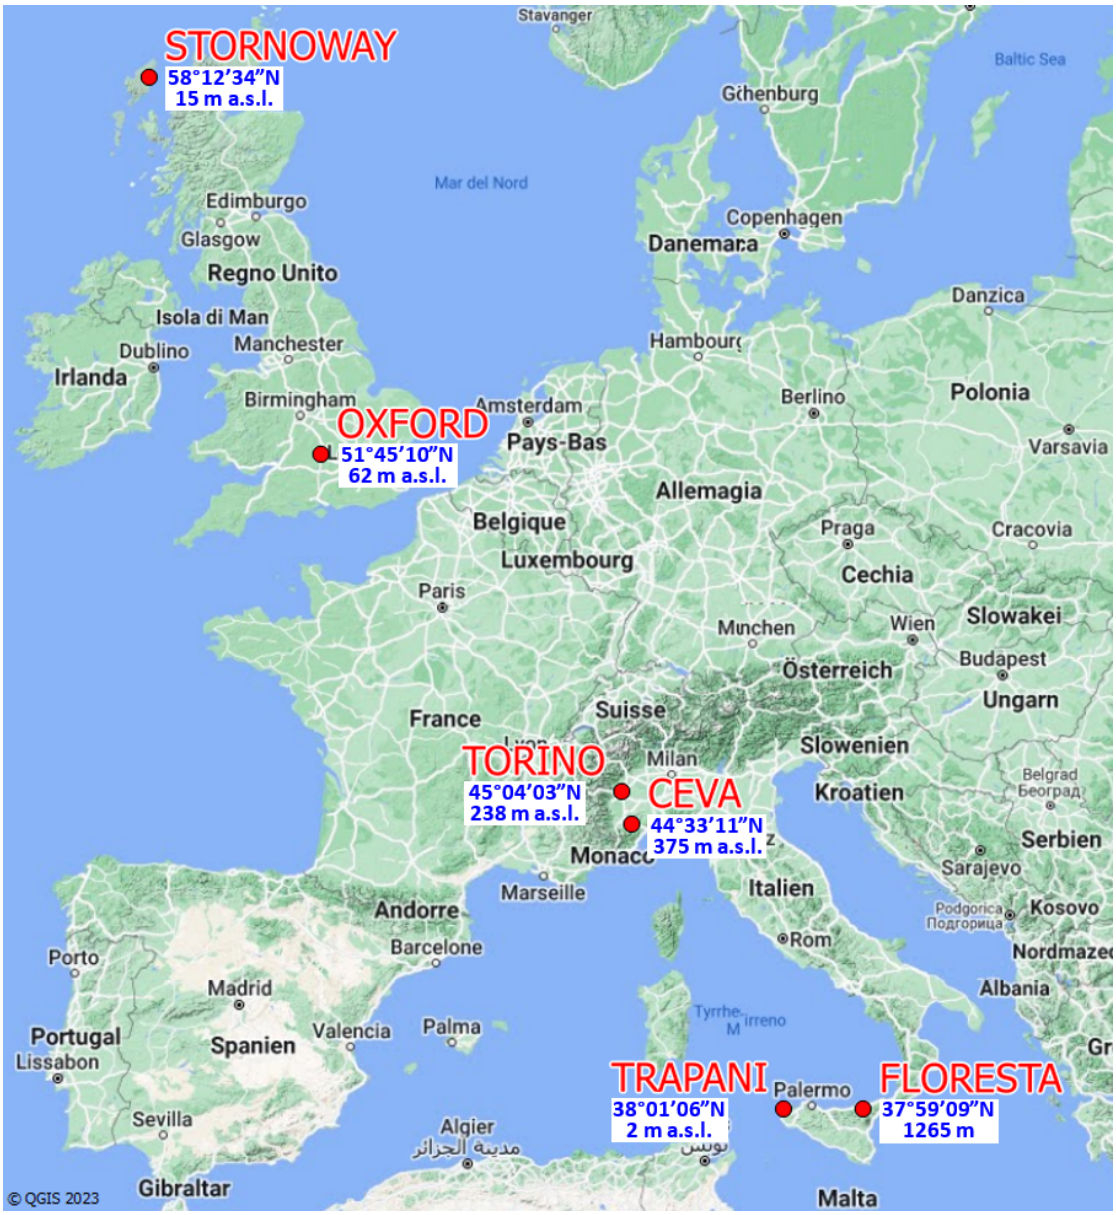

Figure 3.1.1: The six stations under consideration, along with their respective latitudes and elevations above sea level (© QGIS 2023).

warming. Despite this, both stations have similar total annual rainfall (829 mm for TOR and 836 mm for CEV) and similar numbers of rainy days (Fig. 3.1.2a). Oxford (OXF) represents a Northern European station with a relatively low average annual rainfall of 592 mm, evenly distributed throughout the year. In contrast, Stornoway (STW), located in far north-western Scotland, experiences very high rain frequency and a higher annual

rainfall of 1072 mm due to the influence of wet fronts from the Atlantic Ocean. Both UK stations display low seasonality compared to the other stations (see Fig. 3.1.2b).

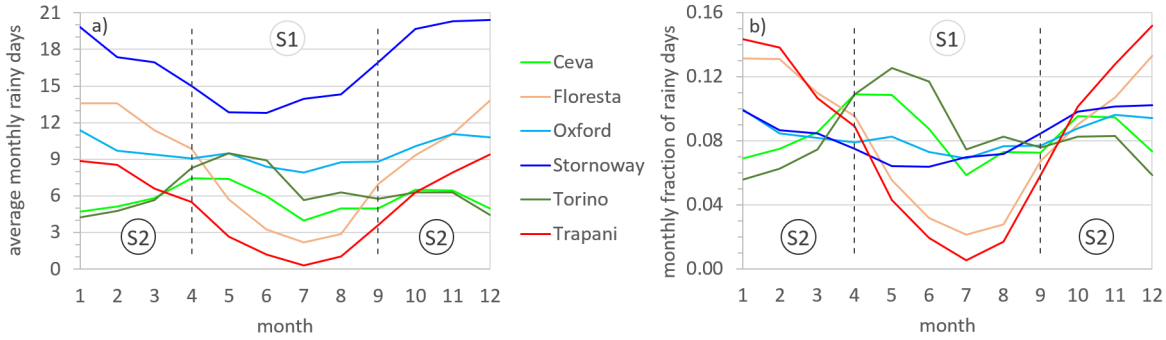

Figure 3.1.2: Time variability of: a) average number of rainy days in each month, and b) fraction of yearly-average rainy days in each month. Dashed lines delimit the two seasons S1 (April - September) and S2 (October-March).

For the stations TRA and FLO, the seasons S1 and S2 distinctly correspond to periods of low and high frequency of rain events, respectively (Fig. 3.1.2b). A similar pattern is observed for OXF and STW, although the differences between the two seasons are less pronounced. Due to the extensive length of the data records, the sample sizes remain substantial even for the two seasonal datasets, as summarised in Table 3.1.1. The sample size is below 500 only for *dch* in TRA for S1, likely due to the many extended dry periods during the dry season.

It is noteworthy that the splitting of the two seasons of CEV and TOR was done differently in a previous paper (Berro et al., 2019). However, in this chapter the same splitting in two six-month seasons is used for the sake of the homogeneity of the present analysis (Fig. 3.1.2a and 3.1.2b).

3.1.1 Preliminary Tests on Observed Records

Before attempting to proceed with either fitting the three discrete distributions of interest or with the comparison of DM and IM, it is necessary to verify if the underlying assumptions are met, or, at least, roughly met, by the available data. That is, we shall check if the rainfall measurements available approximately abide to the conditions of a renewal process and/or of an alternating renewal process. Let us start with the former. In this case, we have to inspect the independence of the interarrival times, which can be done by classical tests on a **it** series. Given the daily frequency of the

Table 3.1.1: Sample sizes of the time variables for the six stations and for the three periods: Y, S1 (from April to September), and S2 (from October to March).

Station	Period	it	ws–ds	wch	dch
CEV	Y	4569	2530	2058	1102
	S1	2335	1357	1159	528
	S2	2203	1154	890	568
FLO	Y	6495	2867	2048	1570
	S1	1966	1053	858	481
	S2	4520	1811	1188	1087
OXF	Y	7933	3920	2559	1915
	S1	3664	1883	1254	902
	S2	4262	2034	1303	1012
STW	Y	14227	4126	2202	2649
	S1	6065	2216	1259	1337
	S2	8145	1909	943	1312
TOR	Y	5171	2726	2145	1289
	S1	3023	1668	1277	726
	S2	2147	1057	868	563
TRA	Y	4064	2256	1755	970
	S1	955	661	574	220
	S2	3108	1594	1181	750

measurements, we anticipated the potential presence of only a lag-1 serial correlation in **it**. Therefore, we decided to compute the sample Kendall, Spearman and Pearson coefficients on the dataset composed by couples of subsequent observations contained in a **it** series and then, for each coefficient, to conduct a permutation test with null hypothesis its equality to zero. We chose a permutation test to adjust for the presence of ties in the data. Table 3.1.2 provides the results. Since either only a weakly or absent correlation is found, we should be able to safely proceed with the intended fittings and with the DM. For the alternating renewal process assumption, we mainly have to verify that the alternating dry spells and wet spells are independent from each other. To achieve this, as above, we have computed the sample Kendall, Spearman and Pearson coefficients and performed permutation tests on the bivariate dataset $\{(ws_i, ds_i)\}_{i=1}^{N_{\mathbf{ws}}}$ of successive dry spells and wet spells built by suitably concatenating the series **ws** and **ds**.

As shown in Table 3.1.3, the measured sample associations are either not significant or very low and significant, the latter result probably caused by the substantial sample size. Again, we shall safely assume that rainfall occurrence process is an alternating renewal process within the IM.

Station	Kendall	Spearman	Pearson
CEV (Y)	0.03	0.039*	0.03
CEV (S1)	0.06*	0.077*	0.021
CEV (S2)	-0.004	-0.005	0.036
FLO (Y)	0.048*	0.058*	0.151*
FLO (S1)	0.044	0.057	0.036
FLO (S2)	0	0	0.03
OXF (Y)	0.043*	0.052*	0.046*
OXF (S1)	0.028	0.034	0.019
OXF (S2)	0.053*	0.063*	0.064*
STW (Y)	0.096*	0.106*	0.078*
STW (S1)	0.05*	0.057*	0.043*
STW (S2)	0.104*	0.113*	0.08*
TOR (Y)	0.018	0.022	0.059*
TOR (S1)	0.043*	0.054*	0.05*
TOR (S2)	-0.014	-0.018	0.036
TRA (Y)	0.06*	0.077*	0.14*
TRA (S1)	0.019	0.027	0.034
TRA (S2)	0.016	0.02	0.06*

Table 3.1.2: Sample Kendall, Spearman and Pearson coefficients computed on the dataset composed by couples of subsequent observations contained in a **it** series for all the stations and subdivisions considered. The asterisk indicates significance at a 0.01 level in a permutation test for the equality to zero of the statistic.

Now, recall that in the DM and IM we hinted at the possibility of deriving the distribution of the volume of rainfall in a wet spell. In this case, according to our assumptions, we should investigate the dependence between the length of a wet spell and the rainfall depths contained therein. To achieve this, firstly, from a series **h** and the associated series **ws**, a bivariate dataset $(\mathbf{ws}, \mathbf{h})$ is built in the following way. Consider the first wet spell observation ws_1 in **ws**, of arbitrary length $m \in \mathbb{N}^+$, and let $h_1, \ldots, h_m$ be the associated rainfall depths observations. Then, the corresponding m couples $(ws_1, h_1), (ws_1, h_2), \ldots, (ws_1, h_m)$ form the first m elements of $(\mathbf{ws}, \mathbf{h})$. Repeating this process for all the elements of **ws** provides $(\mathbf{ws}, \mathbf{h})$. Secondly, as previously done, the sample Kendall, Spearman and Pearson coefficients are computed on $(\mathbf{ws}, \mathbf{h})$ and for each of the three aforementioned test statistics a permutation test is conducted with null hypothesis their equality to zero. The values of the statistics are found in Table 3.1.4, where an asterisk indicates the significance at a 0.01 level with respect to the previously mentioned tests. An inspection of Table 3.1.4 reveals, unsurprisingly, the presence of a significant dependence between the rainfall depths and the

Station	Kendall	Spearman	Pearson
CEV (Y)	-0.014	-0.017	0.005
CEV (S1)	-0.016	-0.02	-0.014
CEV (S2)	-0.038	-0.048	-0.021
FLO (Y)	-0.08*	-0.1*	-0.096*
FLO (S1)	-0.054	-0.068	-0.066*
FLO (S2)	-0.018	-0.022	-0.038
OXF (Y)	-0.015	-0.019	-0.011
OXF (S1)	-0.021	-0.026	-0.018
OXF (S2)	-0.008	-0.01	0.026
STW (Y)	-0.078*	-0.097*	-0.052*
STW (S1)	-0.043*	-0.053	-0.045
STW (S2)	-0.09*	-0.113*	-0.029
TOR (Y)	-0.016	-0.021	-0.007
TOR (S1)	-0.049*	-0.06	-0.069*
TOR (S2)	-0.003	-0.005	-0.003
TRA (Y)	-0.052*	-0.065*	-0.095*
TRA (S1)	-0.061	-0.076	-0.096
TRA (S2)	0.003	0.004	-0.019

Table 3.1.3: Sample Kendall, Spearman and Pearson coefficients computed on the on the bivariate dataset $\{(ws_i, ds_i)\}_{i=1}^{N_{\mathbf{ws}}}$ of successive dry spells and wet spells for all the stations and subdivisions considered. The asterisk indicates significance at a 0.01 level in a permutation test for the equality to zero of the statistic.

corresponding wet spell for some stations. Therefore, the derivation of the distribution of the volume of rainfall in a sequence of rainy days shall not be discussed in the DM and IM and will not be treated in the thesis. A deeper analysis of the association between the length of a wet spell and the rainfall depths contained therein is left for future work.

3.2 Fitting of Interarrival Times and Rainfall Depths

As already mentioned in the beginning of this chapter, we will start this analysis by comparing the fittings of the HLZD, the PSHLZ and the OIHLZD on the available samples of it and h. Note that the PSHLZD has support $\mathbb{N}$ and the rvs it and h naturally have support $\mathbb{N}^+$. To circumvent this, we simply considered a shift of support labels, or, equivalently, the shifted rvs $it-1$ and $h-1$. In all the scenarios, to estimate the parameters of the three considered distributions we have used the respective ML estimations procedures as explained in Sections 2.3.6, 2.4.3 and 2.5.2. To assess the goodness of fit, we have employed a simulated χ^2 goodness-of-fit test. When

Station	Kendall	Spearman	Pearson
CEV (Y)	0.156*	0.207*	0.195*
CEV (S1)	0.128*	0.168*	0.151*
CEV (S2)	0.167*	0.225*	0.206*
FLO (Y)	0.083*	0.114*	0.058*
FLO (S1)	0.098*	0.131*	0.071*
FLO (S2)	0.074*	0.102*	0.045*
OXF (Y)	0.08*	0.105*	0.054*
OXF (S1)	0.06*	0.079*	0.034
OXF (S2)	0.099*	0.13*	0.088*
STW (Y)	0.147*	0.201*	0.148*
STW (S1)	0.087*	0.117*	0.093*
STW (S2)	0.145*	0.201*	0.14*
TOR (Y)	0.14*	0.187*	0.146*
TOR (S1)	0.089*	0.12*	0.076*
TOR (S2)	0.216*	0.288*	0.253*
TRA (Y)	0.059*	0.078*	0.02
TRA (S1)	0.088*	0.109*	0.064
TRA (S2)	0.052*	0.068*	0.025

Table 3.1.4: Sample Kendall, Spearman and Pearson coefficients computed on the dataset $(\mathbf{ws}, \mathbf{h})$ defined as in the end of Section 3.1.1. The asterisk indicates significance at a 0.01 level in a permutation test for the equality to zero of the statistic.

the distributions exhibit a long tail, it is known that the classical χ^2 test might be biased due to the presence of numerous small class sizes (with less than 5 elements) and strong asymmetry. Therefore, we decided to proceed by reconstructing the distribution of the χ^2 statistic under the null hypothesis via Monte Carlo simulation (see, for example, Hope, 1968). For what regards the interarrival times, more detailed results of the goodness of fit tests shall be actually presented in Section 3.3.5. To further inspect the differences between the distributions, as classically done, we have analysed the fitting errors by computing discrepancies between the empirical frequencies and the fitted ones. Since many empirical frequencies are zero (in the tail), the theoretical and empirical cdfs have been considered. The metrics employed are the mean absolute error (MAE) and the mean relative absolute error (MRAE). More precisely, given a sample $\mathbf{x} = (x_1, \dots, x_n)$, we let $\mathrm{MAE}(\mathbf{x}) = \sum_{j=1}^{n} \frac{1}{n} |F_n(x_j) - F(x_j)|$ and $\mathrm{MRAE}(\mathbf{x}) = \sum_{j=1}^{n} \frac{1}{n} |F_n(x_j) - F(x_j)| / F_n(x_j)$ with F_n the empirical cdf and F the fitted cdf.

3.2.1 Interarrival Times

Some statistical characteristics usually observed in it samples are the following: very high variance and skewness, relatively high frequency associated to the observation $it = 1$, monotonically decreasing frequencies with a slowly decaying tail and a drop in the passage from the frequency at $it = 1$ to the one at $it = 2$. An example of empirical frequencies computed from a given sample **it** is shown in Figure 3.2.1. It clearly shows the dominance of the frequencies corresponding to $it = 1$ and $it = 2$, which are particularly meaningful if one is aiming to use the interarrival times as a proxy for the rainfall occurrence. Due to the multitude of rather smaller values in the rightmost part of the histogram, to perform comparisons and enhance visibility, a log-log scale for all the plots showing fitting results has been adopted in this section. In order to explore the flexibility of the PSHZLD, we have compared the fitted PSHLZD with the fitted HLZD and the fitted PSHLZD with the fitted OIHLZD. To summarise the results, we have selected 4 of the 33 available samples since they have been considered particularly meaningful with respect to the whole dataset. The selected samples were STW (Y), TRA (Y), TRA (S1) and TRA (S2). The decision to show here such a relatively small partition of the dataset is justified by the fact that a deeper analysis of the fittings of the HLZD shall be given later in Section 3.3.5.

Fig. 3.2.2 shows plots of the fitted PSHLZ pmf (solid line) and HLZ pmf (dashed line) compared with the empirical frequencies (red dots) for the 4 selected samples. The fitting in both cases is very good. In particular, in the illustrated case of TRA (S1), the PSHLZD succeeds in fitting the drop from $it = 1$ to $it = 2$ whereas the HLZD fails. This seems to usually happen in the drier periods, where this drop is more prominent.

Figure 3.2.3 presents the MAE and MRAE obtained by comparing the fitted cdf's of the PSHLZD (circle) and the HLZD (triangle) with the empirical cdf's computed from the samples **it**. Note that the MAE and the MRAE are generally lower for the PSHLZD. Finally, we noticed that in seemingly all cases, the fitted OIHLZD and PSHLZD have minimal differences and are almost indistinguishable (see Figure 3.2.4 for an example), confirming the great flexibility of the latter distribution.

To conclude the validation analysis, we compared sample means and sample variances with the same theoretical moments of the HLZD and the PSHLZD computed in Section 2.3 and 2.4 respectively. In Table 3.2.1, we show the results for the 4 selected samples. In all cases, the fitted distributions' means agree with the sample ones. For the variances, the

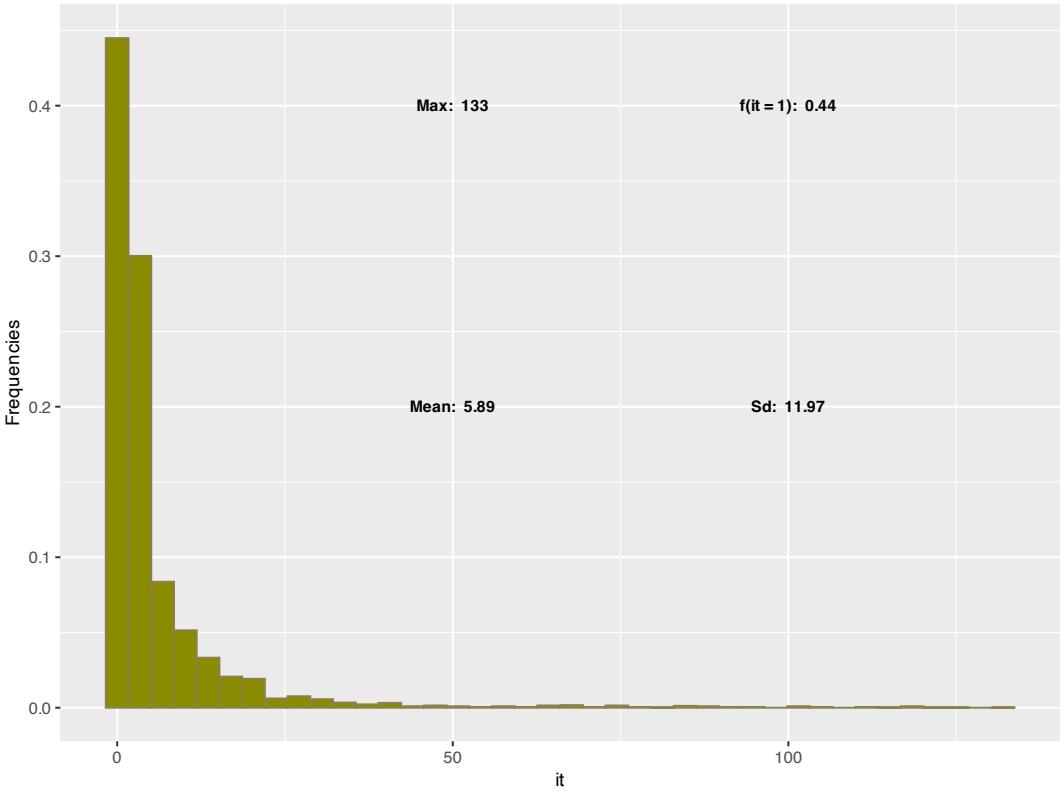

Figure 3.2.1: Histogram computed using the sample **it** from TRA (Y). The range is up to 133. The mode is $it = 1$ with relative frequency 0.44. The mean and the standard deviation are 5.89 and 11.97 respectively.

same happens with a slightly better performance of the PSHLZD, with an exception of TRA (S1). As already said, a more in depth discussion of the fitting results of the HLZ distribution alone shall be given later in Section 3.4.

3.2.2 Rainfall Depths

We now summarize the fitting of the rainfall depths samples $\mathbf{h}^*$ using both the PSHLZD and the HLZD. Recall that given the numerous ties in the data we have decided to treat h as a discrete random variable. In practice, this means that prior to fitting we have taken the integer part of each observation in a series $\mathbf{h}^*$ and then performed the estimation on the resulting discrete series. Consequently, as we continue, the rainfall depths, whenever mentioned either in the form of h or as a series $\mathbf{h}^*$, shall be tacitly understood as the integer part of the original measurements.

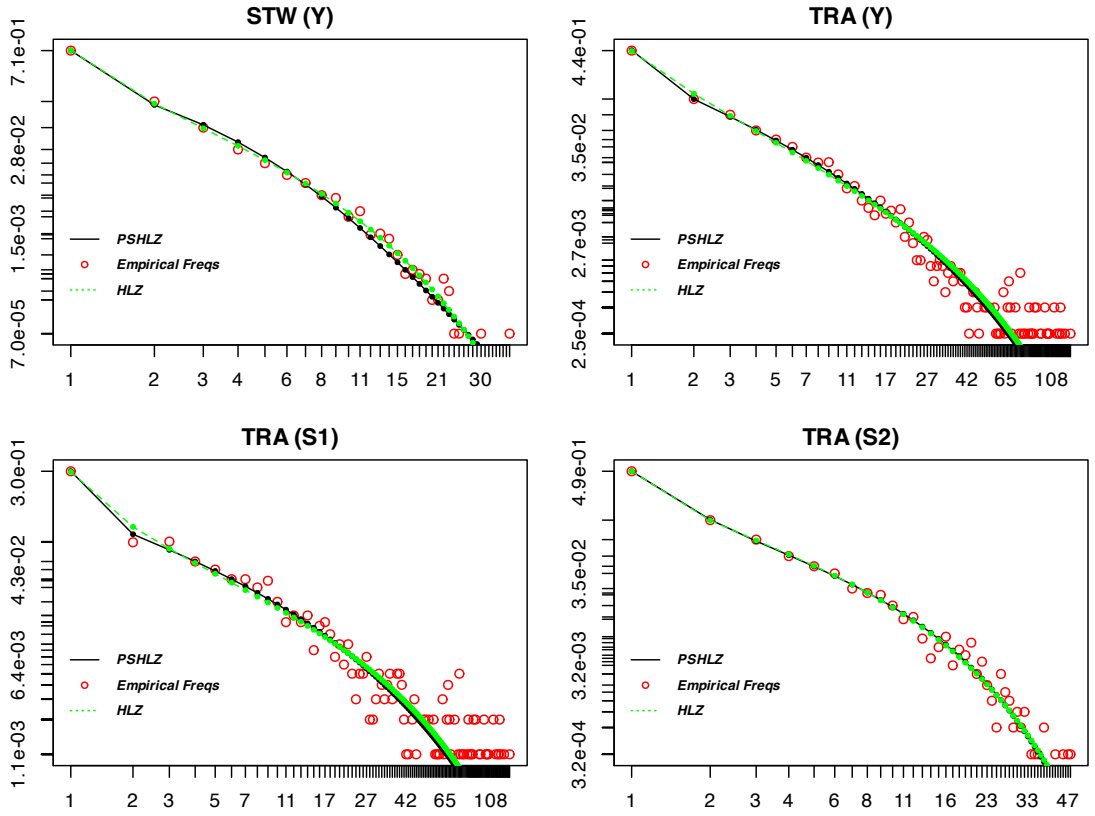

Figure 3.2.2: Log-log plots of the fitted HLZD pmf (green dashed line), the fitted PSHLZ pmf (black solid line) and the empirical frequencies (red dots) for the 4 selected *it* samples STW (Y), TRA (Y), TRA (S1) and TRA (S2).

We decided to omit any comparison with the OIHLZD since its performance is analogous to the one showed for the interarrival times and therefore did not add any significant insights. Fig. 3.2.5 shows a histogram computed from a sample of rainfall depths whose shape appears similar to the one previously shown from a sample **it**. Indeed, we can see the mode in $h=1$, accompanied by a multitude of rather smaller values in the slowly decaying tails. Similarly as before, this prompted us to employ in this section a log-log scale for the plots reporting fitting results. Unlike in the previous section, fitting results will be shown for all the stations and all the subsets, since the fitting of the rainfall depths shall not be investigated further in the second part of this chapter.

Observing the fitting results in Figure's 3.2.6, 3.2.7 and 3.2.8 leads to the following remarks. As with the interarrival times, the fitting appears to be satisfactory, seemingly even better. Additionally, due to the absence of the

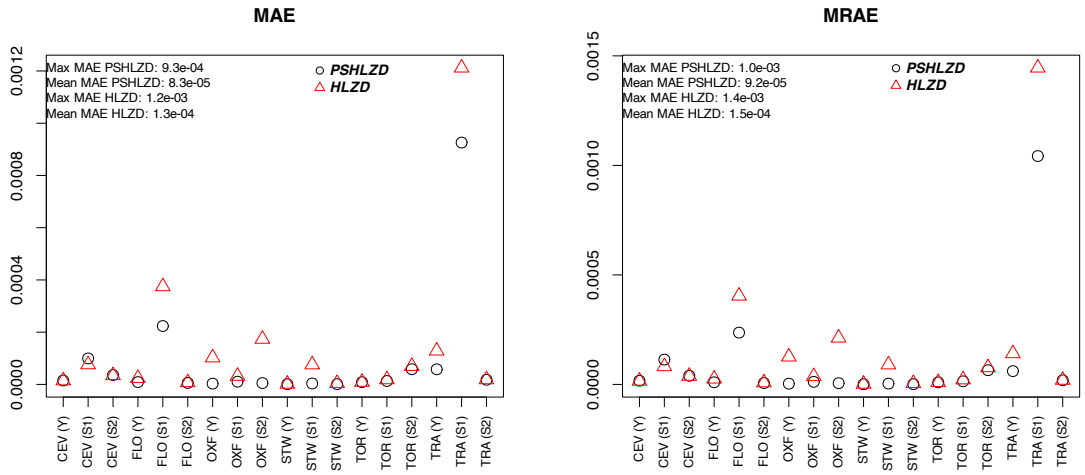

Figure 3.2.3: Dot plots of MAE and MRAE taking as reference the fitted cdf of the PSHLZD (black circle) and of the HLZD (red triangle) for all the samples **it**. The maximum MAE as well as the mean MAE are given in the top left for both the fitting distributions.

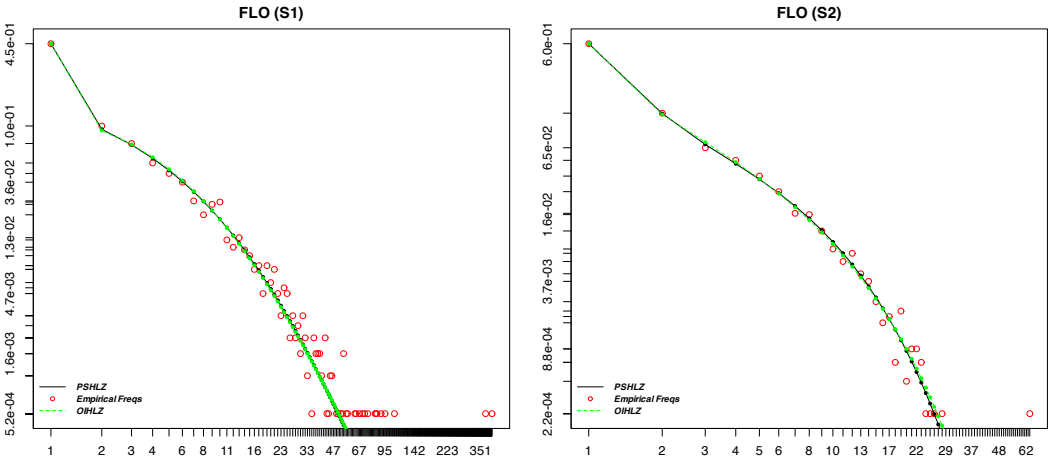

Figure 3.2.4: The fitted PSHLZ pmf (black solid line) is plotted together with the fitted HLZ pmf (green dashed line) and the empirical frequencies (red dots) for the sample **it** of FLO (S1) and FLO (S2).

prominent drop between the first and second frequencies which characterised *it* data, a less marked difference between the performance of the two distributions may be initially inferred. However, by inspecting Fig. 3.2.9, which presents the MAE and the MRAE obtained by comparing the fitted cdf's of the PSHLZD (circle) and the HLZD (triangle) with the empirical cdf's

Table 3.2.1: The sample means and the sample variances for the 4 selected samples are given in the first column. The means and the variances of the fitted HLZD and of the fitted PSHLZD are given in the second and in the third column respectively.

	STW (Y)		
	Sample	HLZ	PSHLZ
Mean	1.822	1.826	1.822
Var	4.374	4.475	4.435
	TRA (Y)		
	Sample	HLZ	PSHLZ
Mean	5.887	5.887	5.887
Var	143.31	125.36	159.14
	TRA (S1)		
	Sample	HLZ	PSHLZ
Mean	13.039	13.039	13.039
Var	469.13	464.87	540.14
	TRA (S2)		
	Sample	HLZD	PSHLZD
Mean	3.792	3.792	3.792
Var	26.7	27.6	27.29

computed from the corresponding $\mathbf{h}^*$, one can see that even in this case the PSHLZD provides an improvement. We end this section by stressing out that the great fitting results showed by both the PSHLZ and the HLZD on samples of h are a proof of the great flexibility enjoyed by both distributions. Indeed, the suitability of these distributions in modelling daily rainfall depths is demonstrated even with the significantly varying rainfall patterns across the stations examined. Importantly, the observed frequencies for h are effectively replicated on the sole annual basis, suggesting that dividing the datasets into sub-periods is not essential for accurately modelling the probabilistic behaviour of h within any of the rainfall regimes studied.

3.3 DM and IM Results

We shall now proceed with the application and comparison of the DM and IM on the available dataset. It must be noted that, only the HLZ family of discrete distributions will be used in the following, even though the previous

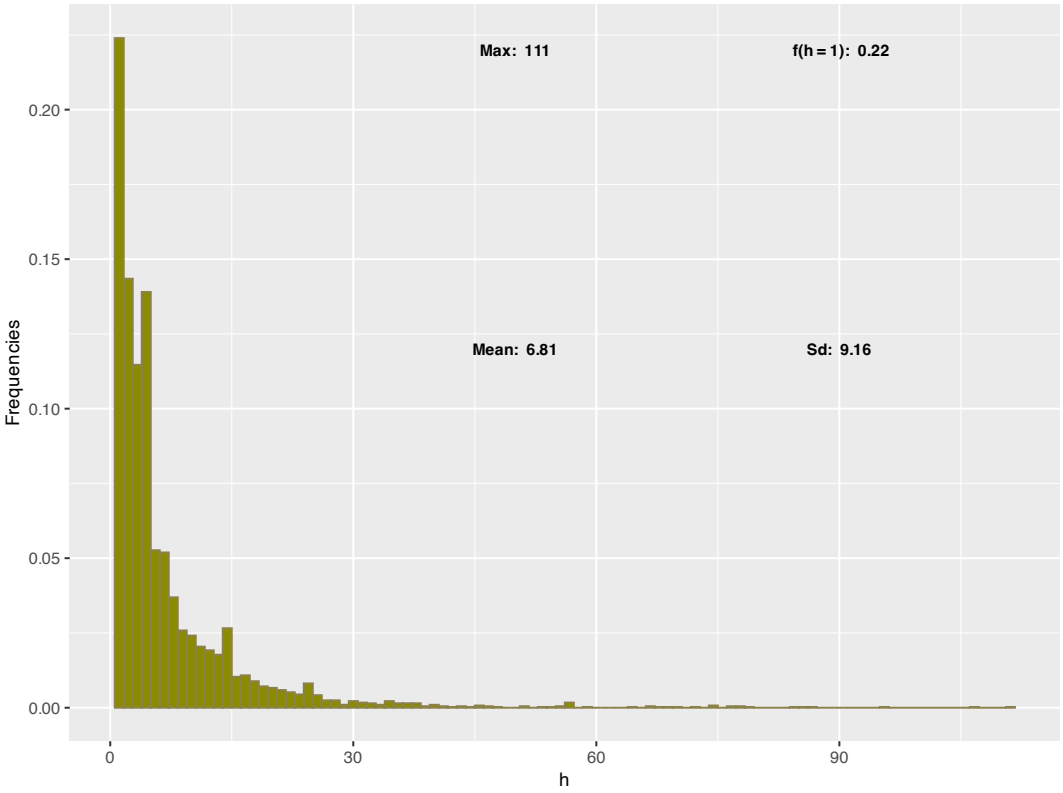

Figure 3.2.5: Histogram computed using the sample $\mathbf{h}^*$ for TRA(Y). 111 is the maximum registered depth. The mode is $h=1$ with relative frequency 0.22. The mean and the standard deviation are 6.81 and 9.16 respectively.

sections hint to a superior fitting of the PSHZLD on some of the available samples. However, the latter has been excluded firstly for a matter of computational complexity, which is slightly increased by the additional effort put into optimising for the added parameter. Secondly, the HLZ family is a nested family and, unlike in the preceding sections where only the full HLZD has been considered, we here shall also use the likelihood ratio as a model selection tool, which rely on the underlying family being nested. Recall that, since a significant dependence between the rainfall depths and the corresponding wet spell has been detected, we will not discuss here the derivation of the distribution of the volume of rainfall in a sequence of rainy days. In the next subsection the methodology is explained more in detail together with the inference procedures employed.

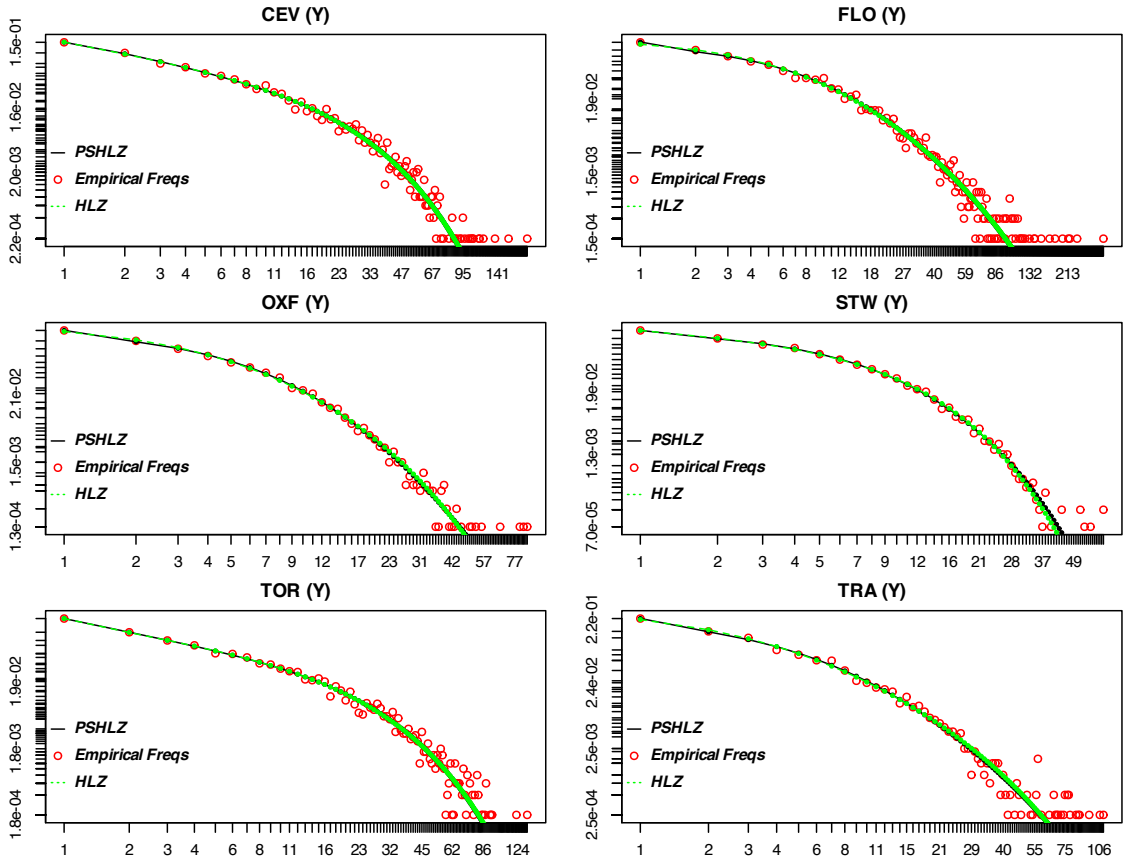

Figure 3.2.6: Log-log plots of the fitted HLZ pmf (green dashed line), the fitted PSHLZ pmf (black solid line) and the h empirical frequencies (red dots) for the subset Y for all the stations.

3.3.1 Direct Method

The DM is structured in the following way. As explained above, it is reasonable to suppose that it follows a 3-parameter HLZD. Given the estimated pmf p_{it} and cdf F_{it} of it, the relationships described in Section 2.6.1 and the discussion therein let us derive the pmfs of ws, ds, wch and dch in the following way. For $k \in \mathbb{N}^+$ we can write

$$p_{ws}(k) = (1 - p_{it}(1))[p_{it}(1)]^{k-1}, \tag{3.1}$$

$$p_{ds}(k) = \frac{p_{it}(k+1)}{1 - p_{it}(1)}, \tag{3.2}$$

$$p_{wch}(k) = \sum_{i=0}^{k-1} \binom{k-1}{i} [p_{it}(2)]^i [p_{it}(1)]^{k-1-i} [1 - F_{it}(2)], \tag{3.3}$$

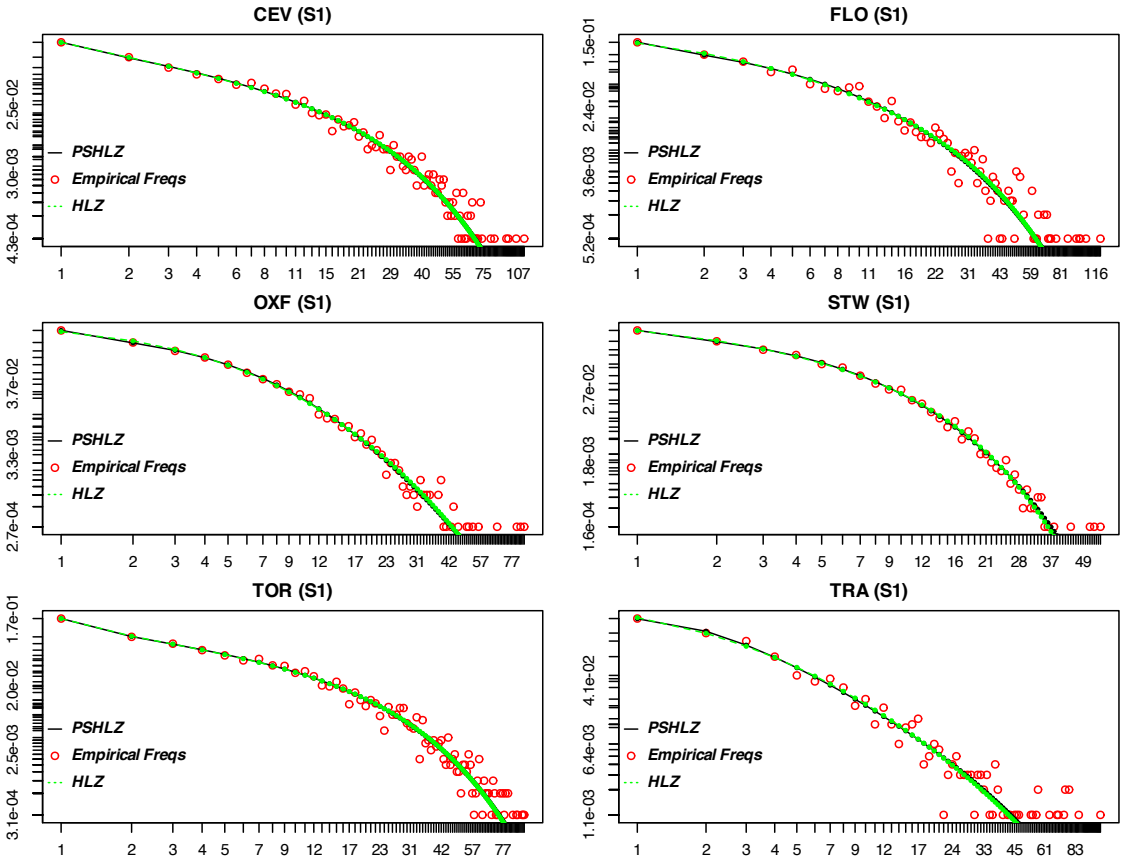

Figure 3.2.7: Log-log plots of the fitted HLZ pmf (green dashed line), the fitted PSHLZ pmf (black solid line) and the h empirical frequencies (red dots) for the subset S1 for all the stations.

$$p_{dch}(k) = \sum_{\pi \vdash m} \frac{l(\pi)!}{r_1! \cdots r_m!} \prod_{\pi' \in \pi} p_{it}(\pi'+1)[1-F_{it}(1)]^{l(\pi)-1} p_{it}(1). \tag{3.4}$$

The estimation of the parameters (θ, s, a) in p_{it} has already been carried out in the first part of this chapter according to Section 2.3.6.

3.3.2 Indirect Method

In the IM, in order to explore the validity of the other distributions in the HLZ family for ws and ds, in addition to the full 3-parameter HLZD we tested another three members. For $v \in \{ws, ds\}$ we have considered

1. The Geometric distribution (ID = 4 in Table 2.1), corresponding to (2.19) with $s = -1$ and thus with functional form

$$p_v(k;\theta) = \theta(1-\theta)^{k-1}, \quad k \in \mathbb{N}^+. \tag{3.5}$$

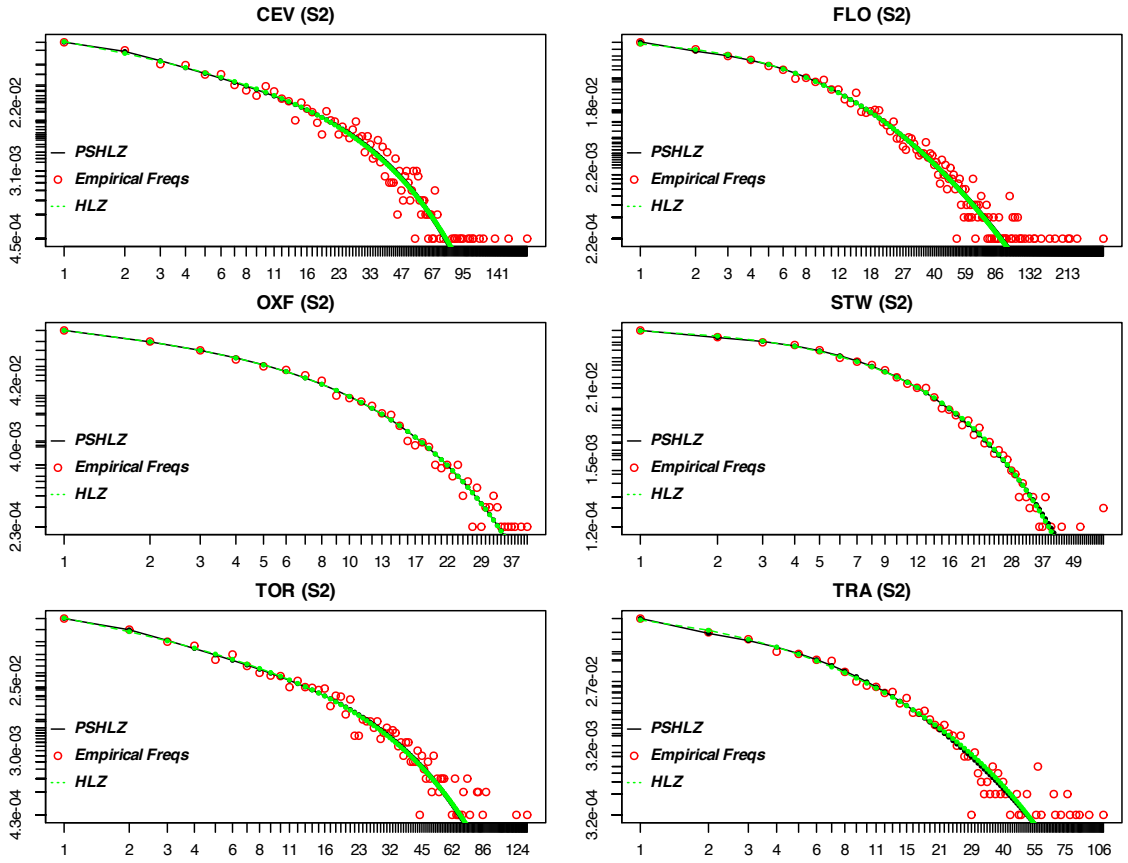

Figure 3.2.8: Log-log plots of the fitted HLZ pmf (green dashed line), the fitted PSHLZ pmf (black solid line) and the h empirical frequencies (red dots) for the subset S2 for all the stations.

2. The Polylog distribution (ID = 2 in Table 2.1), corresponding to (2.19) with $a=0$ and thus with functional form

$$p_v(k;\theta,s)=\frac{\theta^{k-1}}{k^{s+1}\mathrm{Li}_{s+1}(\theta)},\quad k\in\mathbb{N}^+, \tag{3.6}$$

where $\mathrm{Li}_s(\theta)=\Phi(\theta,s+1,1),\quad k\in\mathbb{N}^+$.

3. The 2-parameter Logarithmic distribution (ID = 5 in Table 2.1), corresponding to (2.19) with $s=0$ and thus with functional form

$$p_v(k;\theta,a)=\frac{\theta^{k-1}}{(k+a)\Phi(\theta,1,a+1)},\quad k\in\mathbb{N}^+. \tag{3.7}$$

Then, given the estimated p_{ws} and p_{ds}, the relationships described in Section 2.6.1 and the discussion therein let us derive the pmfs of *it*, *wch* and *dch* in the following way. For $k\in\mathbb{N}^+$ we can write

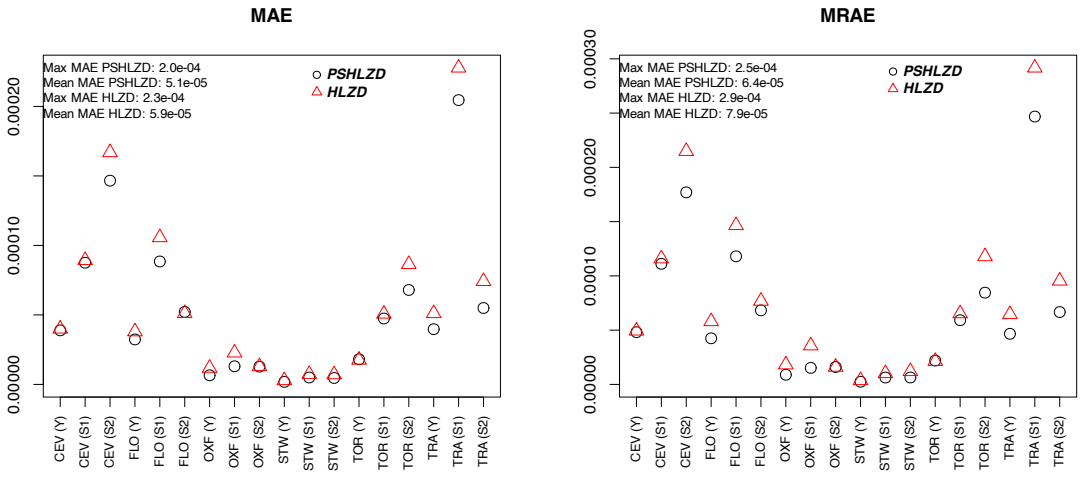

Figure 3.2.9: Dot plots of MAE and MRAE taking as reference the cdf of the PSHLZD (black circle) and of the HLZD (red triangle) for all the h samples. The maximum MAE as well as the mean MAE are given in the top left for both the fitting distributions.

$$p_{it}(k) = \begin{cases} \frac{\mathbb{E}[ws]-1}{\mathbb{E}[ws]} & k=1, \\ p_{ds}(k-1)[1-p_{it}(1)] & k>1, \end{cases} \tag{3.8}$$

$$p_{wch}(k) = \sum_{i=1}^{k} p_{ws}^{i*}(k)[p_{ds}(1)]^{i-1}[1-F_{ds}(1)], \tag{3.9}$$

$$p_{dch}(k) = \sum_{i=1}^{k} p_{ds}^{i*}(k)[p_{ws}(1)]^{i-1}[1-F_{ws}(1)]. \tag{3.10}$$

The selection of a model between the three proposed above for p_{ws} and p_{ds} and the related parameter estimation is described in the following section.

3.3.3 Inference

When selecting a model for ws and ds in the IM, the LLR test as detailed in Section 2.2.6 has been employed within the following model selection

procedure. Let the following reduced parameter spaces

$$\begin{aligned}\Theta_{lerch} &= (0,1)\times(-\infty,+\infty)\times(-1,\infty),\\ \Theta_{polylog} &= (0,1)\times(-\infty,+\infty)\times\{0\},\\ \Theta_{2log} &= (0,1)\times\{0\}\times(-1,\infty),\\ \Theta_{geom} &= (0,1)\times\{-1\}\times\{0\},\end{aligned}$$

corresponding respectively to the distributions with $ID = 1,2,5,4$ in Table 2.1, and fix a series of observations $\mathbf{y}$ of ws or ds. Algorithm 1 can be used to select a distribution between $ID = 1,2,5,4$ in Table 2.1 as a statistical model for the current data $\mathbf{y}$.

Algorithm 1: LL Ratio Algorithm

Initialization

parameter.spaces $\leftarrow \{\Theta_{\text{lerch}}, \Theta_{\text{polylog}}, \Theta_{\text{2log}}, \Theta_{\text{geom}}\}$
LLR $\leftarrow$ [ ]
parameters $\leftarrow$ [ ]

for $i \leftarrow 1$ ***to*** 4 **do**
  $\Theta_0 \leftarrow$ parameter.spaces$[i]$
  $\hat{\boldsymbol{\theta}} \leftarrow \arg\max_{\boldsymbol{\theta}\in\Theta_0} \ell_n(\boldsymbol{\theta}, \boldsymbol{y})$
  parameters$[i] \leftarrow \hat{\boldsymbol{\theta}}$
  LLR$[i] \leftarrow \ell_n(\hat{\boldsymbol{\theta}}, \boldsymbol{y})$
end

Set best $\leftarrow$ lerch

if $-2(LLR[2] - LLR[1]) \leq 3.841$ **then**
  best $\leftarrow$ polylog
end

if $-2(LLR[3] - LLR[1]) \leq 3.841$ ***and*** $LLR[3] \geq LLR[2]$ **then**
  best $\leftarrow$ 2log
end

if $-2(LLR[4] - LLR[1]) \leq 5.991$ ***and***
$2(LLR[4] - \max(LLR[2], LLR[3])) \geq 3.841$ **then**
  best $\leftarrow$ geom
end

return parameters[best], best

The critical values appearing in the second part of Algorithm 1 are the ones corresponding to the χ^2 distribution at a 0.05 level with degrees of freedom equal to the difference between the number of free parameters of the alternative and the null models, as explained in Section 2.2.6. Moreover, to assess the adequacy of the selected member of the Lerch family in reproducing the observed frequencies, we have employed a simulated χ^2 goodness-of-fit test as described in the beginning of Section 3.2.

Before proceeding with the actual fitting results involved in the DM and IM, in the next subsection some additional preliminary statistical analyses involving the other temporal variables are performed and discussed.

3.3.4 Non-Parametric Analyses

The primary statistics for the considered rainfall time variables - it, ws, ds, wch, and dch - are summarized in the box plots in Fig. 3.3.1. These plots present all stations across the three periods Y, S1 and S2. The figure reveals the varied statistical properties of the examined variables and underscores the impact of seasonality on all locations.

It is noteworthy that the STW station exhibits the highest ws and wch statistics, likely attributable to its high frequency of rainfall. These elevated values are naturally offset by the lowest it, ds, and dch statistics across all considered periods. For ws during the S1 period, most stations display a relatively limited range of variability, except for STW and TRA, which exhibit significantly different rainfall regimes.

An in-depth analysis of the relationship between spells and chains can be pursued by observing Fig. 3.3.2. This figure plots the ratios of the observed cumulative frequencies $F_{n,ws}/F_{n,wch}$ (panels a and b) and $F_{n,ds}/F_{n,dch}$ (panels c and d) against the corresponding time variables for the six stations during the two seasons, S1 (panels a and c) and S2 (panels b and d). These ratios provide insights into the relative significance of the derived variables dch and wch compared to ds and ws.

As expected, the ratios $F_{n,ws}/F_{n,wch}$ and $F_{n,ds}/F_{n,dch}$ are both greater than one, with $F_{n,ds}/F_{n,dch}$ ratios typically ranging from 1 to 3, and $F_{n,ws}/F_{n,wch}$ ratios between 1 and 1.7. In general, there is an observable trend from S1 to S2, where $F_{n,ws}/F_{n,wch}$ increases and $F_{n,ds}/F_{n,dch}$ decreases.

The ratios for the CEV and TOR stations remain relatively consistent across both seasons, indicating limited seasonality. On the other hand, STW, which experiences very frequent rainfall, is unique in having a higher $F_{n,ws}/F_{n,wch}$ ratio compared to $F_{n,ds}/F_{n,dch}$.

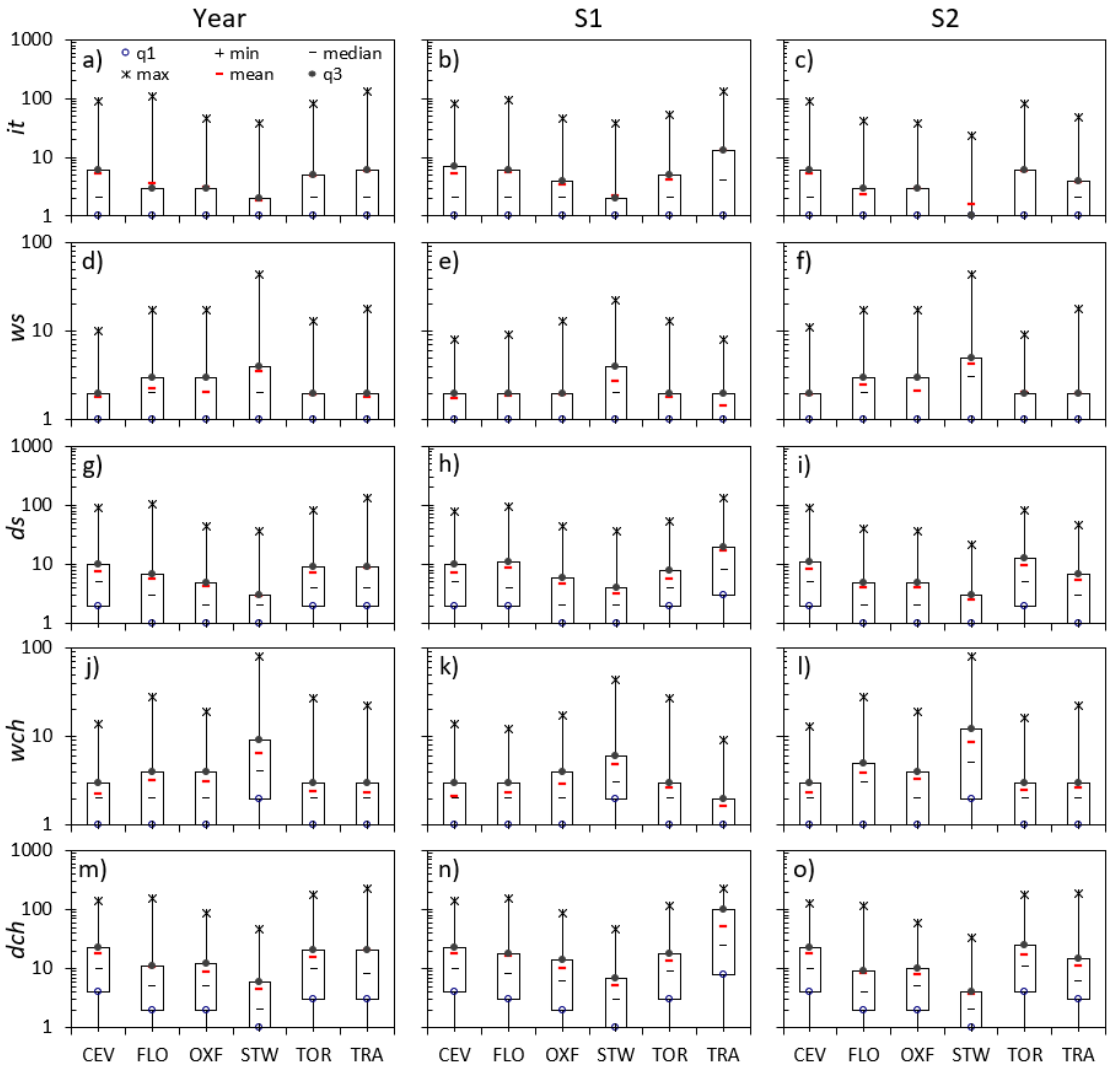

Figure 3.3.1: Box plots of the time variables, it (a – c), ws(d – f), ds(g – i), wch (j – l) and dch (m – o) for all six stations and for the three periods, Y (left column), S1 (central column) and S2 (right column). The variables q1 and q3 identify the first and the third quartiles, respectively.

It is particularly noteworthy that TRA shows very high $F_{n,ds}/F_{n,dch}$ ratios during the S1 period, reflecting the region's significant aridity in western Sicily during this time. Additionally, FLO exhibits substantial differences in $F_{n,ws}/F_{n,wch}$ between the two seasons, further underscoring the seasonal variability in rainfall patterns.

Recall that a key difference between the DM and the IM, as introduced in Section 2.7, lies in modelling the inter-event times as i.i.d. renewal times, which results in a geometric distribution of ws in the DM. This is equivalent

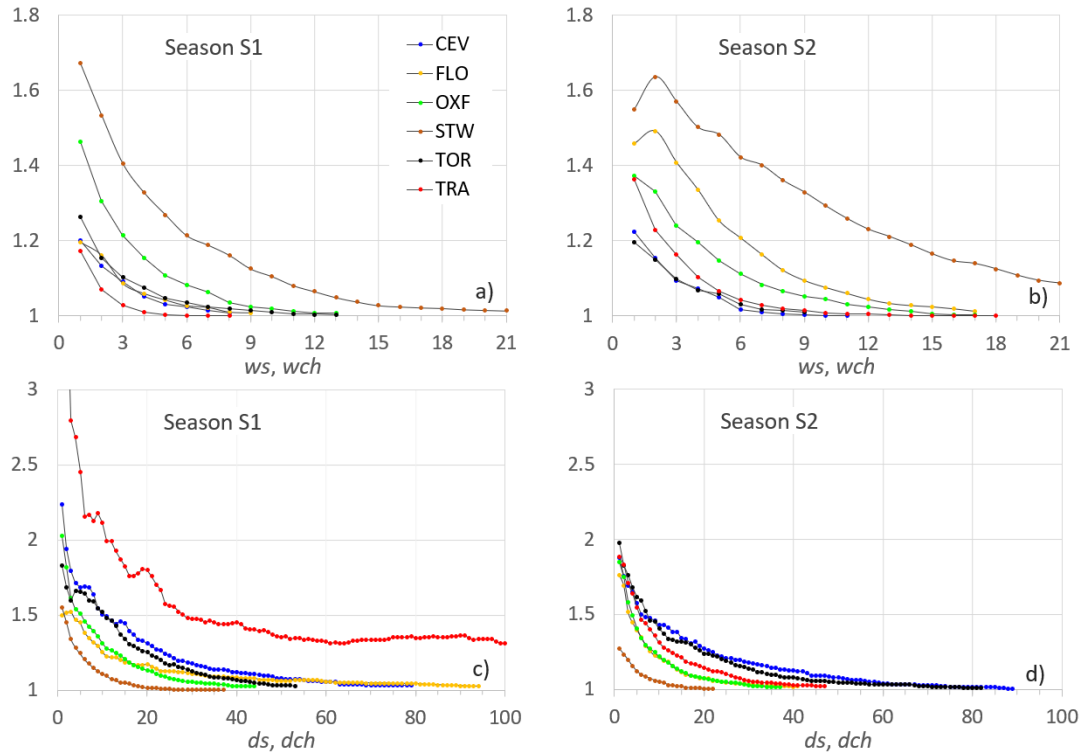

Figure 3.3.2: Ratios between observed cumulative frequencies $F_{n,ws}/F_{n,wch}$ (panels a,b) and $F_{n,ds}/F_{n,dch}$ (panels c,d) for the six stations, and for S1 (panels a,c) and S2 (panels b,d) for all the stations considered.

to making a simple hypothesis on the daily occurrence of rain: supposing that the probability of a rainy day is a constant independent from the weather of the previous days (see, for example, the classic Chatfield, 1966). Formally speaking, this means that the distribution of the ws enjoys the so called lack of memory property. Indeed, in the discrete setting, the only distribution with this characteristic is the geometric distribution. More in detail, we have that ws has a geometric distribution if and only if

$$\mathbb{P}(ws > i+j | ws > i) = \mathbb{P}(ws > j), \quad i,j \in \mathbb{N}. \tag{3.11}$$

Hence, a preliminary analysis of the memory (or lack thereof) property would provide a useful tool for recommending the use of the IM or the DM. Let us consider a time series of observations of ws denoted by **ws**. Let k in $\{\min_{x\in\mathbf{ws}},\ldots,\max_{x\in\mathbf{ws}}-1\}$ and denote with S_k the number of wet spells of length at least k found in **ws**. A way of testing the memoryless property is to study the behaviour of the ratios $\frac{S_{k+1}}{S_k}$ for all $k\in\{\min_{x\in\mathbf{ws}},\ldots,\max_{x\in\mathbf{ws}}-1\}$. The

corresponding theoretical value is

$$\frac{\mathbb{P}(ws > k+1)}{\mathbb{P}(ws > k)} = 1 - r_{ws}(k), \tag{3.12}$$

where r_{ws} is the failure rate as given in (2.26). If $ws \sim HLZ(\theta, a, s)$ as in (2.19), recall that from Proposition 2.3.7 we have

$$r_{ws}(k) = \frac{\theta}{(a+k)^{s+1}\mathcal{T}(\theta, s, a+k-1)}, \quad k \in \mathbb{N}^+. \tag{3.13}$$

Is we set $s = -1$ in the above expression, which corresponds to the $Geom(1-\theta)$ in the family of HLZDs, we get $r_{ws}(k) = \theta$ for any $k \in \mathbb{N}^+$ and thus $\frac{\mathbb{P}(X>k+1)}{\mathbb{P}(X>k)} = 1-\theta$ for any $k \in \mathbb{N}^+$, which is constant as expected. For all the other distributions in the HLZ family r_{ws} is a monotone decreasing function of k from Proposition 2.3.7 when $s > -1$. Thus, the proposed procedure consists in computing these ratios for all $k \in \{\min_{x \in \mathbf{ws}}, \ldots, \bar{k}-1\}$, where $\bar{k} < \max_{x \in \mathbf{ws}}$ is such that the number of wet spells of length at least $\bar{k}$ is still significant, and checking if they approximately constant.

Hence, to assess the memoryless property of ws in our data, Fig. 3.3.3 displays the sequence of the ratio $\frac{S_{k+1}}{S_k}$ for all stations and periods. The series are shown up to ws values with at least 10 observations. The results indicate a roughly constant value for CEV and TOR (Fig. 3.3.3a), a slightly increasing trend for TRA and FLO (Fig. 3.3.3b), and a pronounced increasing trend for OXF (Fig. 3.3.3c) and STW (Fig. 3.3.3d). These patterns suggest that a geometric distribution may not be suitable for modelling the ws records of OXF and STW. The increasing variability in the $\frac{S_{k+1}}{S_k}$ ratios serves as an indicator that the geometric distribution might not adequately capture the characteristics of wet spells for these stations.

3.3.5 DM and IM Comparison

For the series of observations of the six rain gauges, the Lerch family as defined in (2.19) was fitted to the data across three periods, that is the entire year Y, and the two seasonal periods S1 and S2, following both the DM and the IM. For the former, the fitting was already conducted for *it* in the previous sections when considering the full HLZD and shall be further discussed. In the latter, for both ws and ds, a model within the HLZ family was selected by following the procedure in Algorithm 1. The parameters estimated using MLE are presented in Table 3.3.1. It is crucial to emphasize

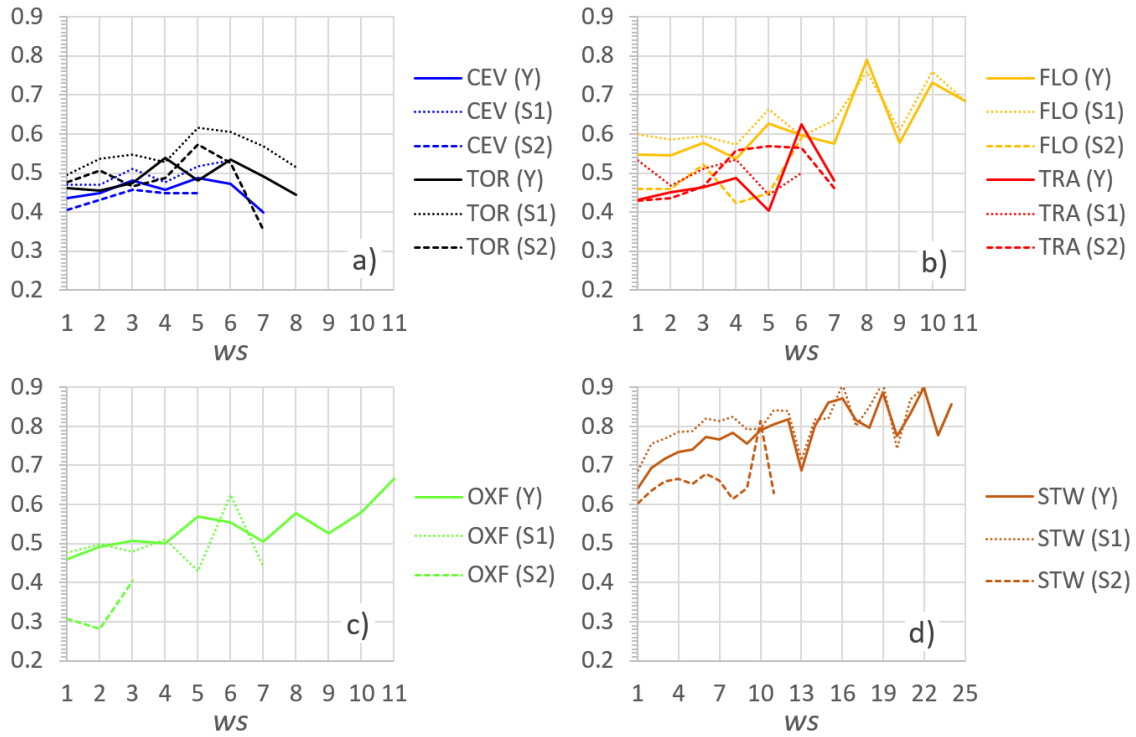

Figure 3.3.3: For the three periods, Y, S1 and S2, plots of the S_{k+1}/S_k ratios versus ws in a) for CEV and TOR, in b) for FLO and TRA, in c) for OXF and in d) for STW.

that for each station and period combination, the parameters listed in Table 3.3.1 align with the special cases of the Lerch family (as detailed in Table 2.1), according to the results of the model selection procedure mentioned above.

In the DM framework, the 3-parameter Lerch distribution is selected for all sites and periods for it. The θ parameter ranges between 0.86 and 0.97, with higher values observed for the entire year and season S1 for TRA. This does not come as a surprise since, as seen in Section 2.3.5, θ governs the heaviness of the tail and we had previously observed that interarrival times samples usually present a remarkable tail. The variation of θ across different periods is minimal for OXF, STW, and, to a lesser extent, CEV, indicating a low degree of seasonality compared to the other stations. As expected, the a values are all negative (Table 3.3.1), reflecting the observed decrease in frequency from $it=1$ to $it=2$. The s parameter is positive for all it fittings, effectively replicating the fact that the mode consistently occurs at $it=1$.

When adapting the IM framework, the geometric distribution (with $s=0$ and $a=1$) was selected in several instances (10 out of 18) to fit ws,

Table 3.3.1: Parameters of the Lerch family of probability distributions fitted on series **it** (DM), **ws** and **ds** (IM) for the six stations and for the three periods, Y, S1 and S2. Values of 0 or 1 for the parameters s and a identify the special cases listed in Table 2.1.

Station	**Variable**	θ			s			a		
		Y	**S1**	**S2**	**Y**	**S1**	**S2**	**Y**	**S1**	**S2**
CEV	it	0.913	0.902	0.923	-0.558	-0.605	-0.524	-0.953	-0.954	-0.958
FLO	it	0.934	0.942	0.858	0.005	-0.300	-0.262	-0.657	-0.833	-0.809
OXF	it	0.902	0.906	0.896	0.069	-0.004	0.122	-0.384	-0.437	-0.342
STW	it	0.867	0.864	0.861	0.348	0.186	0.539	-0.528	-0.489	-0.508
TOR	it	0.919	0.871	0.940	-0.419	-0.607	-0.451	-0.891	-0.942	-0.953
TRA	it	0.970	0.975	0.897	0.164	-0.256	-0.311	-0.364	-0.687	-0.771
CEV	ws	0.446	0.419	0.476	-1	-1	-1	1	1	1
FLO	ws	0.650	0.464	0.599	0	-1	0	3.084	1	1
OXF	ws	0.558	0.486	0.600	-0.739	-1	-0.618	0	1	0
STW	ws	0.843	0.696	0.853	0	-0.676	-0.415	0.921	0	0
TOR	ws	0.473	0.583	0.508	-1	0	-1	1	1.143	1
TRA	ws	0.553	0.308	0.486	0	-1	-1	1.896	1	1
CEV	ds	0.913	0.901	0.922	-0.567	-0.613	-0.533	0	0	0
FLO	ds	0.967	0.953	0.853	0.938	0.4	-0.338	2.889	1.399	0
OXF	ds	0.890	0.880	0.880	-0.173	-0.457	-0.165	0	-0.548	0
STW	ds	0.838	0.861	0.799	0	0	0	0	0	0
TOR	ds	0.918	0.871	0.940	-0.448	-0.616	-0.459	0	0	0
TRA	ds	0.989	0.980	0.892	1.042	0.14	-0.386	3.870	2.062	0

more frequently for the sub-seasons (8 out of 12) than for the entire year. The CEV station uniquely had the geometric distribution chosen for all three periods, whereas this distribution was never selected for STW. To accurately model the probabilistic behaviour of ws, the OXF and STW stations appear to require the polylog-series or the 2-parameter extended log-series distribution (refer to Table 2.1). Interestingly, for TRA and FLO, the geometric distribution is inappropriate for ws over the entire year but is suitable when the dataset is divided into two sub-seasons.

Regarding the ds distribution in the IM framework, two or three parameters are consistently required, except for STW, where the logarithmic distribution ($s=1$ and $a=0$) is selected for all three periods. The polylog-series distribution is selected in 10 out of 18 cases, highlighting its suitability for modelling the probability law of dry spells in various scenarios. Notably, for CEV and TOR, the θ and s values in the polylog distribution for ds closely match those in the Lerch distribution for it. Consequently, for these stations, the additional parameter a in the Lerch distribution effectively serves to accommodate the geometric distribution of ws.

The evaluation of the goodness-of-fit for the selected distributions (for *it* in the DM, and for *ws* and *ds* in the IM) is summarised in Fig. 3.3.4. For all stations and periods (Y, S1, S2), the computed p-values are categorized into four ranges: $(0-0.01)$, $(0.01-0.05)$, $(0.05-0.1)$, and $(0.1-1)$. The light green $(0.05-0.1)$ and dark green $(0.1-1)$ ranges indicate the acceptance range of the null hypothesis.

For a few of the 180 (6 stations $\times$3 periods $\times$5 variables $\times$2 methods) Monte Carlo procedures (19/180), the presence of outliers (high values with very low frequency) bore a strong hint to the necessity of a preliminary frequency smoothing. The latter was performed by uniformly distributing the frequency of the outlier over all values between the observed value and the latest observed non-null frequency.

These cases are indicated by black dots in Fig. 3.3.4. Overall, the figure indicates that the fitting of *it* is satisfactory across all cases for Y, with the exception of FLO, where the p-value slightly exceeds 0.05. Likewise, the results for *ds* are generally good, which is expected given the strong similarity between *it* and *ds* modelling in the DM framework (refer to Proposition 2.6.2). In contrast, the fits for *ws* and *wch* are less satisfactory in several instances, particularly for STW and TRA.

The data on the right side of Fig. 3.3.4 (i.e., IM) indicates a notable reduction in the number of unsatisfactory fits (represented by red and orange classes) when the IM is applied, especially for *ws* and *wch*. Furthermore, analysing the seasonal datasets shows even greater improvement in accurately identifying the probability law of the time variables.

Figures 3.3.5 and 3.3.6 illustrate the cumulative observed frequencies and the corresponding fitted Lerch family cdfs for Y when applying the DM and the IM, respectively. The comparison between the two methods confirms the overall improvement achieved by the IM in cases that were not well-fitted by the DM, such as *ws* and *wch* for STW (compare Fig. 3.3.5b and 3.3.5d with Fig. 3.3.6a and 3.3.6c). This improvement is further supported by the comparison of the p-value classifications (Fig. 3.3.4) for the same variables when the IM is used instead of the DM. This result highlights the fact that adopting the simpler model of a renewal process governed by i.i.d. interarrival times may not always be a valid choice when aiming to also derive the other temporal variables.

Since the previous results show that the primary distinction between the two methods lies in their ability to model *ws* and *wch*, and to a lesser extent *dch*, the plots in Fig. 3.3.7 illustrate the MAE for these variables as

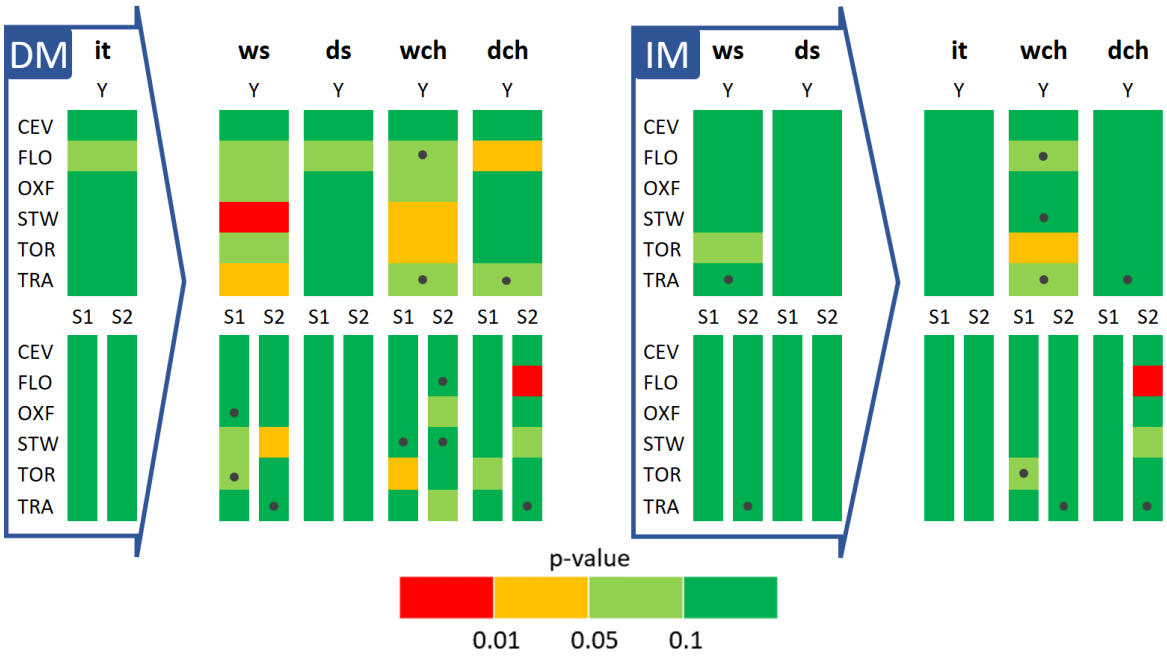

Figure 3.3.4: Summary of the results of the simulated χ^2 test for both the direct (DM, left-side figure) and the indirect (IM, right-side figure) methods. The variables inside the large arrows are the ones fitted in the corresponding method, whereas the other variables are deducted. The p-values (see legend) for all the stations and periods are reported. Black dots indicate that smoothing of observed frequencies was applied to calculate χ^2 ref.

modeled by the two methods for S1 (panels a, c, e) and S2 (panels b, d, f). These results further underscore the superior performance of the IM in most instances, as indicated by the high proportion of points situated below the 1:1 line (where differences are greater in the DM than in the IM).

Overall, the majority of differences cluster along the identity line, indicating that both methods generally perform well. However, substantial discrepancies are evident for STW, OXF, and FLO, which suggests that in these cases, relaxing the standard renewal property in favour of the alternating one is necessary for more accurate modelling.

The Lerch family distribution is also successful in modelling the probability of extreme values for the time variables. The overall consistency of these fittings is illustrated in Fig. 3.3.8, where the empirical 99th percentiles ($Q_{0.99}$) are compared with the estimated percentiles for all stations and periods, applying both the DM (Fig. 3.3.8a) and the IM (Fig. 3.3.8b). In Fig. 3.3.8a, the points generally align closely with the line of perfect agreement, with a few exceptions likely attributable to limited sample sizes, such as those seen at the Sicilian stations during season S1. Only a slight improvement is

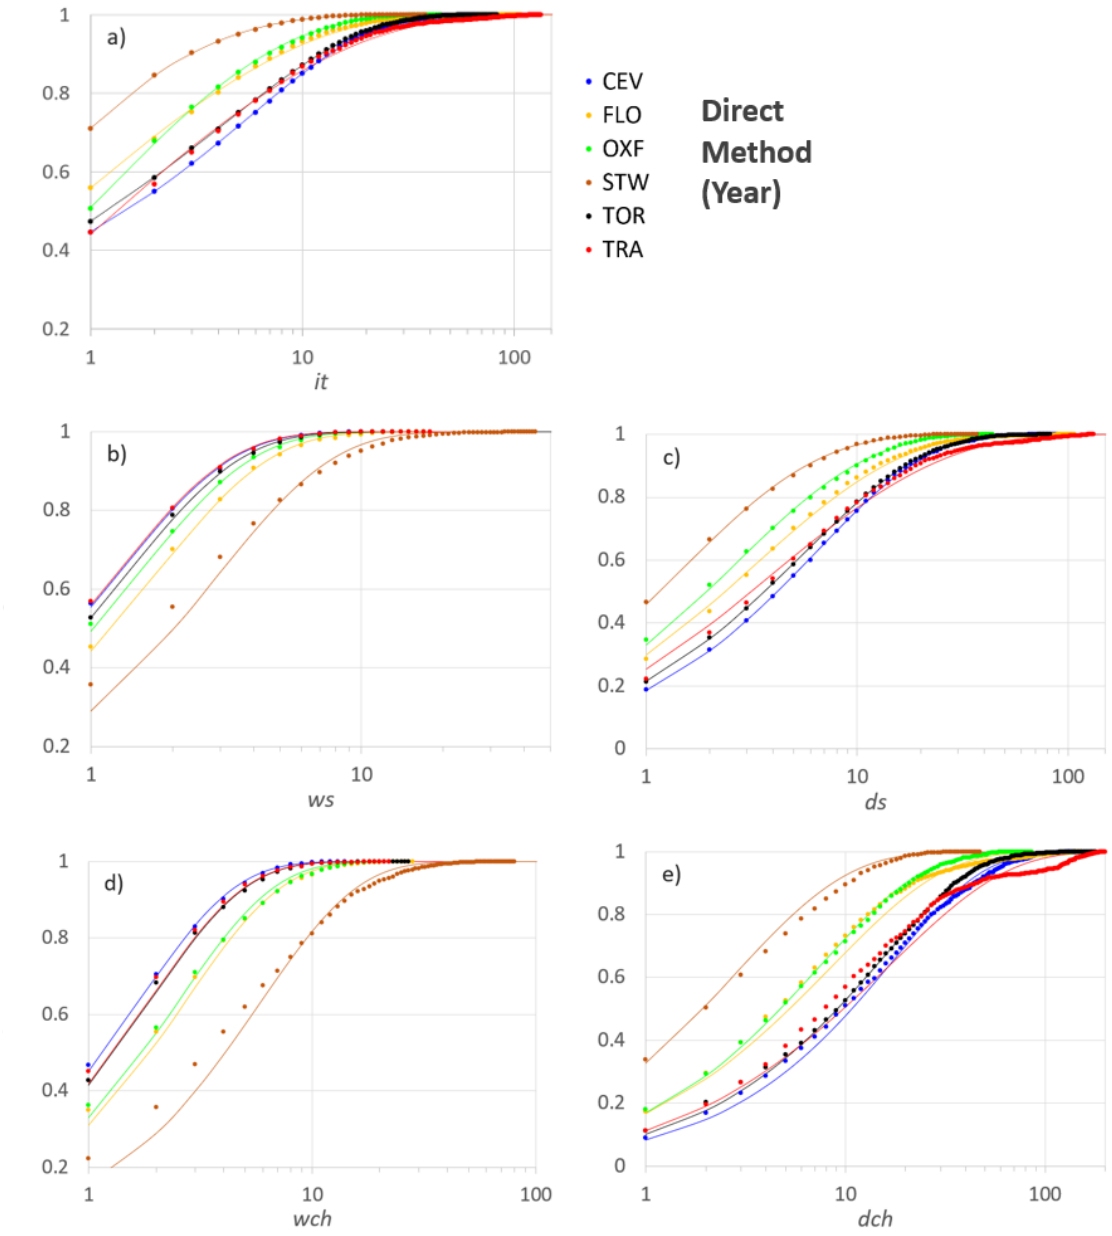

Figure 3.3.5: Observed cumulative frequencies (dots) and fitted cdfs (lines) for the six stations according to the direct method (DM) and for the period Y. The variables on the x-axis are in logarithmic scale.

noticeable when using the IM. These findings imply that the discrepancies observed with the DM in certain cases do not significantly affect the results for the extremes, but rather impact the accuracy for the more frequent data points.

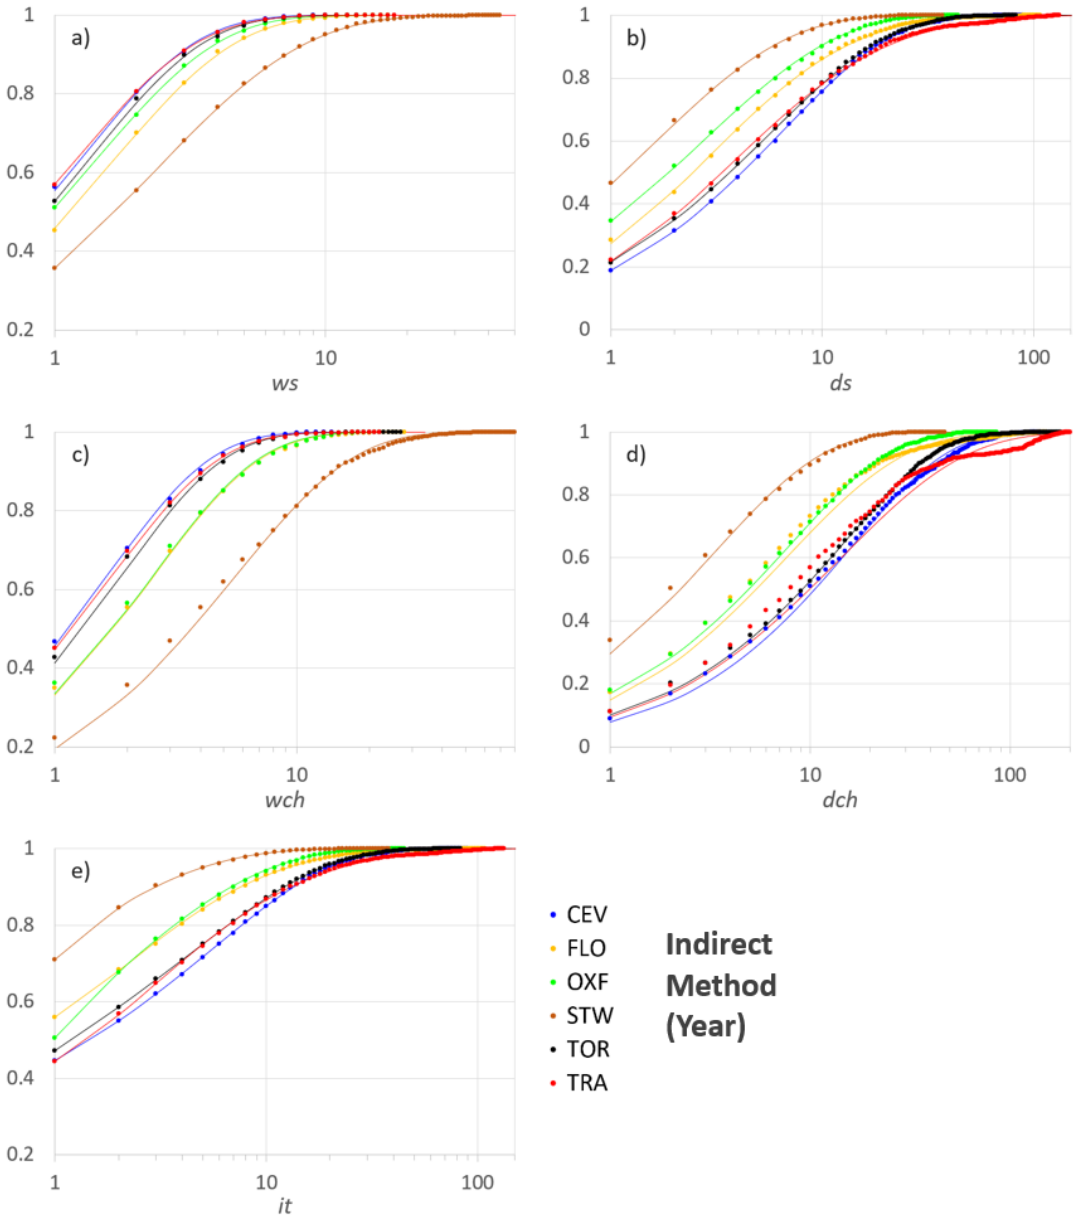

Figure 3.3.6: Observed cumulative frequencies (dots) and fitted cdfs (lines) for the six stations according to the indirect method (IM) and for the period Y. The variables on the x-axis are in logarithmic scale.

3.4 Discussion

As already mentioned, the application of the Lerch distribution to the selected six stations builds on previous studies conducted for stations in Sicily and Piedmont (Agnese et al., 2014; Berro et al., 2019). The adequacy

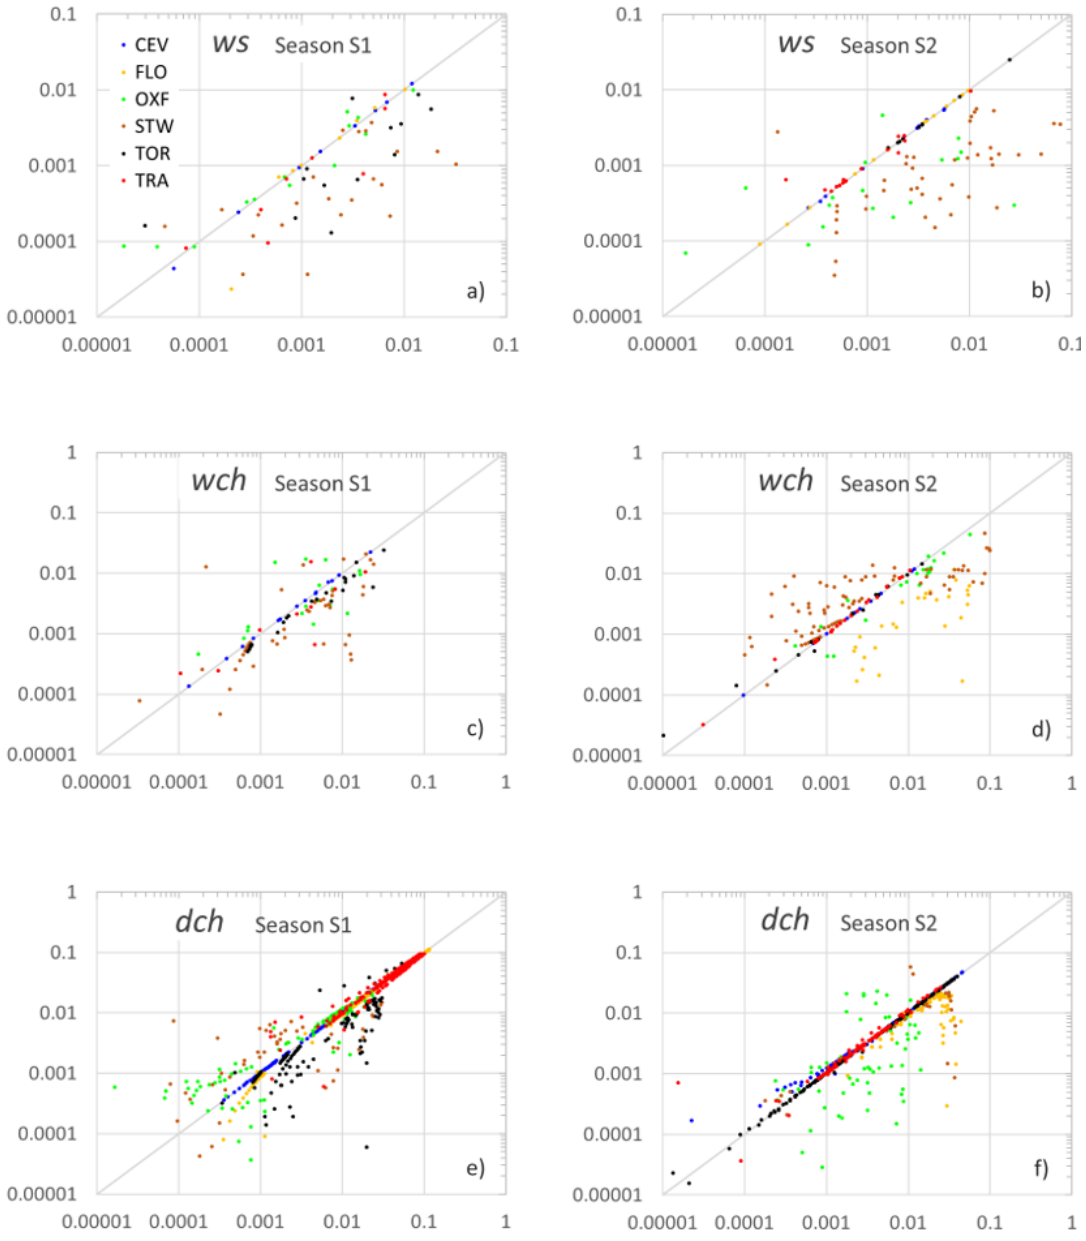

Figure 3.3.7: Scatterplots of the MAE with the DM (x-axis) against the MAE with the IM (y-axis) for ws(a,b), wch (c, d), and dch(e,f) for the six stations.

of this distribution in fitting daily rainfall inter-event times it is confirmed even when considering data recorded at OXF and STW, despite the markedly different rainfall patterns of these latter stations compared to those previously analysed. Notably, as was also previously noted for the rainfall depths, the

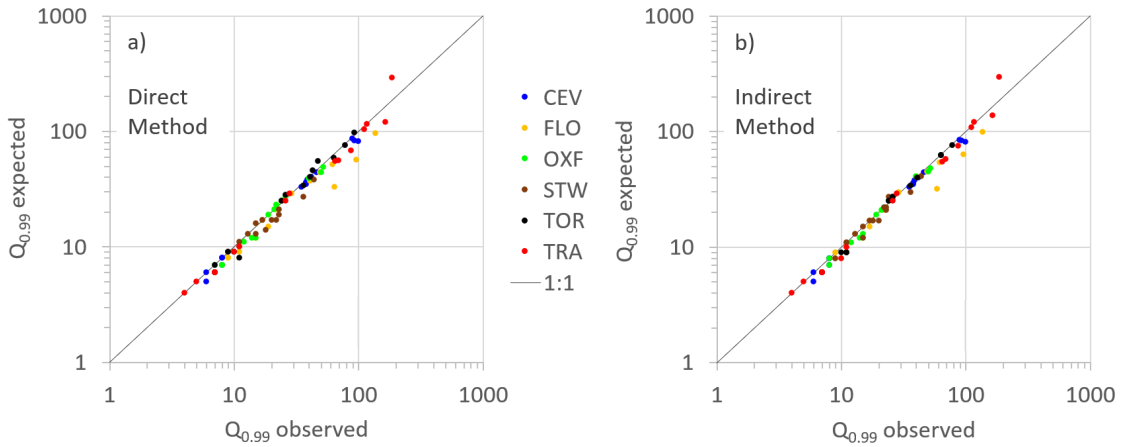

Figure 3.3.8: Comparison between empirical and theoretical quantile $Q_{0.99}$, calculated according to the DM (a) and the IM (b) for all stations, periods, and time variables bundled together.

observed frequencies for *it* are well reproduced at the annual scale, indicating that splitting the datasets into sub-periods is not strictly necessary to accurately model the probabilistic law of *it* in any of the rainfall regimes under consideration.

The scientific literature on the statistical inference of rainfall interarrival times remains relatively sparse, so further evidence of the suitability of the Lerch family to reproduce *it* distributions across a wide range of rainfall regimes is promising for future applications of this methodology. This consistency across diverse climatic conditions underscores the robustness and versatility of the Lerch distribution in modelling rainfall patterns.

However, the results obtained for *ws* and *ds* using the DM indicate that a good fit for *it* does not necessarily ensure an accurate reproduction of the frequencies of these derived quantities. Specifically, the geometric distribution often falls short in adequately describing wet spells. The enhanced performance of the IM, which utilises distributions that relax the memoryless property, suggests the presence of an underlying inner structure in multiday rainfall events. This structure indicates that the probability of rain is not constant within the event itself.

Typically, it is assumed that the internal structure observed in sub-daily (e.g., 10-minute) rainfall records diminishes at daily aggregation levels (see, e.g., Ridolfi et al., 2011). However, the results reported here challenge this assumption for some of the investigated sites, suggesting that the internal dynamics of rainfall events may persist even at daily time scales. This finding

underscores the necessity of employing more flexible modelling approaches, such as the IM, to accurately capture these complexities.

The need to utilise a more complex distribution than the geometric to accurately reproduce the probabilistic structure of *ws* has been emphasised in the literature, particularly for long spells. This was highlighted by Berger and Goossens (1983) for rainfall data in Belgium and by Deni et al. (2010) in Malaysia. Our findings confirm the inadequacy of the geometric distribution for long wet spells (greater than 10 days) that occur with relatively high frequency, as evidenced by the poor performance of the DM for STW.

It is important to emphasise that the unsatisfactory fitting of a geometric distribution does not necessarily indicate the presence of memory in the rainfall series. For instance, the results observed for the stations of TRA and FLO suggest other reasons for the geometric distribution's poor performance in modelling *ws* as derived from *it*. The geometric distribution appears to be a suitable choice when the data are analysed separately for the two seasons. This implies that the complex structure of the *ws* distribution observed over the entire year may not be due to an actual relaxation of the renewal property but rather to a mixing of *ws* samples from two distinct seasons.

However, the geometric distribution seems to perform poorly on STW regardless of seasonality, which is actually quite limited for this station. The STW station seems to represent a case where the memoryless property is violated, as also confirmed by the inspection of the $\frac{S_{k+1}}{S_k}$ ratios.

Splitting the entire dataset into sub-periods seems to improve the performance of fittings for both the DM and the IM methods. This result is particularly relevant for potential applications in ecohydrology models (e.g., D'Odorico, Ridolfi, Porporato, and Rodriguez-Iturbe, D'Odorico et al.; Petrie and Brunsell, 2012) and stochastic weather generators (for instance, Paek et al., 2023). In these fields, accurately representing the climatic component of weather variables (Semenov et al., 1998) requires not only reproducing the overall probabilistic structure of rainfall but also providing detailed information on a seasonal or even sub-seasonal (i.e., monthly) scale.

Studies in regions with climates that do not have distinct monsoon seasons have also highlighted the importance of focusing on specific seasons, such as the dry summer or wet winter periods (Caloiero and Coscarelli, 2020; Paton, 2022; Raymond et al., 2016). For example, Hui Wan and Barrow (2005) demonstrated the need to account for seasonality to accurately reproduce the duration of *ws* in Canada using a Markov chain method. This evidence underscores the value of seasonal analysis in improving the precision

and applicability of rainfall modelling for various operational and research purposes.

Another consequence of the geometric distribution's inadequacy in describing wet periods is the need to account for the daily structure of rainfall when modelling processes such as the seasonal dynamics of soil moisture and vegetation. For instance, Ratan and Venugopal (2013) assessed tropical areas using satellite rainfall data and found that wet spells typically lasted one day in dry regions, whereas durations of 2-4 days were more common in humid areas. Conversely, dry spells tended to last one day in humid regions and 3-4 days in dry regions. These findings emphasise the importance of considering detailed rainfall patterns in models to accurately capture the environmental processes across different climatic conditions.

In some instances, the results obtained for the IM suggest that using the classical geometric distribution for ws and the polylog series for ds provides adequate modelling of the observed frequencies. However, these distributions do not offer significant advantages over the 3-parameter Lerch distribution. Both DM and IM require a similar number of fitting parameters, but IM incurs a higher computational cost due to the need for two independent fittings (for ws and ds) as opposed to a single fitting for it.

An interesting case is STW, where the geometric distribution is never chosen for ws and a 2-parameter distribution is consistently required. Conversely, a 1-parameter distribution is sufficient for ds at this site. This scenario suggests that a reliable fit for both quantities can be achieved without increasing the total number of parameters, compared to using the 3-parameter Lerch distribution. This emphasises the flexibility and efficiency of the IM approach in providing accurate fits while managing computational demands effectively.

Finally, it is worth mentioning that the models proposed in this chapter are local, and thus spatial dependency in parameters may need to be considered when applying these models to multiple stations located in close proximity.

3.5 Conclusions

In this chapter, an analysis was conducted on daily rainfall data spanning a wide range of rainfall regimes across Europe (latitudes $38° - 58°$N) to model the frequency distribution of several key rainfall time variables. Two different methods were employed to investigate the validity of the simpler assumption of the rainfall occurrence process as a renewal process only governed by i.i.d.

interarrival times, which implies a geometric distribution of wet spells. The first method, referred to as the direct method (DM), assumed a geometric distribution of wet spells. The second method, the indirect method (IM), relaxed this assumption by modelling wet spells and dry spells separately, thereby allowing for the possibility of varying rainfall probabilities within the rainfall clusters.

In general, the comparison between the DM and IM methods indicates that the Lerch distribution can be effectively applied to model both interarrival times and dry spells across a wide range of rainfall regimes. However, it may be necessary to conduct a preliminary analysis of the memoryless property (e.g., by examining the $\frac{S_{k+1}}{S_k}$ ratios) to evaluate the reliability of wet spells derived from interarrival times modelled using the DM with the Lerch distribution. If signs of memory are present, the IM method is recommended, as it is better suited to a broader range of conditions, though it may require a larger number of fitting parameters.

The analysis was expanded to incorporate two additional time variables, closely linked to wet and dry spells, known as wet and dry chains. These variables extend the concept of wet and dry spells to sequences that are interrupted by a single non-rainy or rainy day, respectively, representing two measures that could be valuable for practical hydrological applications. The findings for these two chains generally align with the results obtained for the spells, though they also reveal additional challenges in probabilistic modelling, particularly at locations where the sample size may be a limiting factor.

The influence of seasonality on the results was also examined by dividing the data into two 6-month sub-periods. This division generally enhanced the performance of both the DM and IM methods, underscoring that in most locations, the DM applied to seasonal data remains a viable and straightforward approach. The findings of this study could prove useful in simulating scenarios of drought and flood events, given that probabilistic functions, like those utilised in this research, are fundamental to stochastic climate modelling.

Future research aimed at investigating the neighbouring location effects on parameter values could be developed. In other words, spatial dependency may need to be taken into account if dealing with a dense network of stations. Furthermore, as mentioned in the beginning of Section 3.3, we have left out the modelling of the volume of rainfall in a wet spell as we encountered a somehow impactful dependence between the rainfall depths and the length

of the wet spells. In both cases, it could be of great interest to provide a way to model dependence. The concept of copulas is a classical tool in the hydrological context for achieving this. As already thoroughly explained in the introduction, if, for instance, we aim to be consistent with the fully discrete description of the rainfall occurrence process we have opted for, we shall face the well known fact that copulas lose some of their charm when dealing with discrete margins. Most notably, they are not unique anymore. This is what prompted us to explore, both empirically and theoretically, alternative possibilities, with the broader aim of finding tools for the discrete case which are, at least conceptually, analogous to copulas. The results of this line of our research is contained in the next chapter.

Chapter 4

Copula-like Models for Bivariate Discrete Random Vectors

This chapter is based on Kojadinovic and Martini (2024)

Kojadinovic, I. and T. Martini (2024), Copula-like inference for discrete bivariate distributions with rectangular supports, *Electronic Journal of Statistics* 18 (1), 2571–2619.

For ease of explanation and notation, following Geenens (2020), we will focus on bivariate random vectors in this chapter. However, as shall be made clear in the future works section of Chapter 7, an extension of the concepts described from Section 4.2 onwards seems entirely possible and is part of a current project of us. In the following, with copula we will refer to the classical definition (see, e.g., Durante and Sempi, 2015, Definition 1.3.1), which, in the bivariate setting of interest reads

Definition 4.0.1. *A bivariate copula is a function* $C:[0,1]\times[0,1]\to[0,1]$ *that satisfies the following properties:*

1. *For every* $u,v\in[0,1]$,

$$C(u,0)=C(0,v)=0 \quad \textit{and} \quad C(u,1)=u, \quad C(1,v)=v.$$

2. *For every* $u_1,u_2,v_1,v_2\in[0,1]$ *such that* $u_1\leq u_2$ *and* $v_1\leq v_2$,

$$C(u_2,v_2)-C(u_2,v_1)-C(u_1,v_2)+C(u_1,v_1)\geq 0.$$

That is, C is the joint distribution function of a bivariate random vector on $[0,1]^2$ with uniform margins.

Copulas have become ubiquitous statistical tools for describing, analysing and modelling dependence between random variables thanks to the well known theorem of Sklar (1959), which we report here in the bivariate case.

Theorem 4.0.2. *Let F be a joint distribution function of a bivariate random vector (X,Y) with marginal distribution functions F_X and F_Y. Then there exists a copula C such that for all $x,y \in \mathbb{R}$,*

$$F(x,y) = C(F_X(x), F_Y(y)). \tag{4.1}$$

If F_X and F_Y are continuous, then C is unique. Otherwise, C is uniquely determined on $ran(F_X) \times ran(F_Y)$.

Furthermore, if C is a copula and F_X and F_Y are distribution functions, then the function F defined for $x,y \in \mathbb{R}$ by

$$F(x,y) = C(F_X(x), F_Y(y)) \tag{4.2}$$

is a joint distribution function with marginal distribution functions F_X and F_Y.

As we continue, with the *continuous case* (resp. *the discrete case*) we shall refer to (X,Y) having continuous (resp. discrete) margins. The popularity of copulas follows from a natural interpretation of Sklar's theorem: when C is unique, that is, in the continuous case, it is intuitive to see it in the role of the object which "glues" the margins together in order to construct their joint distribution. In this familiar context, through the Probability Integral Transform (PIT), $U = F_X(X)$ and $V = F_Y(Y)$ have uniform distributions on $[0,1]$, and the copula C is their joint distribution. Clearly, one can apply any increasing transformation to X and/or Y and the PIT will give the same result. This is the intuition behind the fact that copulas are invariant under increasing transformations of the margins (see, for instance, Hofert et al., 2018, Theorem 2.4.7). Additionally, it is known that some of the most used dependence concepts are functions of the copula alone in the continuous case (see, e.g., the classic Nelsen, 2006). These properties cement the interpretation and practical use of copulas as 'margin-free' embodiments of dependence in the continuous case. In the case X and/or Y discrete, $\text{ran}(F_X)$ and/or $\text{ran}(F_Y)$ are simply countable subsets of $[0,1]$. It is then clear that the distributions of $U = F_X(X)$ and/or $V = F_Y(Y)$ cannot be uniform

on [0,1]. Thus, their joint distribution cannot be a copula as per Definition 4.0.1. It is actually a subcopula, that is, a function which satisfies the main structural properties of copulas but whose support is only a strict subset of [0,1] containing 0 and 1 (Durante and Sempi, 2015, Definition 2.2.1). Only such a function is uniquely inferrable in the discrete case, according to Sklar's theorem. Any subcopula can be extended into a copula (Durante and Sempi, 2015, Lemma 2.3.4): the gaps in $[0,1]\setminus(\operatorname{Ran}F_X \times \operatorname{Ran}F_Y)$ can be filled in a way preserving the properties of copulas, by, informally speaking, suitably spreading the probability mass of the subcopula on each hyper-rectangle of its grid domain; however there are uncountably infinite many ways of doing so and C in (4.1) remains not identifiable. Furthermore, if one aims to construct measures of dependence based on a particular extension, the choice of the latter will clearly influence the result. The *standard extension* (also known as bilinear extension or checkerboard copula) has been deemed the most natural in the finite discrete case, as described for example in Genest and Nešlehová (2007). However, as noted in the previous reference, quantities based on such an extension still depend on the margins.

An enlightening example of the consequences of the non-uniqueness of the copula in the discrete case is the following simple one (Neslehova, 2004, Section 5.1).

Example 4.0.3. *Suppose $X \sim Bern(1-q_x)$ and $Y \sim Bern(1-q_y)$, where the probabilities $q_x = F_X(0)$ and $q_y = F_Y(0)$ lie in $(0,1)$, and X and Y are possibly dependent. According to Theorem 4.0.2 of Sklar, a copula C associated to $F_{X,Y}$ is only unique on the range $ran(F_X) \times ran(F_Y) = \{0, q_x, 1\} \times \{0, q_y, 1\}$. However, since the values of C are given by simple constraints along the sides of $[0,1]$ (the uniform margins constraints), the underlying copula is uniquely determined in a single point of the interior of the unit square and Sklar's representation reduces down to a single identity*

$$C(q_x, q_y) = F_{X,Y}(0,0). \tag{4.3}$$

Indeed, to recollect the joint bivariate Bernoulli distribution F_{XY}, given the margins and such a copula C, it is sufficient to only use $C(q_x, q_y)$ by

computing

$$\begin{aligned}
\mathbb{P}(X=0,Y=0) &= F_{X,Y}(0,0) = C(F_X(0),F_Y(0)) = C(q_x,q_y),\\
\mathbb{P}(X=0,Y=1) &= F_{X,Y}(0,1) - F_{X,Y}(0,0) =\\
&= C(F(0),F_Y(1)) - C(F(0),F_Y(0)) = q_x - C(q_x,q_y),\\
\mathbb{P}(X=1,Y=0) &= F_{X,Y}(1,0) - F_{X,Y}(0,0) =\\
&= C(F(1),F_Y(0)) - C(F(0),F_Y(0)) = q_y - C(q_x,q_y),\\
\mathbb{P}(X=1,Y=1) &= 1 - C(q_x,q_y) - (q_x - C(q_x,q_y)) - (q - C(q_x,q_y))\\
&= 1 - q_x - q_y + C(q_x,q_y).
\end{aligned}$$

Consequently, on one hand, there exist numerous copulas which lead to one and the same joint distribution, all of which pass through the point $(q_x,q_y,F_{X,Y}(0,0)) \in (0,1)^3$. *For instance, in the case of independence of* X *and* Y *where we must have* $C(q_x,q_y) = q_x q_y$, *while the independence copula naturally satisfies the requirement, many other do so as can be seen in the appendix of Geenens (2020). Any conclusions drawn from such a classical copula-based bivariate Bernoulli model would be highly questionable, as the central element* C *may interchangeably represent independence or dependence of significantly different strengths and characteristics, yet still conform to* $C(q_x,q_y) = q_x q_y$. *On the other hand,* $C(q_x,q_y)$ *depends crucially on the marginal distributions through the point* $(q_x,q_y) = (F_X(0),F_Y(0))$, *breaking down the margin free appeal of copulas.*

We refer the reader to Genest and Nešlehová (2007) and Faugeras (2017) for a detailed investigation of what could happen if one insists in using the tools of continuous copulas in the case of discrete (or even mixed) margins, especially regarding the statistical modelling aspects. Most importantly, they discuss how the usual inferential procedures used in copula modelling may break down in the discrete case, since they are built upon the uniqueness of the underlying copula. In the same references, some strategies to deal with the discrete scenario are explored. The developments contained in Geenens (2020) are a refreshing addition to those tools, accompanied by the novel peculiarity of rejuvenating historical ideas found in the vast literature on contingency tables.

In the upcoming section we try to summarise some of the main ideas on which our results presented in the rest of the chapter build upon. They were initially developed in Geenens (2020), which we closely follow for this introductory part.

4.1 Equivalence Classes of Dependence and Copulas

The first step taken by Geenens (2020) for providing a sound framework where the ideas of copulas can smoothly carry out to the discrete case is an ingenious definition of equivalence classes of dependence, starting from the continuous case. Before proceeding, we shall additionally clarify that the subject of the upcoming discussion will not be generalised or expanded in this thesis, but conceptually forms the backbone of our results and hence we deemed of interest to present it. Let (X,Y) be a continuous vector with distribution F_{XY}, and for simplicity assume that X and Y are supported on $[0,1]$. Indeed, without loss of generality, one can consider that we observe X and Y on the inverse logit scale. Furthermore, suppose that F_{XY} admits a density f_{XY} and marginal densities f_X and f_Y on the unit square. Now let $\mathcal{F}$ be the set of all bivariate probability densities on $[0,1]$, and $\mathcal{S}$ the set of all differentiable strictly increasing functions from $[0,1]$ to $[0,1]$. $\mathcal{S}$ is a group under function composition (denoted $\circ$), as is the cartesian product $\mathcal{S}\times\mathcal{S}$ under element-wise composition (denoted again by $\circ$, as its use shall be clear by the context), i.e. for $(\phi_1,\psi_1),(\phi_2,\psi_2)\in\mathcal{S}\times\mathcal{S}$, $(\phi_1,\psi_1)\circ(\phi_2,\psi_2)=(\phi_1\circ\phi_2,\psi_1\circ\psi_2)$.

For any $(\phi,\psi)\in\mathcal{S}\times\mathcal{S}$, define $g_{\phi,\psi}:\mathcal{F}\to\mathcal{F}$ as

$$g_{\phi,\psi}(f)(u,v)=\frac{f(\phi^{-1}(u),\psi^{-1}(v))}{\phi'(\phi^{-1}(u))\psi'(\psi^{-1}(v))},\quad (u,v)\in\mathcal{J}.$$

Then it is possible to show that $g_{\phi,\psi}$ is a group action of $(\mathcal{S}\times\mathcal{S},\circ)$ on $\mathcal{F}$. This defines orbits for any $f\in\mathcal{F}$, denoted

$$[f]=\{f^*\in\mathcal{F}:\exists(\phi,\psi)\in\mathcal{S}\times\mathcal{S}\text{ such that }f^*=g_{\phi,\psi}(f)\}.$$

These orbits induce an equivalence relation defined as $f_1\sim f_2$ whenever $[f_1]=[f_2]$, for any $f_1,f_2\in\mathcal{F}$. The quotient space $\mathcal{C}=\mathcal{F}/(\mathcal{S}\times\mathcal{S},\circ)$ is the set of all equivalence classes. Note that the denominator just provides a normalization of the resulting function so that it remains a probability density.

Hence, the class $[f_{XY}]$ is made up of all bivariate probability densities obtainable through what we could call 'marginal distortions' of f_{XY}. As the most natural interpretation, what persists through the marginal distortions should precisely be the dependence structure of the random vector. Geenens (2020) proclaims that the elements of $\mathcal{C}$ should be regarded as the most faithful to a 'true copula', because 'they genuinely are the links (*copulae*

in Latin) which cement marginals inside bivariate densities'. However, to avoid confusion with other classically established definitions of copulas, he suggests that $[f]$ should instead be named the 'nucleus' of a distribution f.

Then, it is clear that to retrieve the traditional definition of a continuous copula from this framework the group action that needs to applied is g_{F_X,F_Y}. See that

$$g_{F_X,F_Y}(f)(u,v) = \frac{f_{XY}(F_X^{-1}(u),F_Y^{-1}(v))}{f_X(F_X^{-1}(u))f_Y(F_Y^{-1}(v))}, \quad u,v \in [0,1]$$

is the density c of a copula for continuous random variables, as can be derived from (4.1). Furthermore, note that, albeit providing an universally understood tool for achieving this, the PIT is unnecessary for defining the nucleus, which contains all the dependence structure of a random vector.

Viewing the effect of $g_{\phi,\psi}$ as a mere reassignment of probability mass, this construction can thus be adapted *mutatis mutandis* to the discrete scenario, as detailed below. Before proceeding, let us underline that the upcoming developments are restricted to the finitely supported case. A suggestion for an extension to the infinite countable case is proposed in (Geenens, 2020, Section 8). However, a detailed treatment of such a yet unexplored scenario is out of the scope of this thesis. Nevertheless, it is clear that it could be subject of future investigations (see Section 7.2 of Chapter 7).

Let (X,Y) be a discrete random vector where X (resp. Y) can take r (resp. s) distinct values for some strictly positive integers r and s. As the framework we expose depends only on the values of the probability mass function (pmf) of (X,Y), without loss of generality, we shall consider, for notational convenience only, that (X,Y) takes its values in $I_{r,s} = [r] \times [s]$, where $[r] = \{1,\dots,r\}$ and $[s] = \{1,\dots,s\}$. Denote by p the pmf of (X,Y). It is immediate that p can be identified with a $r \times s$ matrix whose elements are $p_{ij} = \mathbb{P}(X=i,Y=j)$ for $(i,j) \in I_{r,s}$. Denote $p^{[1]} = (p_1^{[1]},p_2^{[1]},\dots,p_r^{[1]})$ and $p^{[2]} = (p_1^{[2]},p_2^{[2]},\dots,p_s^{[2]})$ the marginal distributions of p. That is, $p_i^{[1]} = \sum_{j=1}^{s} p_{ij} = \mathbb{P}(X=x_i)$ for i in $[r]$ (resp. $p_j^{[2]} = \sum_{i=1}^{r} p_{ij} = \mathbb{P}(Y=y_j)$ for j in $[s]$). It will be assumed throughout that $p_i^{[1]} > 0$ and $p_j^{[2]} > 0$ $\forall i,j$; i.e., no row or column of p is identically null. Furthermore we let $\text{supp}(p) = \{(i,j) \in I_{r,s} : p_{ij} > 0\}$ and when $\text{supp}(p) = I_{r,s}$ we say that p has a rectangular support. The elements of $I_{r,s} \setminus \text{supp}(p)$ are known as the structural zeroes of p and they play an important role in the dependence structure of (X,Y), as shall become clearer later.

The matrix form of p is commonly seen in the vast and historical literature

on contingency tables. Therein, the *odds ratio matrix* defined as the $(r-1)\times(s-1)$ matrix $\omega(p)$ whose elements are

$$\omega(p)_{ij} = \frac{p_{11}p_{ij}}{p_{1j}p_{i1}}, \quad (i,j) \in ([r]\setminus\{1\}) \times ([s]\setminus\{1\}), \tag{4.4}$$

is classically considered to encode the dependence between X and Y; see Agresti (2013, Section 2.4), Kateri (2014, p 43), Rudas (2018, p 123) or Geenens (2023, Section 4).

Remark 4.1.1. Clearly, when p has some structural zeroes, undefined cases of $\frac{0}{0}$ may arise in (4.4). As done in Geenens (2020), with a slight lack of rigor, we consider two odds ratios matrices to be equal when all of their well defined entries are equal. □

Note that $\omega(p)$ is margin-free in a certain sense, that is, unaffected by monotone increasing transformation of the margins of p in the form of multiplication of constants. To see this, simply consider two positive real vectors $c=(c_1,\dots,c_r)$ and $d=(d_1,\dots,d_s)$ and consider a diagonal matrix D whose entries are d (resp. a diagonal matrix C whose entries are c). Then, let

$$p^* = \frac{C\,p\,D}{\|C\,p\,D\|_1},$$

which corresponds to multiplying component per component the i-th row of p by c_i for $i\in[r]$ (resp. the j-th column of p by d_j for $j\in[s]$) and normalising with the elementwise 1-norm $\|\cdot\|_1$ of the resulting matrix. Then for any (i,j) in $I_{r,s}$ we have that

$$\omega(p^*)_{ij} = \frac{p_{11}c_1d_1p_{ij}d_ic_j}{p_{1j}d_1c_jp_{i1}d_ic_1} = \frac{p_{11}p_{ij}}{p_{1j}p_{i1}} = \omega(p)_{ij}.$$

In an entirely similar way as what has been done in the continuous case, we now build the corresponding equivalence classes of dependence in the present finite discrete scenario. For a positive integer m, define $\mathcal{D}_{m\times m}$ as the set of all diagonal $m\times m$ matrices whose diagonal entries are positive. $\mathcal{D}_{m\times m}$ is a group under matrix multiplication. For any positive integers r and s, $\mathcal{D}_{r\times r}\times\mathcal{D}_{s\times s}$ is a group under element-wise matrix multiplication. Furthermore, for now let $\mathcal{P}_{r\times s}$ be the set of bivariate pmfs on $I_{r,s}$.

Then for any $\phi\in\mathcal{D}_{r\times r}$ and any $\psi\in\mathcal{D}_{s\times s}$, we define the function $g_{\phi,\psi}\colon \mathcal{P}_{r\times s}\to\mathcal{P}_{r\times s}$ as

$$g_{\phi,\psi}(p) = \frac{\phi\,p\,\psi}{\|\phi\,p\,\psi\|_1}. \tag{4.5}$$

Similar as above, the matrices ϕ and ψ respectively multiply the rows and columns of p, distorting the margins. The corresponding orbit of p is given by

$$[p] = \{p^* \in \mathcal{P}_{R\times S} : \exists(\phi,\psi) \in \mathcal{D}_{r\times r} \times \mathcal{D}_{s\times s} \text{ such that } p^* = g_{\phi,\psi}(p)\}.$$

The equivalence relation and classes are defined analogously, with $p \sim p^*$ if and only if $[p] = [p^*]$. Further, define a limit point of $[p]$ as an element of $\mathcal{P}_{R\times S}$ which can be written as

$$\frac{(\prod_{k=1}^{\infty} \Phi_k)\, p\, (\prod_{k=1}^{\infty} \Psi_k)}{\|(\prod_{k=1}^{\infty} \Phi_k)\, p\, (\prod_{k=1}^{\infty} \Psi_k)\|_1} \tag{4.6}$$

for some sequences of matrices $\{\Phi_k\}_{k\geq 1} \in \mathcal{D}_{r\times r}$ and $\{\Psi_k\}_{k\geq 1} \in \mathcal{D}_{s\times s}$. Let $\bar{[p]}$ be the closure of $[p]$, that is, the union of $[p]$ and its limit points. The role of the latter shall eventually be clear in connection with the structural zeroes of p.

Having removed the marginal effects, we can again regard the orbits $[p]$ as equivalence classes of dependence. Indeed, fix any $\phi \in \mathcal{D}_{r\times r}$ and any $\psi \in \mathcal{D}_{s\times s}$. For any two $p, p^* \in \mathcal{P}_{R\times S}$, we have that $p \sim p^*$ implies that $\omega(p) = \omega(p^*)$, as all odds ratios are invariant under marginal distortions induced by $g_{\phi,\psi}$ as shown previously. The same holds even in the presence of any undefined elements in the odds ratio matrix, as $g_{\Phi,\Psi}$ leaves the zeros of p unaffected. Consequently, it is natural to again call $[p]$ the nucleus of the pmf p. If all entries of $\omega(p)$ are defined and positive, that is in the case of rectangular support of p, then $\omega(p) = \omega(p^*)$ implies $[p] = [p^*]$. However, one may find two $p_1, p_2 \in \mathcal{P}_{R\times S}$ with $\omega(p_1) = \omega(p_2)$ but $[p_1] \neq [p_2]$ when $\text{supp}(p_1) \neq \text{supp}(p_2)$, that is, when p_1 and p_2 show a different pattern of structural zeros. This is a strong hint to the preponderant role of structural zeros on the dependence structure. Clearly, the fact that some values of X and Y are not allowed simultaneously (that is, any $p_{ij} = 0$) can be understood as a strong form of interaction. Moreover, under our assumption of positive marginals of p, any $p_{ij} = 0$ automatically prevents independence of (X, Y). To see this, simply note that $p_{ij} = 0$ cannot equate $p_i^{[1]} p_j^{[2]}$. In such circumstances, we can infer the presence of a *regional dependence* effect analogous to the continuous counterpart (Holland and Wang, 1987).

Having built the equivalence classes, a representative of dependence can be selected as the pmf u in $[p]$ having uniform margins, in the spirit of proposing a candidate which is as margin free as possible. It is then natural to provide the following definition.

Definition 4.1.2. A bivariate $r \times s$ discrete copula is a bivariate discrete distribution of a vector (U,V) supported on $\{r\} \times \{s\}$, where $\{r\} = \left\{\frac{1}{r+1}, \frac{2}{r+1}, \ldots, \frac{r}{r+1}\right\}$ and $\{s\} = \left\{\frac{1}{s+1}, \frac{2}{s+1}, \ldots, \frac{s}{s+1}\right\}$, whose marginal distributions are discrete uniform on $\{r\}$ and $\{s\}$. The associated *copula pmf* is thus a bivariate discrete pmf on $\{r\} \times \{s\}$, such that for all $i \in \{s\}$, $\sum_{j=1}^{s} u_{ij} = \frac{1}{r}$, and for all $j \in \{s\}$, $\sum_{i=1}^{r} u_{ij} = \frac{1}{s}$, where $u_{ij} = \mathbb{P}(U = i, V = j)$. □

Remark 4.1.3. The latter is not new and the given definition of discrete copula parallels definitions which can be found for example in Kolesarova et al. (2006) and Perrone et al. (2019). These two references are part of the literature on discrete copulas which has been discussed in the introduction. However, as said there, discrete copulas have been mainly applied in the continuous case. Additionally, as shall become even more clear in the remaining part of the chapter, our developments only concern the copula pmf. That is why further connections to this literature shall not be explored in this thesis, and, for now, we simply borrow the first part of Definition 4.1.2, leaving such investigations for possible future work.

□

Before proceeding further, we revisit Example 4.0.3 in light of the new concepts we have just introduced, following Example 1 of Geenens (2024).

Example 4.1.4. *Recall the setup of Example 4.0.3 and additionally consider the corresponding bivariate pmf p, defined on $\{0,1\}^2$ for consistency with the previous example. It can be described by the following (2×2) contingency table*

$x \backslash y$	0	1	
0	p_{00}	p_{01}	$p_0^{[1]}$
1	p_{10}	p_{11}	$p_1^{[1]}$
	$p_0^{[2]}$	$p_1^{[2]}$	1

,

which follows the notation introduced above. Referring to Example 4.0.3, we actually have $p_0^{[1]} = q_x$ (resp. $p_0^{[2]} = q_y$) and $p_1^{[1]} = 1 - q_x$ (resp. $p_1^{[2]} = 1 - q_y$). However, for the sake of generality of this example, the notation for the marginals shall be left as in the table above.

As $\sum_{(i,j) \in \{0,1\}^2} p_{ij} = 1$, supposing the margins are unknown, there are initially three free parameters for p. Once the marginal values $p_1^{[1]}$ and $p_1^{[2]}$ are fixed, only one free parameter is left: it should be the one describing the dependence inside p, as it is the only value needed to recover the joint pmf

when its margins are fixed. If the dependence parameter is to be 'margin-free', it must be (any one-to-one function of) the odds-ratio

$$\omega(p) = \frac{p_{00}p_{11}}{p_{10}p_{01}},$$

according to our discussion above. Thus, it seems reasonable to identify the dependence structure to the value of $\omega(p)$ *in this case. As previously shown, Sklar's representation here reduces down to a single identity*

$$p_{00} = C(1-p_1^{[1]}, 1-p_1^{[2]}). \tag{4.7}$$

As computed in Example 4.0.3, for given $p_1^{[1]}$ *and* $p_1^{[2]}$*, this value is, indeed, enough to identify by substitution the other values* p_{01}, p_{10} *and* p_{11} *- hence, the whole pmf* p*. The odds ratio* $\omega(p)$ *can then be rewritten as*

$$\begin{aligned}&\omega_C(p_1^{[1]}, p_1^{[2]})\\ &\quad = \frac{C(1-p_1^{[1]}, 1-p_1^{[2]})(C(1-p_1^{[1]}, 1-p_1^{[2]}) + p_{10} + p_{01} - 1)}{(1-p_1^{[1]} - C(1-p_1^{[1]}, 1-p_1^{[2]}))(1-p_1^{[2]} - C(1-p_1^{[1]}, 1-p_1^{[2]}))}.\end{aligned} \tag{4.8}$$

For a given copula C*, this is a continuous function of* $(p_1^{[1]}, p_1^{[2]}) \in (0,1)^2$*. In fact, the only copula* C *guaranteeing* $\omega_C(p_1^{[1]}, p_1^{[2]})$ *to be constant in* $p_1^{[1]}$ *and* $p_1^{[2]}$ *is the Plackett copula, which was precisely designed to achieve that (Plackett, 1965). Thus, it is now even more evident that we may easily construct two bivariate Bernoulli distributions using the same copula* C*, but showing very different dependence structures for different pairs of marginal parameters* $(p_1^{[1]}, p_1^{[2]})$*. Once again, it is evident that equating copulas with dependence structures in this context is not reasonable.*

Furthermore, we can directly write (4.7) and (4.8) in terms of the unique subcopula H *of* p *as well, by setting*

$$p_{00} = H(1-p_1^{[1]}, 1-p_1^{[2]})$$

and the odds ratio $\omega(p)$ *as*

$$\begin{aligned}&\omega_H(1-p_1^{[1]}, 1-p_1^{[2]})\\ &\quad = \frac{H(1-p_1^{[1]}, 1-p_1^{[2]})(H(1-p_1^{[1]}, 1-p_1^{[2]}) + p_{10} + p_{01} - 1)}{(1-p_1^{[1]} - H(1-p_1^{[1]}, 1-p_1^{[2]}))(1-p_1^{[2]} - H(1-p_1^{[1]}, 1-p_1^{[2]}))}.\end{aligned} \tag{4.9}$$

A subtle yet significant distinction between (4.8) and (4.9), which justifies the different notation used, is that, unlike C, H is not defined elsewhere in the interior of $[0,1]^2$ than at $(1-p_1^{[1]},1-p_1^{[2]})$. Therefore, a notion of 'varying the marginal parameters while maintaining the same subcopula' would be clearly unfounded - changing the margins would, by definition, require considering a completely different subcopula. In fact, both H and C are margin-dependent to the exact same extent, but, as beautifully said in (Geenens, 2024, Example 1), "the 'blanket' nature of copulas, which cover the whole $[0,1]^2$ and always have uniform margins, may give the dangerously comfortable feeling that it is not the case for C".

To represent the dependence structure of p in a margin-free way, and therefore imitate the role of copulas in continuous settings, we aim to compute the bivariate Bernoulli distribution with the same odds-ratio as p (that is, belonging to $[p]$), but with "uninformative" uniform margins, as advocated in the framework exposed before this example. It is a simple algebraic exercise to show that, for any given $\omega>0$,

$u\backslash v$	0	1	
0	$\frac{1}{2}\frac{\omega}{1+\sqrt{\omega}}$	$\frac{1}{2}\frac{1}{1+\sqrt{\omega}}$	$\frac{1}{2}$
1	$\frac{1}{2}\frac{1}{1+\sqrt{\omega}}$	$\frac{1}{2}\frac{\omega}{1+\sqrt{\omega}}$	$\frac{1}{2}$
	$\frac{1}{2}$	$\frac{1}{2}$	1

is such a distribution, and it is unique. This is what could be called the Bernoulli copula pmf, as seen in (Geenens, 2020, Section 5), where more details can be found. Although evidently not a 'copula' according to the classical meaning, this Bernoulli copula pmf enjoys all the pleasant properties which make copulas successful in continuous cases. For example, computing its Pearson correlation coefficient would amount to computing a discrete margin free analogue of Spearman's rho (see the end of Section 4.4).

Consider now a pmf p and the associated equivalence class of dependence $[p]$. Natural questions arise regarding existence and especially uniqueness of a copula pmf $u\in[p]$, according to the broader aim of providing a copula-like framework. Moreover, when these two conditions hold, the actual computation of u has an equal interest. We briefly delineate the ideas contained in Geenens (2020) relative to this issue, as more details will be given in the upcoming main sections of this chapter. Following its line of thought, we start by the computation of u. It is clear that a proposed tool must only entail transformations of the type seen in (4.5) or (4.6). With this in mind, Geenens (2020) has cleverly suggested that the iterative proportional fitting

procedure (IPFP) fills the role in practice. This procedure takes the form of an algorithm (actually of several equivalent algorithms) designed to adjust the elements of an input matrix so that it satisfies specified row and column sums. When the input matrix is the representation of a pmf and the target rows and columns sums are selected as uniform, the connection is clear. Since its introduction by Kruithof (1937) for the calculation of telephone traffic, the IPFP turned out to play a major role in a surprisingly large number of different scientific contexts (see, e.g., Idel, 2016, for a comprehensive review). The adjustment is obtained by alternatively normalising the rows and columns of the starting pmf p in matrix form so that at each step of algorithm at least one of the required marginal constraints is satisfied. A little thought reveals that these operations corresponds to alternatively left and right multiplying p by diagonal matrices, keeping the odds ratios constant and making clear that each element of the iteration will belong to $[p]$, exactly as in (4.6). The limit of these iterations, when it exists, will correspond to the copula pmf u associated to p. It is now clear that the question of existence and uniqueness of u its intimately connected with the convergence of the IPFP. The key concept is that in such a circumstance, the IPFP actually converges to the I-projection of p (in the sense of Csiszár, 1975) on a Fréchet class of pmfs with uniform margins (see Proposition 4.2.5). That is the reason why we are willing to say that, ultimately, I-projections on a Fréchet class of pmfs can be thought as the underlying central ingredient of Geenens (2020)'s proposals. Indeed, as shall be shown, they are the cornerstone which paves the way to provide a Sklar like decomposition for bivariate discrete finitely supported distributions. Conditions for their existence and uniqueness, extensively studied in the literature, can be directly translated into existence and uniqueness conditions for copula pmfs and they are deeply related to the structural zeroes of the input pmf p. This will become clearer in Sections 4.2 and 4.3.

Given this framework, we are interested in its exploitation to set forth models for bivariate discrete random vectors in the spirit of copula modelling for continuous random vectors. In other words, we will not delve further into the concept of equivalence classes of dependence, but we shall explore the statistical modelling consequences of the interesting ideas that it has conveyed. We are then considering a typical statistical setup, where the pmf p of (X,Y) is not known, and instead, we have n (not necessarily independent) samples $(X_1,Y_1),\ldots,(X_n,Y_n)$ of (X,Y) available for inference on p. Traditional methods for modelling p when r and s are not excessively

large include log-linear models (see, e.g., Agresti, 2013; Kateri, 2014; Rudas, 2018) and association models (see, e.g., Goodman, 1985; Kateri, 2014). In this chapter, we will investigate a different category of (semi-parametric or parametric) models based on the possibility of decomposing the unknown pmf p into its two univariate margins and a bivariate pmf with uniform margins, namely the copula pmf we have discussed above. Indeed, another key interpretation of Geenens (2020) is that of establishing a strong case for such a decomposition.

A first contribution of this chapter, which holds independent significance, is the formulation of a result concerning the differentiability of an I-projection on a Fréchet class with respect to the input pmf. This result is essential for examining the asymptotics of the inference procedures proposed in the statistical portion of this chapter.

A second contribution, building upon the framework of Geenens (2020) delineated above, is the statement of a corollary based on known properties of I-projections. This corollary establishes that, under specific conditions on the bivariate pmf p, the latter can be decomposed into its two univariate margins and a bivariate pmf with uniform margins, that is the *copula pmf* u we have mentioned before. Analogous to the modelling of continuous multivariate distributions using copulas (see, e.g., Hofert et al., 2018, and the references therein), which utilises a well-known theorem of Sklar (1959), the resulting decomposition suggests first modelling the margins and the copula pmf separately, and subsequently combining the resulting estimates using an appropriate I-projection to obtain a parametric or semi-parametric estimate of p.

Exploiting the aforementioned decomposition of bivariate pmfs, a third contribution of this chapter is the investigation of nonparametric and parametric estimation procedures as well as goodness-of-fit tests for the underlying copula pmf. It is important to note that these analyses are conducted under the assumption of strict positivity of the initial unknown pmf p. Indeed, I-projections may not always exist, and to ensure the existence of the aforementioned copula-like decomposition of p, we will additionally assume in the statistical section of this chapter that $p_{ij} > 0$ for all $(i,j) \in I_{r,s}$. That is, that p has rectangular support. Interestingly, this assumption can be considered the discrete analog of the assumption of strict positivity (within the interior of the unit square) of the copula density, which is frequently made when modelling multivariate continuous distributions using copulas. It should be noted that, as will be discussed in more detail in our concluding

remarks, the assumption of rectangular support for p could be replaced with alternative conditions provided certain practical challenges are addressed. Fortunately, many applications appear to be compatible with the assumption of rectangular support.

The remaining part of this chapter is structured as follows. In the following section, after a review of the main properties of I-projections and the IPFP, we present differentiability results for I-projections on Fréchet classes. Next, given these recalled facts, we are able to shed light on the question of existence and uniqueness of a copula pmf u. In Section 4.4, we outline conditions under which a copula-like decomposition of bivariate pmfs based on I-projections on Fréchet classes exists. The next section is dedicated to the nonparametric and parametric estimation of copula pmfs and provides related asymptotic results as well as finite-sample findings based on simulations. Goodness-of-fit tests for copula pmfs are subsequently examined, both theoretically and empirically. A data example is provided in Section 4.7, followed by concluding remarks in the final section. For an easier reading, all proofs are relegated to a series of appendices.

4.2 I-Projections on Fréchet Classes and the IPFP

After introducing additional notation to lighten the upcoming discussions, we recollect the concept of I-projection as defined by Csiszár (1975) and discuss its relationship with the IPFP when the I-projection is performed on a Fréchet class. We conclude this section by presenting differentiability results for I-projections on Fréchet classes, which may hold independent interest.

4.2.1 Notation

Let $\mathbb{R}^{r\times s}$ denote the set of all $r\times s$ real matrices. It's important to distinguish between $\mathbb{R}^{r\times s}$ and $\mathbb{R}^{rs}$, where the latter represents the set of all rs-dimensional real vectors. For a matrix $x\in\mathbb{R}^{r\times s}$, the element in row $i\in[r]=\{1,\dots,r\}$ and column $j\in[s]=\{1,\dots,s\}$ will be denoted as x_{ij}. Additionally, for the row and column sums of x, we use the conventional notation

$$x_{i+}=\sum_{j=1}^{s}x_{ij},\qquad i\in[r],\qquad\text{and}\qquad x_{+j}=\sum_{i=1}^{r}x_{ij},\qquad j\in[s],$$

respectively. Moreover, we define the set Γ as

$$\Gamma = \{x \in \mathbb{R}^{r\times s} : x_{ij} \geq 0, x_{i+} > 0, x_{+j} > 0 \text{ for all } (i,j) \in I_{r,s} \text{ and } \sum_{(i,j)\in I_{r,s}} x_{ij} = 1\}. \quad (4.10)$$

In this context, Γ represents the set of all bivariate probability mass functions (pmfs) on $I_{r,s}$ whose univariate margins are strictly positive. As mentioned in the introduction, the developments that follow rely solely on r, s, and pmf values. Thus, focusing on pmfs on $I_{r,s}$ is done without loss of generality. Indeed, any pmf of interest, defined on the Cartesian product of two sets of reals with cardinalities r and s, respectively, can be "relocated" onto $I_{r,s}$.

As we proceed, to differentiate bivariate pmfs from other real matrices, we will consistently use lowercase letters for the former and uppercase letters for the latter. Additionally, given a bivariate pmf x in Γ, its first and second univariate margins will be denoted by $x^{[1]}$ and $x^{[2]}$, respectively. According to our notation, we have

$$x_i^{[1]} = x_{i+}, \qquad i \in [r], \qquad \text{and} \qquad x_j^{[2]} = x_{+j}, \qquad j \in [s].$$

Next, let a and b be fixed univariate pmfs on $[r]$ and $[s]$, respectively, such that $a_i > 0$ for all $i \in [r]$ and $b_j > 0$ for all $j \in [s]$. We define

$$\Gamma_{a,\cdot} = \{x \in \Gamma : x^{[1]} = a\} \quad (\text{resp. } \Gamma_{\cdot,b} = \{x \in \Gamma : x^{[2]} = b\}) \quad (4.11)$$

as the subset of Γ in (4.10) containing bivariate pmfs whose first (resp. second) margin is a (resp. b). Consequently,

$$\Gamma_{a,b} = \{x \in \Gamma : x^{[1]} = a \text{ and } x^{[2]} = b\} = \Gamma_{a,\cdot} \cap \Gamma_{\cdot,b} \quad (4.12)$$

is the *Fréchet class* of all bivariate pmfs whose first margin is a and second margin is b.

4.2.2 I-Projections on Fréchet Classes

To introduce the concept of I-projection as defined by Csiszár (1975), we first need to define the *information divergence*, also known as the *Kullback-Leibler divergence* or *relative entropy*, for a bivariate pmf $y \in \Gamma$ with respect to a bivariate pmf $x \in \Gamma$. This is given by

$$D(y|x) = \begin{cases} \sum_{(i,j)\in I_{r,s}} y_{ij} \log \frac{y_{ij}}{x_{ij}} & \text{if } \operatorname{supp}(y) \subset \operatorname{supp}(x), \\ \infty & \text{otherwise,} \end{cases} \quad (4.13)$$

where, for any $x \in \Gamma$, $\text{supp}(x) = (i,j) \in I_{r,s} : x_{ij} > 0$, with the conventions that $0 \log 0 = 0$ and $0 \log \frac{0}{0} = 0$.

The next result directly follows from Theorem 2.1 of Csiszár (1975) and the remark after Theorem 2.2 in the same reference.

Proposition 4.2.1 (Existence of I-projections)**.** *Let $x \in \Gamma$ and let $\mathcal{S}$ be a closed convex subset of Γ. Suppose furthermore that there exists a pmf $y \in \mathcal{S}$ such that $\text{supp}(y) \subset \text{supp}(x)$. Then, there exists a unique $y^* \in \mathcal{S}$ such that $y^* = \arg\inf_{y \in \mathcal{S}} D(y|x)$. Moreover, we have that $\text{supp}(y) \subset \text{supp}(y^*)$ for every $y \in \mathcal{S}$ such that $\text{supp}(y) \subset \text{supp}(x)$.*

The bivariate pmf y^* mentioned in the above proposition is commonly referred to as the *I-projection* of the initial bivariate pmf x onto $\mathcal{S}$.

As explained in Section 4.1, the approach proposed by Geenens (2020), on which our chapter hinges upon, relies on I-projections. Let a and b be fixed univariate pmfs on $[r]$ and $[s]$, respectively, such that $a_i > 0$ for all $i \in [r]$ and $b_j > 0$ for all $j \in [s]$. Let $\Gamma_{a,b}$, as defined in (4.12), be the Fréchet class of all bivariate pmfs whose first margin is a and second margin is b. Note that $\Gamma_{a,b}$ is a closed and convex subset of Γ. Moving forward, we shall denote the I-projection of a pmf $x \in \Gamma$ onto $\Gamma_{a,b}$, if it exists, as:

$$\mathcal{I}_{a,b}(x) = \arg\inf_{y \in \Gamma_{a,b}} D(y|x). \tag{4.14}$$

The following proposition, which, for example, follows directly from Corollary 3.3 in Csiszár (1975), provides a more explicit form for the I-projection of $x \in \Gamma$ onto the Fréchet classes $\Gamma_{a,b}$ when it contains a pmf whose support is equal to that of x.

Proposition 4.2.2. *Let $x \in \Gamma$. If there exists $y \in \Gamma_{a,b}$ such that $\text{supp}(y) = \text{supp}(x)$, then there exists two diagonal matrices $D_1 \in \mathbb{R}^{r \times r}$ and $D_2 \in \mathbb{R}^{s \times s}$ such that $\mathcal{I}_{a,b}(x) = y^* = D_1 x D_2$ and $\text{supp}(y^*) = \text{supp}(x)$. In this case, following Pretzel (1980), x and y^* are said to be diagonally equivalent.*

As can be seen from Propositions 4.2.1 and 4.2.2, the primary practical challenge before attempting to project a pmf $x \in \Gamma$ onto $\Gamma_{a,b}$ is to test for the existence of a bivariate pmf $y \in \Gamma_{a,b}$ such that $\text{supp}(y) \subset \text{supp}(x)$. Necessary and sufficient conditions for this have been recently restated in Theorem 1 of Brossard and Leuridan (2018), accompanied by a useful characterisation which shall be used later. These conditions appear to date back to Corollary 3 in Bacharach (1965). Furthermore, a well detailed description of them can also be found in Theorems 2 and 3 of Rothblum and Schneider (1989). Let

$x \in \Gamma$ be a bivariate pmf on $I_{r,s}$ and let P_x be the corresponding probability measure on $I_{r,s}$. Define

$$N_x = \{A \times B : A \subset [r], B \subset [s], P_x(A \times B) = 0\} \tag{4.15}$$

to be the set of rectangular subsets of $I_{r,s}$ on which P_x is null (or, equivalently, with a slight abuse of notation, x is null). Furthermore, let P_a (resp. P_b) be the probability measures on $[r]$ (resp. $[s]$) corresponding to the target univariate pmf a (resp. b). With this notation, we can state the following.

Proposition 4.2.3 (Testing for the existence of I-projections)**.** *Then, there exists* $y \in \Gamma_{a,b}$ *such that* $\operatorname{supp}(y) \subset \operatorname{supp}(x)$ *if and only if*

$$P_a(A) \leq P_b([s] \setminus B) \text{ for all } R \times C \in N_x. \tag{4.16}$$

and there exists $y \in \Gamma_{a,b}$ *such that* $\operatorname{supp}(y) = \operatorname{supp}(x)$ *if and only equalities in* (4.16) *happen only for* $A \times B$ *such that* $P_x([r] \setminus A \times [s] \setminus B) = 0$.

Remark 4.2.4. A detailed discussion of Proposition 4.2.2 and providing its proof would require delving deeply in the literature on the so-called matrix patterns, to which some of the references cited just above belong to. Indeed, according to, for instance, the definition given in (Brualdi, 2006, p 3), when employing the tools of matrix analysis $\operatorname{supp}(x)$ should be called the pattern of x. Such a literature investigates the existence and computation of matrices with given patterns and adhering to specific constraints. For example, given marginal sums. In this case, the problem is known as matrix scaling and the IPFP has been used both as a computational tool and as a way to prove the existence of such a scaling (see, e.g., Sinkhorn, 1974). The connections with our setup are then clear. Since such an investigation is out of scope of this thesis, we refer the reader to Appendix A of Idel (2016) for a nice review and the references therein for a deeper investigation (such as Brualdi, 2006). □

4.2.3 The Iterative Proportional Fitting Procedure

In practice, to perform an I-projection on $\Gamma_{a,b}$, one can employ the IPFP. Also known as Sinkhorn's algorithm or matrix scaling, the IPFP aims to adjust the elements of a matrix so that it satisfies specified row and column sums. For a detailed overview of the procedure, its variants, and its diverse applications, the reader is referred to Pukelsheim (2014), Idel (2016), and Brossard and Leuridan (2018).

In principle, the IPFP can be applied to any input matrix in $\mathbb{R}^{r\times s}$ with nonnegative elements and strictly positive row and column sums. Since this chapter has a probabilistic focus, the IPFP will be described specifically for input matrices $x\in\Gamma$ that can be interpreted as bivariate pmfs. In this context, the goal of the procedure is to adjust an input bivariate pmf $x\in\Gamma$ so that it has a as its first margin and b as its second margin.

The IPFP with target margins a and b consists of applying repeatedly two transformations. The first one corresponds to the map $\mathcal{R}_a$ from Γ in (4.10) to $\Gamma_{a,\cdot}$ in (4.11) defined by

$$\mathcal{R}_a(x)=\begin{bmatrix} \frac{a_1x_{11}}{x_{1+}} & \cdots & \frac{a_1x_{1s}}{x_{1+}} \\ \vdots & & \vdots \\ \frac{a_rx_{r1}}{x_{r+}} & \cdots & \frac{a_rx_{rs}}{x_{r+}} \end{bmatrix}, \qquad x\in\Gamma, \tag{4.17}$$

while the second one corresponds to the map $\mathcal{C}_b$ from Γ to $\Gamma_{\cdot,b}$ in (4.11) defined by

$$\mathcal{C}_b(x)=\begin{bmatrix} \frac{b_1x_{11}}{x_{+1}} & \cdots & \frac{b_sx_{1s}}{x_{+s}} \\ \vdots & & \vdots \\ \frac{b_1x_{r1}}{x_{+1}} & \cdots & \frac{b_sx_{rs}}{x_{+s}} \end{bmatrix}, \qquad x\in\Gamma. \tag{4.18}$$

Given an input matrix $x\in\Gamma$, the first map $\mathcal{R}_a(x)$ rescales the rows of x so that its first univariate margin equals a. Similarly, the second map $\mathcal{C}_b(x)$ rescales the columns of x so that its second univariate margin equals b.

Let $N\in\mathbb{N}$. Using the above defined maps, the $(2N+1)$th step of the IPFP with target margins a and b can then be expressed as

$$\mathcal{I}_{2N+1,a,b}(x)=\mathcal{R}_a\circ(\mathcal{C}_b\circ\mathcal{R}_a)^{(N)}(x), \qquad x\in\Gamma, \tag{4.19}$$

while the $2N$th step, $N\geq 1$, can be expressed as

$$\mathcal{I}_{2N,a,b}(x)=(\mathcal{C}_b\circ\mathcal{R}_a)^{(N)}(x), \qquad x\in\Gamma, \tag{4.20}$$

where the superscript (N) denotes composition N times and composition 0 times is taken to be the identity. We will say that the IPFP converges for a starting bivariate pmf $x\in\Gamma$ if the sequence $\{\mathcal{I}_{N,a,b}(x)\}_{N\geq 1}$ converges and we then denote the resulting adjusted pmf (which is in $\Gamma_{a,b}$ as shall become clearer below) by

$$\mathcal{I}_{a,b}(x)=\lim_{N\to\infty}\mathcal{I}_{N,a,b}(x). \tag{4.21}$$

When working in practice with the IPFP, if it converges for the bivariate pmf x, for some user-defined $\varepsilon>0$, $\lim_{N\to\infty}\mathcal{I}_{N,a,b}(x)$ is approximated by

$\mathcal{I}_{N,a,b}(x)$, where N (which depends on x) is an integer such that

$$\|\mathcal{I}_{N,a,b}(x) - \mathcal{I}_{N-1,a,b}(x)\|_{1,1} \leq \varepsilon, \tag{4.22}$$

where $\|x\|_{1,1} = \sum_{(i,j)\in I_{r,s}} |x_{ij}|$. Results on the number of iterations required for the previous condition to be satisfied are discussed for instance in Kalantari et al. (2008), Pukelsheim (2014) and Chakrabarty and Khanna (2021).

The next proposition, which, for instance, immediately follows from Pukelsheim (2014, Theorem 3) and Brossard and Leuridan (2018, Theorems 2 and 3), provides the fundamental relationship between I-projections on Fréchet classes and the IPFP.

Proposition 4.2.5. *(Link between I-projections on Fréchet classes and the IPFP) Let $x \in \Gamma$. The sequence $\{\mathcal{I}_{N,a,b}(x)\}_{N\geq 1}$ converges if and only if there exists $y \in \Gamma_{a,b}$ such that $\operatorname{supp}(y) \subset \operatorname{supp}(x)$ and, if $\lim_{N\to\infty} \mathcal{I}_{N,a,b}(x)$ exists, it is equal to $\mathcal{I}_{a,b}(x)$ in (4.14).*

In other words, the IPFP of a bivariate pmf x converges (to the I-projection of x on the corresponding Fréchet class) if and only if the I-projection of x on the corresponding Fréchet class exists.

4.2.4 Differentiability Results for I-Projections on Fréchet Classes

Theorem 3.3 in Gietl and Reffel (2017) implies a form of continuity of an I-projection on a Fréchet class when it exists (see Lemma 4.A.2 in Appendix 4.A for a precise statement). To examine the asymptotics of the statistical inference procedures proposed later in this chapter, we will also need an I-projection on a Fréchet class, when it exists, to be differentiable in a related sense, as will become evident in Section 4.5. This differentiability property has not been thoroughly investigated in the literature, although it can be connected to the work Jiménez-Gamero et al. (2011) as explained in Remark 4.2.7 below.

The aim of this section is to present a differentiability result, roughly speaking, concerning certain strictly positive submatrices of the input matrix. To make the statement precise, we first need to introduce additional notation.

For any $S \subset I_{r,s}$, $S \neq \emptyset$, let vec_S be the map from $\mathbb{R}^{r\times s}$ to $\mathbb{R}^{|S|}$ which, given a matrix $y \in \mathbb{R}^{r\times s}$, returns a column-major vectorization of y ignoring the elements y_{ij} such that $(i,j) \notin S$. With some abuse of notation, we can write

$$\operatorname{vec}_S(y) = (y_{ij})_{(i,j)\in S}, \qquad y \in \mathbb{R}^{r\times s}. \tag{4.23}$$

Let vec_S^{-1} be the map from $\mathbb{R}^{|S|}$ to $\mathbb{R}^{r\times s}$ which, given a vector v in $[0,1]^{|S|}$ expressed (with some abuse of notation) as $(v_{ij})_{(i,j)\in S}$, returns a matrix in $\mathbb{R}^{r\times s}$ whose element at position $(i,j)\in S$ is equal to v_{ij} and whose element at position $(i,j)\notin S$ is left unspecified. Then, for any two matrices $y,y'\in\mathbb{R}^{r\times s}$, write $y\overset{S}{=}y'$ if $y_{ij}=y'_{ij}$ for all $(i,j)\in S$. Furthermore, recall the definitions of Γ in (4.10) and $\Gamma_{a,b}$ in (4.12), and, for any $A,B\subset T\subset I_{r,s}$, $A,B\neq\emptyset$, let

$$\Gamma_T=\{y\in\Gamma:\text{supp}(y)=T\}, \tag{4.24}$$

$$\Lambda_{B,T}=\{w\in(0,1)^{|B|}:\text{ there exists } y\in\Gamma_T \text{ s.t. } \text{vec}_B^{-1}(w)\overset{B}{=}y\}, \tag{4.25}$$

$$\Gamma_{a,b,T}=\{y\in\Gamma_{a,b}:\text{supp}(y)=T\}, \tag{4.26}$$

$$\Lambda_{a,b,A,T}=\{z\in(0,1)^{|A|}:\text{ there exists } y\in\Gamma_{a,b,T} \text{ s.t. } \text{vec}_A^{-1}(z)\overset{A}{=}y\}. \tag{4.27}$$

Finally, for any function G from $\mathbb{R}^k\times\mathbb{R}^m$ to $\mathbb{R}^l$ differentiable at $(u,v)\in\mathbb{R}^k\times\mathbb{R}^m$, we shall use the notation

$$\partial_1 G(u,v)=\begin{bmatrix}\left.\frac{\partial G_1(w,v)}{\partial w_1}\right|_{w=u} & \cdots & \left.\frac{\partial G_1(w,v)}{\partial w_k}\right|_{w=u}\\ \vdots & & \vdots\\ \left.\frac{\partial G_l(w,v)}{\partial w_1}\right|_{w=u} & \cdots & \left.\frac{\partial G_l(w,v)}{\partial w_k}\right|_{w=u}\end{bmatrix}$$

and

$$\partial_2 G(u,v)=\begin{bmatrix}\left.\frac{\partial G_1(u,w)}{\partial w_1}\right|_{w=v} & \cdots & \left.\frac{\partial G_1(u,w)}{\partial w_m}\right|_{w=v}\\ \vdots & & \vdots\\ \left.\frac{\partial G_l(u,w)}{\partial w_1}\right|_{w=v} & \cdots & \left.\frac{\partial G_l(u,w)}{\partial w_m}\right|_{w=v}\end{bmatrix}.$$

The following proposition, proven in Appendix 4.A, is notationally complex but relies on a straightforward concept: a bivariate pmf with support T can be expressed in terms of $|B|=|T|-1$ elements and a pmf with support T and univariate margins a and b can be expressed in terms of $|A|<|B|$ elements. Consequently, the I-projection of a pmf in Γ_T on $\Gamma_{a,b,T}$ can be regarded as a map from an open subset of $(0,1)^{|B|}$ to $(0,1)^{|A|}$. This perspective allows us to regard the I-projection as a map related to an unconstrained optimization problem and, eventually, to apply the implicit function theorem (see, e.g., Fitzpatrick, 2009, Theorem 17.6, p 450).

Proposition 4.2.6. *Let $T\subset I_{r,s}$, $|T|>1$, such that $\Gamma_{a,b,T}$ in (4.26) is nonempty, and assume that there exists a nonempty subset $A\subsetneq T$ such that:*

1. *for any* $y \in \Gamma_{a,b,T}$, *the elements* y_{ij}, $(i,j) \in T \setminus A$, *can be recovered from the elements* y_{ij}, $(i,j) \in A$, *using the* $r+s$ *constraints* $y_{i+} = a_i$, $i \in [r]$, *and* $y_{+j} = b_j$, $j \in [s]$,

2. *the set* $\Lambda_{a,b,A,T}$ *in* (4.27) *is an open subset of* $\mathbb{R}^{|A|}$.

Let B *be any subset of* T *such that* $|B| = |T| - 1$ *and let* d *be the one-to-one map from* $\Lambda_{B,T}$ *in* (4.25) *to* Γ_T *in* (4.24) *which, given* $w \in \Lambda_{B,T}$, *returns the unique* $y \in \Gamma_T$ *obtained by recovering the only unspecified element of* $\text{vec}_B^{-1}(w) \in \Gamma_T$ *using the constraint* $\sum_{(i,j)\in T} y_{ij} = 1$. *Then, the map* $\text{vec}_A \circ \mathcal{I}_{a,b} \circ d$ *from* $\Lambda_{B,T}$ *to* $\Lambda_{a,b,A,T}$, *where* $\mathcal{I}_{a,b}$ *is defined in* (4.14), *is differentiable at any* $\text{vec}_B(x)$, $x \in \Gamma_T$, *with Jacobian matrix at* $\text{vec}_B(x)$ *equal to*

$$-[\partial_1 \partial_1 H(\text{vec}_A(\mathcal{I}_{a,b}(x)) \| \text{vec}_B(x))]^{-1} \partial_2 \partial_1 H(\text{vec}_A(\mathcal{I}_{a,b}(x)) \| \text{vec}_B(x)).$$

In the above centered display, H *is the map from* $\Lambda_{a,b,A,T} \times \Lambda_{B,T}$ *to* $[0, \infty)$ *defined by*

$$H(z\|w) = D(c(z)||d(w)), \qquad z \in \Lambda_{a,b,A,T}, w \in \Lambda_{B,T}, \tag{4.28}$$

where D *is defined in* (4.13) *and* c *is the one-to-one map from* $\Lambda_{a,b,A,T}$ *in* (4.27) *to* $\Gamma_{a,b,T}$ *in* (4.26) *which, given* $z \in \Lambda_{a,b,A,T}$, *returns the unique* $y \in \Gamma_{a,b,T}$ *obtained by recovering the unspecified elements of* $\text{vec}_A^{-1}(z) \in \Gamma_{a,b,T}$ *using the* $r+s$ *constraints* $y_{i+} = a_i$, $i \in [r]$ *and* $y_{+j} = b_j$, $j \in [s]$.

Remark 4.2.7. Following the suggestions of a Referee when revising Kojadinovic and Martini (2024), we explored the connections of the statistical results to be stated in Section 4.5 with *minimum divergence estimation* (see Remark 4.5.4 below for more details). After exploring the related literature, we realised that Lemma 1 in Jiménez-Gamero et al. (2011) is connected with Proposition 4.2.6 above. Upon some reflection, one can however see that the aforementioned lemma cannot be used to obtain a differentiability result for an I-projection on a Fréchet class. Furthermore, its proof seems incomplete as the assumptions considered in Jiménez-Gamero et al. (2011) do not seem to guarantee a key matrix invertibility required to apply the implicit function theorem (see (17.17) in Fitzpatrick, 2009, Theorem 17.6, p 450). This will be further explored in the next chapter. □

4.3 Existence and Uniqueness of the Copula Pmf

Having reviewed I-projections and the IPFP, while also providing their relationship, we are ready to restate Geenens (2020)'s Theorem 6.1 regarding

the existence and uniqueness of a copula pmf u associated to a pmf p, after properly defining u. Additionally, we shall report and expand its interesting discussion about the role of the structural zeroes of p in this regard. Let $u^{[1]}$ (resp. $u^{[2]}$) be the univariate pmf of the uniform distribution on $[r]$ (resp. $[s]$) and let $\Gamma_{\text{unif}} = \Gamma_{u^{[1]},u^{[2]}}$ be the Fréchet class of all bivariate pmfs on $I_{r,s}$ whose univariate margins are $u^{[1]}$ and $u^{[2]}$, respectively. The I-projection of a pmf $x \in \Gamma$ on Γ_{unif}, if it exists, will be denoted by

$$\mathcal{U}(x) = \mathcal{I}_{u^{[1]},u^{[2]}}(x), \tag{4.29}$$

where $\mathcal{I}_{u^{[1]},u^{[2]}}$ is defined as in (4.14) with $a = u^{[1]}$ and $b = u^{[2]}$. Hence, given p a bivariate pmf on $I_{r,s}$ with strictly positive margins, if it exists, the associated copula pmf u is defined by

$$u = \mathcal{U}(p). \tag{4.30}$$

In other words, u is the I-projection of p on the Fréchet class of all bivariate pmfs on $I_{r,s}$ whose univariate margins are $u^{[1]}$ and $u^{[2]}$.

Remark 4.3.1. As done in Geenens (2020, Definition 6.1), recall that in Section 4.1 we initially defined copula pmfs as bivariate pmfs on $\{(i/(r+1), j/(s+1)) : (i,j) \in I_{r,s}\} \subset [0,1]^2$, by analogy with bivariate copulas which are defined on $[0,1]^2$. However, from expression (4.30) onwards, we have chosen to define them as bivariate pmfs on $I_{r,s}$, for simplicity and notational convenience, as already explained in the previous sections. Indeed, we stress out again that the developments contained in this chapter do not depend on the values taken by the underlying bivariate random vector, which take the role of mere labels, but solely on the pmf values. □

The next result relies on combining Propositions 4.2.1, 4.2.2 and 4.2.3. It is proven in Appendix 4.B.

Theorem 4.3.2 (Theorem 6.1 in Geenens (2020))**.** *Let p be a bivariate pmf on $I_{r,s}$ with strictly positive margins and consider N_p as defined in (4.15).*

(a) Suppose that, for all $(A \times B) \in N_p$, $\dfrac{|A|}{r} + \dfrac{|B|}{s} < 1$. Then, u exists unique with $\text{supp}(u) = \text{supp}(p)$.

(b) Suppose that, for all $(A \times B) \in N_p$,

$$\frac{|A|}{r} + \frac{|B|}{s} \leq 1,$$

with $\frac{|\hat{A}|}{r}+\frac{|\hat{B}|}{s}=1$ *for some* $(\hat{A}\times\hat{B})\in N_p$.

(i) If, for all $(\tilde{A}\times\tilde{B})\in N_p$ *such that* $\frac{|\tilde{A}|}{r}+\frac{|\tilde{B}|}{s}=1$, *we have that*

$$([r]\setminus\tilde{A}\times[s]\setminus\tilde{B})\in N_p,$$

then u *exists unique with* $\mathrm{supp}(u)=\mathrm{supp}(p)$.

(ii) If there exists $(\tilde{A}\times\tilde{B})\in N_p$ *such that* $\frac{|\tilde{A}|}{r}+\frac{|\tilde{B}|}{s}=1$ *and*

$$([r]\setminus\tilde{A}\times[s]\setminus\tilde{B})\notin N_p,$$

then u *exists unique with* $\mathrm{supp}(u)$ *strictly contained in* $\mathrm{supp}(p)$.

(c) Suppose that there exists $(\tilde{A}\times\tilde{B})\in N_p$ *such that* $\frac{|\tilde{A}|}{r}+\frac{|\tilde{B}|}{s}>1$. *Then,* u *does not exist.*

From the latter theorem, it is now clear how the structural zeros of p determine the fate of u. Case (a) represents the straightforward scenario, encompassing all pmfs on $I_{r,s}$ without structural zeros, and even some with structural zeros as long as these zeros are not too prominent in the sense given by the theorem.

Case (b) is more critical. In scenario (b)(i), according to the vast literature briefly mentioned in Remark 4.2.4, it is possible to reorganise the matrix representation of p into a block-diagonal form through permutations of its rows and columns. Each sub-block of non-zero elements in p satisfies the strict inequalities in Proposition 4.2.3 and can be adjusted independently when modifying the margins with the IPFP. Basically, this allows u to be still expressed in the form (4.5). Conversely, in scenario (b)(ii), the matrix form of p cannot be reorganized into a block-diagonal form. To satisfy the uniform margins constraints, new zeros must be introduced, making u a limit point of $[p]$ in the sense of (4.6). From statement 2 (c) in Theorem 1 of Brossard and Leuridan (2018), we can say that the support of u is in the complement in $\mathrm{supp}(p)$ of the union of all the rectangular subsets $A^c\times B^c$ over all non-empty subsets $A\times B$ in $I_{r,s}$ such that, pardoning the abuse of notation, p is null on them and $p^{[1]}(A)=p^{[2]}([s]\setminus B)$. Simply put, the new zeros are added on

$$\bigcup_{(A\times B)\in\mathcal{N}_p}([r]\setminus A\times[s]\setminus B),$$

where

$$\mathcal{N}_p = \{(A \times B) \in N_p : p^{[1]}(A) = p^{[2]}([s] \setminus B)\}.$$

Abiding by Remark 4.1.1, it holds that $\omega(u) = \omega(p)$, making u the unique copula pmf of p once again.

Finally, case (c) highlights the absence of a copula pmf when the structural zeros form a substantial portion of p. Since these zeros cannot be converted into positive values by the IPFP and are fixed, there are insufficient degrees of freedom to adjust the margins effectively. In this scenario, the dependency between X and Y is heavily influenced by the structural zeros, rendering an approach based on odds ratios essentially ineffective.

These observations allow us to assert the following.

Corollary 4.3.3 (Existence and uniqueness of the copula pmf)**.** *A bivariate pmf p possesses a unique copula pmf u if and only if there exists a pmf v with uniform margins such that* $\operatorname{supp}(v) \subset \operatorname{supp}(p)$*. Equivalently, the same happens if and only if* $\frac{|A|}{r} + \frac{|B|}{s} \leq 1$ *for all* $(A \times B) \in N_p$*. By definition, the copula pmf u satisfies $\omega(u) = \omega(p)$ and* $\operatorname{supp}(u) \subset \operatorname{supp}(p)$*. If, additionally, we have that* $\operatorname{supp}(v) = \operatorname{supp}(p)$*, then* $\operatorname{supp}(u) = \operatorname{supp}(p)$*.*

4.4 Copula-Like Decomposition of Bivariate Pmfs

We shall now state a copula-like decomposition for a bivariate pmf on $I_{r,s}$ and strictly positive univariate margins, using the results given in the previous sections. The next proposition, proven in Appendix 4.C, will play, in the next sections, a role analog to Sklar's theorem when modelling continuous multivariate distributions using copulas.

Proposition 4.4.1 (Copula-like decomposition of bivariate pmfs)**.** *Let p be a bivariate pmf on $I_{r,s}$ with strictly positive univariate margins $p^{[1]}$ and $p^{[2]}$. Then, the following two statements are equivalent:*

1. *There exists a bivariate pmf v on $I_{r,s}$ with uniform margins such that* $\operatorname{supp}(v) = \operatorname{supp}(p)$*.*

2. *There exists a unique bivariate pmf u on $I_{r,s}$ with uniform margins such that*

$$p = \mathcal{I}_{p^{[1]},p^{[2]}}(u). \tag{4.31}$$

   Furthermore, the unique bivariate pmf u on $I_{r,s}$ with uniform margins in (4.31) is given by $u = \mathcal{U}(p)$, where $\mathcal{U}$ is defined in (4.29).

Example 4.4.2. *In the setting of the previous proposition, let us illustrate the fact that, if Assertion 1 does not hold, then (4.31) does not hold. Assume that p is the pmf on $I_{3,3}$ represented by the matrix*

$$p = \begin{bmatrix} \frac{1}{7} & \frac{1}{7} & \frac{1}{7} \\ \frac{1}{7} & 0 & \frac{1}{7} \\ \frac{1}{7} & 0 & \frac{1}{7} \end{bmatrix}.$$

With $u^{[1]} = u^{[2]}$ the uniform pmf on I_3, note that (4.16) is satisfied for $x = p$, $a = u^{[1]}$ and $b = u^{[2]}$ with equality holding for $A = \{2,3\}$ and $B = \{2\}$ and such that $P_x([r] \setminus A \times [s] \setminus B) = P_x(\{1\} \times \{1,3\}) = \frac{2}{7} \neq 0$. Therefore, according to Proposition 4.2.3, there exists $v \in \Gamma_{unif}$ with $\operatorname{supp}(v) \subset \operatorname{supp}(p)$, but there does not exist any $v \in \Gamma_{unif}$ such that $\operatorname{supp}(v) = \operatorname{supp}(p)$. Hence, from Proposition 4.2.1, we know that $u = \mathcal{U}(p)$ exists. From Proposition 4.2.5, we can compute u by using the IPFP via (4.19) and (4.20). It can be verified that, for any $N \in \mathbb{N}$ (that is, after a row scaling step of the IPFP),

$$\mathcal{R}_{u^{[1]}} \circ (\mathcal{C}_{u^{[2]}} \circ \mathcal{R}_{u^{[1]}})^{(2N)}(p) = \frac{1}{3} \begin{bmatrix} \frac{1}{3(N+1)} & \frac{1+3N}{3(N+1)} & \frac{1}{3(N+1)} \\ \frac{1}{2} & 0 & \frac{1}{2} \\ \frac{1}{2} & 0 & \frac{1}{2} \end{bmatrix},$$

where $\mathcal{R}_{u^{[1]}}$ and $\mathcal{C}_{u^{[2]}}$ are defined as in (4.17) and (4.18), respectively, with $a = b = u^{[1]} = u^{[2]}$, and that, for any $N > 0$ (that is, after a column scaling step of the IPFP),

$$(\mathcal{C}_{u^{[2]}} \circ \mathcal{R}_{u^{[1]}})^{(2N)}(p) = \frac{1}{3} \begin{bmatrix} \frac{1}{1+3N} & 1 & \frac{1}{1+3N} \\ \frac{3N}{2(1+3N)} & 0 & \frac{3N}{2(1+3N)} \\ \frac{3N}{2(1+3N)} & 0 & \frac{3N}{2(1+3N)} \end{bmatrix}.$$

Letting N tend to ∞ for both sequences, we obtain

$$u = \mathcal{U}(p) = \begin{bmatrix} 0 & \frac{1}{3} & 0 \\ \frac{1}{6} & 0 & \frac{1}{6} \\ \frac{1}{6} & 0 & \frac{1}{6} \end{bmatrix}$$

and we see that $\operatorname{supp}(u) \subsetneq \operatorname{supp}(p)$. Furthermore, $\mathcal{I}_{p^{[1]},p^{[2]}}(u)$ does not exist as a consequence of Proposition 4.2.3 since (4.16) with $x = u$, $a = p^{[1]}$ and $b = p^{[2]}$ is contradicted for $A = \{1\}$ and $B = \{1,3\}$. Thus, the decomposition in (4.31) does not hold.

Akin to Sklar's theorem, we see that in (4.31) the copula pmf u of p and the marginal pmfs $p^{[1]}$ and $p^{[2]}$ are "glued together" via an I-projection (on

$\Gamma_{p^{[1]},p^{[2]}}$) to obtain the bivariate pmf p. As already said in Section 4.1, and now formalised as a consequence of Proposition 4.2.2 or Corollary 4.3.3, u and p share the same odds ratio matrix. We want to further stress out that, as can actually be verified, again from Proposition 4.2.2, that any bivariate pmf $y \in \Gamma_{a,b}$ such that $\text{supp}(y) = \text{supp}(u)$ (assuming that it exists) obtained by an I-projection of u on $\Gamma_{a,b}$ would keep the odds ratio matrix constant. The same happens if $\text{supp}(y)$ is strictly contained in $\text{supp}(u)$, admitting a slight lack of rigor by following Remark 4.1.1. As thoroughly explained in Geenens (2020), and, as already mentioned in Section 4.1, the unique copula pmf u of p is simply a natural (yet arbitrarily chosen) representative of the equivalence class of all bivariate pmfs with the same odds ratio matrix. That is, according to the vast literature on contingency tables, with the same dependence structure.

In the continuous context, it is known that a continuous bivariate distribution can be summarised in terms of dependence by a moment of the underlying copula (such as Spearman's rho or Kendall's tau – see, e.g., Hofert et al., 2018, Chapter 2 and the references therein). Similarly, when the decomposition in (4.31) holds, a bivariate pmf p could be summarized by a moment of its copula pmf u. In the spirit of Spearman's rho, aGeenens (2020) proposed Pearson's linear correlation $\text{cor}(U,V)$ when (U,V) has pmf u as a possible summary of u. Geenens (2020, Section 6.6) suggested naming this quantity *Yule's coefficient*, due to its analogy with Yule's colligation coefficient (Yule, 1912) for 2 by 2 contingency tables. More specifically, let Υ be the map from Γ_{unif} to $[-1,1]$ defined, for any $v \in \Gamma_{\text{unif}}$, by

$$\Upsilon(v) = 3\sqrt{\frac{(r-1)(s-1)}{(r+1)(s+1)}}\left(\frac{4}{(r-1)(s-1)}\sum_{(i,j)\in I_{r,s}}(i-1)(j-1)\,v_{ij} - 1\right). \quad (4.32)$$

Yule's coefficient of p (or u) as proposed by Geenens (2020) is then simply

$$\rho = \text{cor}(U,V) = \Upsilon \circ \mathcal{U}(p) = \Upsilon(u). \quad (4.33)$$

Recall that p is the pmf of (X,Y). Note that ρ does not coincide with $\text{cor}(X,Y)$ in general.

Various other moments of u could be considered. As possible alternatives to ρ, we shall additionally use coefficients based on Goodman's and Kruskal's gamma (Goodman and Kruskal, 1954) and Kendall's tau b (see, e.g., Kendall and Gibbons, 1990) of (U,V) when (U,V) has pmf u. The following proposition, proven in Appendix 4.C, presents the expressions of Goodman's and Kruskal's gamma and Kendall's tau b for bivariate pmfs on $I_{r,s}$ with uniform margins.

Proposition 4.4.3. *For any pmf $v \in \Gamma_{unif}$, Goodman's and Kruskal's gamma of v can be expressed as*

$$G(v) = \frac{2\kappa(v) - 1 + 1/r + 1/s - \|v\|_2^2}{1 - 1/r - 1/s + \|v\|_2^2}, \tag{4.34}$$

and Kendall's tau b of v can be expressed as

$$T(v) = \frac{\sqrt{rs}\{2\kappa(v) - 1 + 1/r + 1/s - \|v\|_2^2\}}{\sqrt{(r-1)(s-1)}}, \tag{4.35}$$

where

$$\kappa(v) = 2 \sum_{\substack{i \in \{2,\ldots,r\} \\ j \in \{1,\ldots,s-1\}}} \sum_{\substack{i' \in \{1,\ldots,i-1\} \\ j' \in \{j+1,\ldots,s\}}} v_{ij} v_{i'j'}.$$

By analogy with Yule's coefficient in (4.33), we define the *gamma coefficient* of p (or u) and the *tau coefficient* of p (or u) to be

$$\gamma = G \circ \mathcal{U}(p) = G(u) \qquad \text{and} \qquad \tau = T \circ \mathcal{U}(p) = T(u), \tag{4.36}$$

respectively, where the maps G and T are defined in (4.34) and (4.35), respectively. It is important to keep in mind that these coefficients coincide with Goodman's and Kruskal's gamma and Kendall's tau b of p only when p has uniform margins. It is easy to verify that all three coefficients ρ, γ and τ are equal to 1 (resp. -1) when $r = s$ and the copula pmf u is a diagonal (resp. anti-diagonal) matrix. As explained in Geenens (2020, Sections 5 and 6), in this case, u corresponds to the upper (resp. lower) Fréchet bound for a square copula pmf.

4.5 Estimation of Copula Pmfs

Let us return to the setting considered previously at the end of Section 4.1, with the additional assumption of strict positivity of p. We are interested in modelling the bivariate pmf p of a discrete random vector (X,Y) with (rectangular) support $I_{r,s}$, that is, $\text{supp}(p) = I_{r,s}$. To achieve this, we have at our disposal n (not necessarily independent) copies $(X_1,Y_1), \ldots, (X_n,Y_n)$ of (X,Y). Note that there exists a trivial bivariate pmf with uniform margins with support equal to $I_{r,s}$: it is the independence copula pmf π defined by $\pi_{ij} = 1/(rs)$, $(i,j) \in I_{r,s}$. We can therefore immediately apply Proposition 4.4.1 to obtain the decomposition of p given in (4.31).

Equation (4.31) suggests proceeding in three steps to form a parametric or semi-parametric estimate of p, by analogy with estimation procedures classically employed in the context of copula modelling for continuous random variables (see, e.g., Hofert et al., 2018, Chapter 4 and the references therein). For instance, to obtain a parametric estimate of p, one could proceed as follows:

1. Estimate the univariate margins $p^{[1]}$ and $p^{[2]}$ of p parametrically; let $p^{[1,\alpha^{[n]}]}$ and $p^{[2,\beta^{[n]}]}$ be the resulting estimates.

2. Estimate the copula pmf u parametrically; let $u^{[\theta^{[n]}]}$ be the resulting estimate.

3. Form a parametric estimate of p in (4.31) via an I-projection as

$$p^{[\alpha^{[n]},\beta^{[n]},\theta^{[n]}]} = \mathcal{I}_{p^{[1,\alpha^{[n]}]},p^{[2,\beta^{[n]}]}}(u^{[\theta^{[n]}]}).$$

Remark 4.5.1. Using the notation of Section 4.2.3 and given $x \in \Gamma$ with $\operatorname{supp}(x) = I_{r,s}$, Theorem 6.2 of Gietl and Reffel (2017) implies that the IPFP of x depends continuously on x and on the underlying target marginal pmfs a and b. Given this, under a natural assumption of strict positivity of the parametric model for u (see Section 4.5.2 below), Proposition 4.2.5 and the continuous mapping theorem imply that $p^{[\alpha^{[n]},\beta^{[n]},\theta^{[n]}]} = \mathcal{I}_{p^{[1,\alpha^{[n]}]},p^{[2,\beta^{[n]}]}}(u^{[\theta^{[n]}]})$ is a consistent estimator of p if $p^{[1,\alpha^{[n]}]}$, $p^{[2,\beta^{[n]}]}$ and $u^{[\theta^{[n]}]}$ are consistent estimators of $p^{[1]}$, $p^{[2]}$ and u, respectively. □

The first estimation step above can be carried out using classical approaches in statistics. The focus of this section is on addressing the second step. However, as the parametric estimation of u will turn out to be strongly related to a natural nonparametric estimator of u, we first define and explore the latter. For the rest of this chapter, all convergences are as $n \to \infty$ unless otherwise stated.

4.5.1 Nonparametric Estimation

A straightforward approach to obtaining a nonparametric estimator of u is to use the plugin principle. Given that a natural estimator of p is $\hat{p}^{[n]}$, defined as

$$\hat{p}_{ij}^{[n]} = \frac{1}{n} \sum_{k=1}^{n} \mathbf{1}(X_k = i, Y_k = j), \qquad (i,j) \in I_{r,s}, \tag{4.37}$$

a meaningful estimator of u would simply be $\mathcal{U}(\hat{p}^{[n]})$, where $\mathcal{U}$ is defined in (4.29). However, upon closer inspection, this estimator may not always exist when n is small. Indeed, the fact that $\text{supp}(p) = I_{r,s}$ does not necessarily imply that $\text{supp}(\hat{p}^{[n]}) = I_{r,s}$ for all n. Furthermore, there is no guarantee that there exists a bivariate pmf $v^{[n]}$ on $I_{r,s}$ with uniform margins such that $\text{supp}(v^{[n]}) \subset \text{supp}(\hat{p}^{[n]})$, which is necessary to ensure that $\mathcal{U}(\hat{p}^{[n]})$ exists. To address this issue, we consider a smoothed version of $\hat{p}^{[n]}$ in (4.37). Although many solutions could be considered for the smoothed version of $\hat{p}^{[n]}$(see, e.g., Simonoff, 1995), we opt for one of the simplest approaches and consider the estimator $p^{[n]}$ of p defined by

$$p_{ij}^{[n]} = \frac{1}{n+1}\left\{\sum_{k=1}^{n} \mathbf{1}(X_k = i, Y_k = j) + q_{ij}\right\}, \qquad (i,j) \in I_{r,s}, \tag{4.38}$$

where q is a pmf on $I_{r,s}$ with $\text{supp}(q) = I_{r,s}$. This can be equivalently rewritten in terms of $\hat{p}^{[n]}$ in (4.37) and q as

$$p^{[n]} = \frac{n}{n+1}\hat{p}^{[n]} + \frac{1}{n+1}q. \tag{4.39}$$

In our numerical experiments, we considered two possibilities for q: the independence copula pmf π and another natural candidate, the pmf $\hat{p}^{[n,1]}\hat{p}^{[n,2],\top}$ which has copula pmf π and the same margins as $\hat{p}^{[n]}$ in (4.37). Note that the latter choice is only meaningful when both $\hat{p}^{[n,1]}$ and $\hat{p}^{[n,2]}$ are strictly positive; if this is not the case, a practitioner could decide to reduce the cardinalities r or s of the marginal supports. In our simulations, both choices for q yielded very similar results when applicable. For simplicity, we will take q equal to π as we proceed.

The estimator of u that we consider is then

$$u^{[n]} = \mathcal{U}(p^{[n]}) \tag{4.40}$$

and, by analogy with classical copula modelling, could be called the *empirical copula pmf*.

Denote convergence in probability with the arrow $\xrightarrow{\text{P}}$. The consistency of $u^{[n]}$ can be immediately deduced from the consistency of $\hat{p}^{[n]}$ in (4.37), the continuity of the I-projection of a strictly positive bivariate pmf on a Fréchet class (see Lemma 4.A.2 in Appendix 4.A) and the continuous mapping theorem.

Proposition 4.5.2 (Consistency of $u^{[n]}$). *Assume that* $\hat{p}^{[n]} \xrightarrow{\text{P}} p$ *in* $\mathbb{R}^{r\times s}$, *where* $\hat{p}^{[n]}$ *is defined in* (4.37). *Then,* $u^{[n]} \xrightarrow{\text{P}} u$ *in* $\mathbb{R}^{r\times s}$.

The next proposition, proven in Appendix 4.D, gives the limiting distribution of $\sqrt{n}(u^{[n]}-u)$ in $\mathbb{R}^{r\times s}$. It is mostly a consequence of Proposition 4.2.6 and the delta method (see, e.g., van der Vaart, 1998, Theorem 3.1).

Proposition 4.5.3 (Limiting distribution of $\sqrt{n}(u^{[n]}-u)$)**.** *Assume that* $\sqrt{n}(\hat{p}^{[n]}-p)\rightsquigarrow Z_p$ *in* $\mathbb{R}^{r\times s}$*, where* $\hat{p}^{[n]}$ *is defined in* (4.37)*, the arrow* $\rightsquigarrow$ *denotes weak convergence and* Z_p *is a random element of* $\mathbb{R}^{r\times s}$*. Then*

$$\sqrt{n}(u^{[n]}-u)=\mathcal{U}_p'(\sqrt{n}(\hat{p}^{[n]}-p))+o_P(1),$$

where

- $\mathcal{U}_p'$ *is the map from* $\mathbb{R}^{r\times s}$ *to* $\mathbb{R}^{r\times s}$ *defined by*

$$\mathcal{U}_p'(h)=\mathrm{vec}^{-1}(J_{u,p}\mathrm{vec}(h)),\qquad h\in\mathbb{R}^{r\times s},$$

- vec *is the operator defined as in* (4.23) *with* $S=I_{r,s}$*,*
- $J_{u,p}$ *is the* $rs\times rs$ *matrix given by*

$$J_{u,p}=-KL_u^{-1}M_pN, \tag{4.41}$$

- K *is the* $rs\times(r-1)(s-1)$ *matrix given by*

$$\begin{bmatrix} Q & 0 & \dots & 0\\ 0 & Q & \dots & 0\\ \vdots & \vdots & \ddots & \vdots\\ 0 & 0 & \dots & Q\\ -Q & -Q & \dots & -Q\end{bmatrix} \text{ where } Q=\begin{bmatrix} 1 & 0 & \dots & 0\\ 0 & 1 & \dots & 0\\ \vdots & \vdots & \ddots & \vdots\\ 0 & 0 & \dots & 1\\ -1 & -1 & \dots & -1\end{bmatrix}\in\mathbb{R}^{r\times(r-1)},$$

- L_u *is the* $(r-1)(s-1)\times(r-1)(s-1)$ *matrix whose element at row* $k+(r-1)(l-1)$*,* $(k,l)\in I_{r-1,s-1}$*, and column* $i+(r-1)(j-1)$*,* $(i,j)\in I_{r-1,s-1}$*, is given by*

$$\frac{\mathbf{1}(i=k,j=l)}{u_{kl}}+\frac{\mathbf{1}(j=l)}{u_{rl}}+\frac{\mathbf{1}(i=k)}{u_{ks}}+\frac{1}{u_{rs}}, \tag{4.42}$$

- M_p *is the* $(r-1)(s-1)\times(rs-1)$ *matrix whose element at row* $k+(r-1)(l-1)$*,* $(k,l)\in I_{r-1,s-1}$*, and column* $i+r(j-1)$*,* $(i,j)\in I_{rs}\setminus\{(i,j)\}$*, is given by*

$$-\frac{\mathbf{1}(i=k,j=l)}{p_{kl}}+\frac{\mathbf{1}(i=r,j=l)}{p_{rl}}+\frac{\mathbf{1}(i=k,j=s)}{p_{ks}}+\frac{1}{p_{rs}}, \tag{4.43}$$

- *and* N *is the* $(rs-1)\times rs$ *matrix given by*

$$\begin{bmatrix} 1 & 0 & \dots & 0 & 0 \\ 0 & 1 & \dots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \\ 0 & 0 & \dots & 1 & 0 \end{bmatrix}.$$

Consequently,

$$\sqrt{n}(u^{[n]}-u) \rightsquigarrow \mathcal{U}'_p(Z_p) \text{ in } \mathbb{R}^{r\times s}.$$

Moreover, when $(X_1,Y_1),\dots,(X_n,Y_n)$ *are independent copies of* (X,Y), $\mathrm{vec}(Z_p)$ *is a* rs*-dimensional centered normal random vector with covariance matrix* $\Sigma_p = \mathrm{diag}(\mathrm{vec}(p)) - \mathrm{vec}(p)\mathrm{vec}(p)^\top$ *and* $\mathrm{vec}(\mathcal{U}'_p(Z_p))$ *is a* rs*-dimensional centered normal random vector with covariance matrix* $J_{u,p}\Sigma_p J_{u,p}^\top$.

Remark 4.5.4. From (4.14) and (4.29), the empirical copula pmf $u^{[n]}$ defined in (4.40) can be rewritten more explicitly as

$$u^{[n]} = \operatorname*{arg\,inf}_{v\in\Gamma_{\mathrm{unif}}} D(v\|p^{[n]}). \tag{4.44}$$

As noted by a Referee when revising Kojadinovic and Martini (2024), the estimator mentioned above is a *minimum divergence estimator* (see, e.g., Read and Cressie, 1988; Morales et al., 1995; Basu et al., 2011). Such estimators have been explored for families of divergences (such as power-divergences in Read and Cressie (1988) and ϕ-divergences in Morales et al. (1995)) which include the Kullback–Leibler divergence as a particular case. When based on the Kullback–Leibler divergence, and using our notation, these estimators can be expressed as

$$v^{[n]} = \operatorname*{arg\,inf}_{v\in\Pi_0} D(v\|p^{[n]}), \tag{4.45}$$

where Π_0 is a subset of interest of Γ in (4.10). The strong similarity between (4.44) and (4.45) insinuates that asymptotic results for $u^{[n]}$ should follow from asymptotic results for $v^{[n]}$ upon taking $\Pi_0 = \Gamma_{\mathrm{unif}}$. However, an inspection of the appendices of, for instance, Read and Cressie (1988) reveals that asymptotic results for $v^{[n]}$ are typically obtained under the *null hypothesis* that $p\in\Pi_0$ (which is fully meaningful given the focus on goodness-of-fit testing in the aforementioned reference). In other words, upon taking Π_0 equal to Γ_{unif}, typical asymptotic results for minimum divergence estimators would allow us to obtain the asymptotics of $u^{[n]}$ under the

assumption that $p \in \Gamma_{\text{unif}}$. This would obviously not be of much interest since the approach studied in this chapter implicitly assumes that the unknown bivariate pmf p does not generally have uniform margins. Nevertheless, the connection with minimum divergence estimators indicates that the results stated in Proposition 4.5.3 should be a special case of asymptotic results for minimum divergence estimators under the *alternative hypothesis* that $p \notin \Pi_0$, i.e., *under misspecification*. After a literature review, we found only one reference addressing this issue: it is the work of Jiménez-Gamero et al. (2011). Specifically, in principle, Theorem 1 in the aforementioned reference could serve as an important building block for an alternative proof of Proposition 4.5.3 above. However, its proof seems incomplete as its relies on a lemma whose proof appears incomplete as already mentioned in Remark 4.2.7. The completion of this proof is one of the reasons behind the developments contained in our submitted paper Geenens et al. (2024) and shall be illustrated in Chapter 5. □

Recall the definitions of the moments ρ, γ and τ of p (or u) considered at the end of Section 4.4 in (4.33) and (4.36). Natural estimators of the latter then simply follow from the plugin principle and are

$$\rho^{[n]} = \Upsilon(u^{[n]}), \qquad \gamma^{[n]} = G(u^{[n]}) \qquad \text{and} \qquad \tau^{[n]} = T(u^{[n]}), \tag{4.46}$$

respectively, where $u^{[n]}$ is defined in (4.40) and the maps Υ, G and T are defined in (4.32), (4.34) and (4.35), respectively.

As we continue, for an arbitrary map $\eta : \Gamma \to \mathbb{R}$, we define its gradient $\dot{\eta}$ to be the usual gradient written with respect to the standard column-major vectorization of its $r \times s$ matrix argument, that is,

$$\dot{\eta} = \left(\frac{\partial \eta}{\partial x_{11}}, \dots, \frac{\partial \eta}{\partial x_{r1}}, \dots, \frac{\partial \eta}{\partial x_{1s}}, \dots, \frac{\partial \eta}{\partial x_{rs}} \right).$$

The following result is then an immediate corollary of Propositions 4.5.2 and 4.5.3, the continuous mapping theorem and the delta method.

Corollary 4.5.5 (Asymptotics of moment estimators)**.** *If* $\hat{p}^{[n]} \xrightarrow{\text{P}} p$ *in* $\mathbb{R}^{r \times s}$, *where* $\hat{p}^{[n]}$ *is defined in* (4.37), *then* $\rho^{[n]} \xrightarrow{\text{P}} \rho$, $\gamma^{[n]} \xrightarrow{\text{P}} \gamma$ *and* $\tau^{[n]} \xrightarrow{\text{P}} \tau$ *in* $\mathbb{R}$. *If, additionally,* $\sqrt{n}(\hat{p}^{[n]} - p)$ *converges weakly in* $\mathbb{R}^{r \times s}$, *then*

$$\begin{aligned}
\sqrt{n}(\rho^{[n]} - \rho) &= \dot{\Upsilon}(u)^\top J_{u,p} \sqrt{n}\,\text{vec}(\hat{p}^{[n]} - p) + o_P(1), \\
\sqrt{n}(\gamma^{[n]} - \gamma) &= \dot{G}(u)^\top J_{u,p} \sqrt{n}\,\text{vec}(\hat{p}^{[n]} - p)) + o_P(1), \\
\sqrt{n}(\tau^{[n]} - \tau) &= \dot{T}(u)^\top J_{u,p} \sqrt{n}\,\text{vec}(\hat{p}^{[n]} - p)) + o_P(1),
\end{aligned}$$

where $J_{u,p}$ is defined in (4.41). Consequently, when (X_1,Y_1), …, (X_n,Y_n) are independent copies of (X,Y), the sequences $\sqrt{n}(\rho^{[n]}-\rho)$, $\sqrt{n}(\gamma^{[n]}-\gamma)$ and $\sqrt{n}(\tau^{[n]}-\tau)$) are asymptotically centered normal with variances

$$\dot{\Upsilon}(u)^\top J_{u,p}\Sigma_p J_{u,p}^\top \dot{\Upsilon}(u),$$
$$\dot{G}(u)^\top J_{u,p}\Sigma_p J_{u,p}^\top \dot{G}(u),$$
$$\dot{T}(u)^\top J_{u,p}\Sigma_p J_{u,p}^\top \dot{T}(u),$$

respectively, where $\Sigma_p = \mathrm{diag}(\mathrm{vec}(p)) - \mathrm{vec}(p)\mathrm{vec}(p)^\top$.

4.5.2 Parametric Estimation

Let

$$\mathcal{J} = \{u^{[\theta]} : \theta \in \Theta\} \tag{4.47}$$

be a bivariate parametric family of copula pmfs on $I_{r,s}$, where Θ is an open subset of $\mathbb{R}^m$ for some strictly positive integer m. Because the assumption $\mathrm{supp}\,(p) = I_{r,s}$ implies that $\mathrm{supp}\,(u) = I_{r,s}$, the family $\mathcal{J}$ is naturally assumed to satisfy the following: for any $\theta \in \Theta$, $u_{ij}^{[\theta]} > 0$ for all $(i,j) \in I_{r,s}$. In this section, we assume that the unknown copula pmf u in (4.31) belongs to $\mathcal{J}$, that is, there exists $\theta_0 \in \Theta$ such that $u = u^{[\theta_0]}$, and our aim is to address the estimation of θ_0.

Before considering two estimation approaches, note that several examples of parametric copula pmfs can be found in Section 7 of Geenens (2020). Assuming a rectangular support for p (and thus u), it is particularly meaningful to follow one of the suggestions therein and construct the family $\mathcal{J}$ from a parametric family $\{C_\theta : \theta \in \Theta\}$ of classical bivariate copulas with strictly positive densities on $(0,1)^2$ such that, for any $\theta \in \Theta$ and $(i,j) \in I_{r,s}$,

$$u_{ij}^{[\theta]} = C_\theta\left(\frac{i}{r},\frac{j}{s}\right) - C_\theta\left(\frac{i}{r},\frac{j-1}{s}\right) - C_\theta\left(\frac{i-1}{r},\frac{j}{s}\right) + C_\theta\left(\frac{i-1}{r},\frac{j-1}{s}\right). \tag{4.48}$$

Clearly, this way of proceeding is fully meaningful only if the resulting family $\mathcal{J}$ in (4.47) is identifiable, that is, $u^{[\theta]} \neq u^{[\theta']}$ whenever $\theta \neq \theta'$. To verify that a family $\mathcal{J}$ is identifiable, it thus suffices to check that, for any $\theta \neq \theta'$, there exists $(i,j) \in I_{r,s}$ such that $u_{ij}^{[\theta]} \neq u_{ij}^{[\theta']}$. Note that the quantity $u_{ij}^{[\theta]}$ in (4.48) is actually the C_θ-volume of the rectangle $((i-1)/r, i/r] \times ((j-1)/s, j/s]$ (see, e.g., Hofert et al., 2018, Section 2.1). Non-identifiability thus occurs when a change in θ leaves the C_θ-volumes of all the rectangles $((i-1)/r, i/r] \times ((j-1)/s, j/s]$, $(i,j) \in I_{r,s}$, unchanged. Since the construction in (4.48) is based on

classical copula families (which are identifiable), arguably, non-identifiability of the resulting family $\mathcal{J}$ in (4.47) should be the exception rather than the rule in particular as r and s get larger. For a parametric copula family where C_θ has an explicit expression, identifiability of the family $\mathcal{J}$ can be checked analytically by replacing C_θ by its expression in (4.48). This is done for example in Geenens (2020, Section 7.1) for the Farlie–Gumbel–Morgenstern family. For parametric copula families for which C_θ is not explicitly available (such as elliptical copula families), identifiability could be checked numerically.

Method-of-moments Estimation

We shall assume in this subsection that $\mathcal{J}$ is a one-parameter family, that is, $m=1$. Given $\mathcal{J}$, let g_ρ, g_γ and g_τ be the functions defined, for any $\theta\in\Theta$, by

$$g_\rho(\theta)=\Upsilon(u^{[\theta]}),\qquad g_\gamma(\theta)=G(u^{[\theta]})\qquad\text{and}\qquad g_\tau(\theta)=T(u^{[\theta]}),\tag{4.49}$$

where the maps Υ, G and T are defined in (4.32), (4.34) and (4.35), respectively. Method-of-moments estimators based on Yule's coefficient, the gamma coefficient or the tau coefficient can be used if the functions g_ρ, g_γ and g_τ are one-to-one. In that case, corresponding estimators of θ_0 are simply given by

$$\theta_\rho^{[n]}=g_\rho^{-1}(\rho^{[n]}),\qquad \theta_\gamma^{[n]}=g_\gamma^{-1}(\gamma^{[n]})\qquad\text{and}\qquad \theta_\tau^{[n]}=g_\tau^{-1}(\tau^{[n]})\tag{4.50}$$

where $\rho^{[n]}$, $\gamma^{[n]}$ and $\tau^{[n]}$ are the estimators of ρ, γ and τ, respectively, defined in (4.46).

The following result is then an immediate consequence of Corollary 4.5.5, the continuous mapping theorem and the delta method.

Corollary 4.5.6 (Asymptotics of method-of-moments estimators)**.** *Assume that the functions g_ρ, g_γ and g_τ in (4.49) are one-to-one and that g_ρ^{-1}, g_γ^{-1} and g_τ^{-1} are continuously differentiable at $\rho_0=\Upsilon(u^{[\theta_0]})$, $\gamma_0=G(u^{[\theta_0]})$ and $\tau_0=T(u^{[\theta_0]})$. If $\hat{p}^{[n]}\xrightarrow{\mathrm{P}}p$ in $\mathbb{R}^{r\times s}$, then $\theta_\rho^{[n]}\xrightarrow{\mathrm{P}}\theta_0$, $\theta_\gamma^{[n]}\xrightarrow{\mathrm{P}}\theta_0$ and $\theta_\tau^{[n]}\xrightarrow{\mathrm{P}}\theta_0$ in $\mathbb{R}$.*

If, additionally, $\sqrt{n}(\hat{p}^{[n]}-p)$ *converges weakly in* $\mathbb{R}^{r\times s}$*, then*

$$\sqrt{n}(\theta_\rho^{[n]}-\theta_0)=\frac{\dot{\Upsilon}(u)^\top J_{u,p}\sqrt{n}\,\mathrm{vec}(\hat{p}^{[n]}-p)}{g'_\rho(\theta_0)}+o_P(1),$$
$$\sqrt{n}(\theta_\gamma^{[n]}-\theta_0)=\frac{\dot{G}(u)^\top J_{u,p}\sqrt{n}\,\mathrm{vec}(\hat{p}^{[n]}-p)}{g'_\gamma(\theta_0)}+o_P(1),$$
$$\sqrt{n}(\theta_\tau^{[n]}-\theta_0)=\frac{\dot{T}(u)^\top J_{u,p}\sqrt{n}\,\mathrm{vec}(\hat{p}^{[n]}-p)}{g'_\tau(\theta_0)}+o_P(1),$$

where $J_{u,p}$ *is defined in* (4.41).

Maximum Pseudo-likelihood Estimation

In this subsection, we assume that $\mathcal{J}$ in (4.47) is a multi-parameter family, meaning $m\geq 1$, with Θ an open subset of $\mathbb{R}^m$. Recall that we are working under the assumption that there exists $\theta_0\in\Theta$ such that $u=u^{[\theta_0]}$, and our goal is to estimate the unknown parameter vector θ_0. It is crucial to note that we do not have at our disposal observed counts from the bivariate pmf $u=u^{[\theta_0]}$. Instead, we only have access to the observed counts $n\hat{p}^{[n]}$ from p, where $\hat{p}^{[n]}$ is defined in (4.37). To perform maximum likelihood estimation, we would thus additionally need to postulate marginal parametric models for $p^{[1]}$ and $p^{[2]}$ and obtain an estimate of θ_0 as a by-product of the estimation of all the parameters of a model for p (see Remark 4.5.8 below). Rather the following such a complex approach, it is conceptually simpler to consider minimum divergence estimators of the form $\check{\theta}^{[n]}=\arg\inf_{\theta\in\Theta}D(u^{[\theta]}\|u^{[n]})$ or

$$\theta^{[n]}=\operatorname*{arg\,inf}_{\theta\in\Theta}D(u^{[n]}\|u^{[\theta]})=\operatorname*{arg\,sup}_{\theta\in\Theta}\sum_{(i,j)\in I_{r,s}}u_{ij}^{[n]}\log u_{ij}^{[\theta]},\tag{4.51}$$

where D is the Kullback–Leibler divergence defined in (4.13). Notice that if the numbers in $nu^{[n]}$ were counts obtained from a random sample from $u=u^{[\theta_0]}$,

$$\bar{L}^{[n]}(\theta)=n\sum_{(i,j)\in I_{r,s}}u_{ij}^{[n]}\log u_{ij}^{[\theta]},\qquad\theta\in\Theta,\tag{4.52}$$

would be the log-likelihood of the model. As the numbers in $nu^{[n]}$ are only a proxy to observed counts from $u^{[\theta_0]}$, the estimator in (4.51), which can be rewritten as

$$\theta^{[n]}=(\theta_1^{[n]},\ldots,\theta_m^{[n]})=\operatorname*{arg\,sup}_{\theta\in\Theta}\bar{L}^{[n]}(\theta)=\operatorname*{arg\,sup}_{\theta\in\Theta}\frac{1}{n}\bar{L}^{[n]}(\theta),\tag{4.53}$$

is a maximum *pseudo-likelihood* estimator of θ_0. The aim of this section is to derive its consistency and its asymptotic normality.

Remark 4.5.7. (Connection to minimum divergence estimators in multinomial models). Let $\mathcal{F}=\{p^{[\delta]}:\delta\in\Delta\}$ be a parametric family of bivariate pmfs, where Δ is a open subset of $\mathbb{R}^d$ for some strictly positive integer d. Assume additionally that $(X_1,Y_1),\ldots,(X_n,Y_n)$ is a random sample from p (which implies that $n\text{vec}(\hat{p}^{[n]})$ is a multinomial random vector with parameters n and $\text{vec}(p)$, where $\hat{p}^{[n]}$ is defined in (4.37) and vec is the operator defined as in (4.23) with $S=I_{r,s}$) and that there exists a $\delta_0\in\Delta$ such that $p=p^{[\delta_0]}$. From (4.51), we see that the maximum pseudo-likelihood estimator in (4.53) bears a strong resemblance with the estimator $\delta^{[n]}=\arg\inf_{\delta\in\Delta}D(p^{[n]}\|p^{[\delta]})$. The latter belongs to the classes of minimum divergence estimators of δ_0 studied for instance in Read and Cressie (1988), Morales et al. (1995) or Basu et al. (2011). Because the numbers in $nu^{[n]}$ are not observed counts from $u^{[\theta_0]}$, the consistency and the asymptotic normality of (4.53) cannot unfortunately be directly deduced from the asymptotic results stated in the aforementioned references. □

Remark 4.5.8. (Connection to maximum likelihood estimators). Let $\mathcal{M}_1=\{p^{[1,\alpha]}:\alpha\in A\}$ (resp. $\mathcal{M}_2=\{p^{[2,\beta]}:\beta\in B\}$) be a univariate parametric family of pmfs on $[r]$ (resp. $[s]$), where A (resp. B) is an open subset of $\mathbb{R}^{m_1}$ (resp. $\mathbb{R}^{m_2}$) for some strictly positive integer m_1 (resp. m_2). The families $\mathcal{M}_1$ and $\mathcal{M}_2$ are further naturally assumed to satisfy the following: for any $(\alpha,\beta)\in A\times B$, $p_i^{[1,\alpha]}>0$ for all $i\in[r]$ and $p_j^{[2,\beta]}>0$ for all $j\in[s]$. Having (4.31) in mind, one can combine the previous marginal parametric assumptions with (4.47) to form a parametric model for p as

$$\mathcal{P}=\{p^{[\alpha,\beta,\theta]}=\mathcal{I}_{p^{[1,\alpha]},p^{[2,\beta]}}(u^{[\theta]}):(\alpha,\beta,\theta)\in A\times B\times\Theta\}. \tag{4.54}$$

Under the assumption that there exists $(\alpha_0,\beta_0,\theta_0)\in(A,B,\Theta)$ such that $p=p^{[\alpha_0,\beta_0,\theta_0]}$, the estimation of θ_0 is then a by-product of estimating the entire parameter vector $(\alpha_0,\beta_0,\theta_0)$. When $(X_1,Y_1),\ldots,(X_n,Y_n)$ is a random sample from p, this could be obtained by maximising the log-likelihood of the model in (4.54), which can be written as

$$L^{[n]}(\alpha,\beta,\theta)=n\sum_{(i,j)\in I_{r,s}}\hat{p}_{ij}^{[n]}\log p_{ij}^{[\alpha,\beta,\theta]}=n\sum_{(i,j)\in I_{r,s}}\hat{p}_{ij}^{[n]}\log\mathcal{I}_{p^{[1,\alpha]},p^{[2,\beta]}}(u^{[\theta]})_{ij},$$

where $\hat{p}_n$ is defined in (4.37). Practically, this optimisation is expected to be computationally intensive as it would typically require numerous executions

of the IPFP. Moreover, similar to the indirect estimation of the parameter vector of a classical parametric copula by maximum likelihood estimation (see, e.g., Hofert et al., 2018, Chapter 4 and the references therein), the resulting estimate of θ_0 may be influenced by potential misspecification of the univariate families $\mathcal{M}_1$ and $\mathcal{M}_2$. A less computationally demanding "two-stage" approach involves first estimating α_0 (resp. β_0) by $\alpha^{[n]}$ (resp. $\beta^{[n]}$) from the first (resp. second) component sample $X_1,\dots,X_n$ (resp. $Y_1,\dots,Y_n$) and then maximising the following log-pseudo-likelihood:

$$\begin{aligned}\tilde{L}^{[n]}(\theta) &= n \sum_{(i,j)\in I_{r,s}} \hat{p}_{ij}^{[n]} \log p_{ij}^{[\alpha^{[n]},\beta^{[n]},\theta]} \\ &= n \sum_{(i,j)\in I_{r,s}} \hat{p}_{ij}^{[n]} \log \mathcal{I}_{p^{[1,\alpha^{[n]}]},p^{[2,\beta^{[n]}]}}(u^{[\theta]})_{ij}.\end{aligned}$$

From a computational standpoint, the maximisation of $\tilde{L}^{[n]}$ is expected to be significantly more expensive than that of $\bar{L}^{[n]}$ in (4.52) because the former typically requires numerous executions of the IPFP. Additionally, as before, the resulting estimate of θ_0 may be influenced by potential misspecification of the univariate families $\mathcal{M}_1$ and $\mathcal{M}_2$. This further supports the argument for maximising $\bar{L}^{[n]}$ in (4.52), which, in essence, is analogous to the log-pseudo-likelihood of Genest et al. (1995) in the current discrete context. □

The following result, proven in Appendix 4.E using Theorem 5.7 of van der Vaart (1998), provides conditions under which the estimator $\theta^{[n]}$ in (4.53) is consistent.

Proposition 4.5.9 (Consistency of the maximum pseudo-likelihood estimator)**.** *Assume that the family $\mathcal{J}$ is identifiable and that there exists $\lambda \in (0,1)$ such that, for any $\theta \in \Theta$ and $(i,j) \in I_{r,s}$, $u_{ij}^{[\theta]} \geq \lambda$. Then, if $\hat{p}^{[n]} \xrightarrow{\mathrm{P}} p$ in $\mathbb{R}^{r\times s}$, where $\hat{p}^{[n]}$ is defined in (4.37), $\theta^{[n]} \xrightarrow{\mathrm{P}} \theta_0$ in $\mathbb{R}^m$.*

Remark 4.5.10. The requirement that there exists $\lambda \in (0,1)$ such that, for any $\theta \in \Theta$ and $(i,j) \in I_{r,s}$, $u_{ij}^{[\theta]} \geq \lambda$ might seem overly restrictive. For instance, it can be shown that this condition will not be satisfied by families $\mathcal{J}$ in (4.47) constructed via (4.48) when C_θ is a Gumbel–Hougaard copula (see, e.g., Hofert et al., 2018, Chapter 2 and the references therein) if the parameter space Θ is taken as $(1,\infty)$. However, it will be met if the parameter space is restricted to $(1,M)$ for any fixed large real M. □

To state conditions under which the estimator $\theta^{[n]}$ in (4.53) is asymptotically normal, we need to introduce additional notation. For any $\theta \in \Theta$ and $(i,j) \in I_{r,s}$, let $\ell_{ij}^{[\theta]} = \log u_{ij}^{[\theta]}$, let

$$\dot{u}_{ij,k}^{[\theta]} = \frac{\partial u_{ij}^{[\theta]}}{\partial \theta_k}, k \in I_m = \{1,\ldots,m\}, \qquad \dot{u}_{ij}^{[\theta]} = \left(\dot{u}_{ij,1}^{[\theta]},\ldots,\dot{u}_{ij,m}^{[\theta]}\right), \tag{4.55}$$

and let

$$\ddot{u}_{ij,kl}^{[\theta]} = \frac{\partial^2 u_{ij}^{[\theta]}}{\partial \theta_k \partial \theta_l}, k,l \in I_m, \qquad \ddot{u}_{ij}^{[\theta]} = \begin{bmatrix} \ddot{u}_{ij,11}^{[\theta]} & \ldots & \ddot{u}_{ij,1m}^{[\theta]} \\ \vdots & & \vdots \\ \ddot{u}_{ij,m1}^{[\theta]} & \ldots & \ddot{u}_{ij,mm}^{[\theta]} \end{bmatrix}. \tag{4.56}$$

Similarly, for any $\theta \in \Theta$ and $(i,j) \in I_{r,s}$, let

$$\dot{\ell}_{ij,k}^{[\theta]} = \frac{\partial \log u_{ij}^{[\theta]}}{\partial \theta_k} = \frac{\dot{u}_{ij,k}^{[\theta]}}{u_{ij}^{[\theta]}}, k \in I_m, \qquad \dot{\ell}_{ij}^{[\theta]} = \left(\dot{\ell}_{ij,1}^{[\theta]},\ldots,\dot{\ell}_{ij,m}^{[\theta]}\right) = \frac{\dot{u}_{ij}^{[\theta]}}{u_{ij}^{[\theta]}}, \tag{4.57}$$

let

$$\ddot{\ell}_{ij,kl}^{[\theta]} = \frac{\partial^2 \ell_{ij}^{[\theta]}}{\partial \theta_k \partial \theta_l} = \frac{\ddot{u}_{ij,kl}^{[\theta]}}{u_{ij}^{[\theta]}} - \frac{\dot{u}_{ij,k}^{[\theta]} \dot{u}_{ij,l}^{[\theta]}}{(u_{ij}^{[\theta]})^2}, \qquad k,l \in I_m, \tag{4.58}$$

and let

$$\ddot{\ell}_{ij}^{[\theta]} = \begin{bmatrix} \ddot{\ell}_{ij,11}^{[\theta]} & \cdots & \ddot{\ell}_{ij,1m}^{[\theta]} \\ \vdots & & \vdots \\ \ddot{\ell}_{ij,m1}^{[\theta]} & \cdots & \ddot{\ell}_{ij,mm}^{[\theta]} \end{bmatrix} = \frac{\ddot{u}_{ij}^{[\theta]}}{u_{ij}^{[\theta]}} - \frac{\dot{u}_{ij}^{[\theta]} \dot{u}_{ij}^{[\theta],\top}}{(u_{ij}^{[\theta]})^2}, \tag{4.59}$$

where $\dot{u}_{ij}^{[\theta]}$ and $\ddot{u}_{ij}^{[\theta]}$ are defined in (4.55) and (4.56), respectively. Using the fact that $\sum_{(i,j)\in I_{r,s}} u_{ij}^{[\theta]} = 1$ implies that $\sum_{(i,j)\in I_{r,s}} \dot{u}_{ij,k}^{[\theta]} = 0$ and $\sum_{(i,j)\in I_{r,s}} \ddot{u}_{ij,kl}^{[\theta]} = 0$ for all $k,l \in I_m$, we obtain from (4.57) and (4.58) that

$$\sum_{(i,j)\in I_{r,s}} u_{ij}^{[\theta]} \dot{\ell}_{ij,k}^{[\theta]} = 0 \text{ for all } k \in I_m,$$

and that

$$\sum_{(i,j)\in I_{r,s}} u_{ij}^{[\theta]} \ddot{\ell}_{ij,kl}^{[\theta]} = -\sum_{(i,j)\in I_{r,s}} u_{ij}^{[\theta]} \dot{\ell}_{ij,k}^{[\theta]} \dot{\ell}_{ij,l}^{[\theta]} \text{ for all } k,l \in I_m.$$

Using the notation defined in (4.57) and (4.59), the previous two centered displays can be rewritten as the vector identity $\mathbb{E}_\theta(\dot{\ell}_{(U,V)}^{[\theta]}) = 0$ and the matrix identity

$$\mathbb{E}_\theta(\ddot{\ell}_{(U,V)}^{[\theta]}) = -\mathbb{E}_\theta(\dot{\ell}_{(U,V)}^{[\theta]} \dot{\ell}_{(U,V)}^{[\theta],\top}), \tag{4.60}$$

respectively, where (U,V) has pmf $u^{[\theta]}$ and $\mathbb{E}_\theta$ denotes the expectation with respect to $u^{[\theta]}$. In other words, in the discrete setting under consideration, unsurprisingly, we recover the classical identities that occur under regularity conditions in the context of classical maximum likelihood estimation (see, e.g., van der Vaart, 1998, Section 5.5, p 63).

The following result is proven in Appendix 4.E along the lines of the proof of Theorem 5.21 in van der Vaart (1998).

Proposition 4.5.11 (Asymptotic normality of the maximum pseudo-likelihood estimator)**.** *Assume that* $\theta^{[n]} \xrightarrow{\mathrm{P}} \theta_0$ *in* $\mathbb{R}^m$ *and, furthermore, that, for any* $(i,j) \in I_{r,s}$, $\theta \mapsto \ell_{ij}^{[\theta]}$ *is twice differentiable at any* $\theta \in \Theta$ *and that the matrix* $\mathbb{E}_{\theta_0}(\ddot{\ell}_{(U,V)}^{[\theta_0]})$, *where* (U,V) *has pmf* $u^{[\theta_0]}$, *is invertible. Then, if* $\sqrt{n}(\hat{p}^{[n]} - p)$ *converges weakly in* $\mathbb{R}^{r \times s}$, *we have that*

$$\sqrt{n}(\theta^{[n]} - \theta_0) = \{\mathbb{E}_{\theta_0}(\dot{\ell}_{(U,V)}^{[\theta_0]} \dot{\ell}_{(U,V)}^{[\theta_0],\top})\}^{-1} \dot{\ell}^{[\theta_0]} J_{u,p} \sqrt{n}\,\mathrm{vec}(\hat{p}^{[n]} - p) + o_P(1),$$

where $\dot{\ell}^{[\theta_0]}$ *is the* $m \times rs$ *matrix whose column* $i + r(j-1)$, $(i,j) \in I_{r,s}$, $\dot{\ell}_{ij}^{[\theta_0]}$ *is defined as in* (4.57) *with* $\theta = \theta_0$ *and* $J_{u,p}$ *is defined in* (4.41). *Consequently, when* $(X_1,Y_1),\ldots,(X_n,Y_n)$ *are independent copies of* (X,Y), *the sequence* $\sqrt{n}(\theta^{[n]} - \theta_0)$ *is asymptotically centered normal with covariance matrix*

$$\{\mathbb{E}_{\theta_0}(\dot{\ell}_{(U,V)}^{[\theta_0]} \dot{\ell}_{(U,V)}^{[\theta_0],\top})\}^{-1} \dot{\ell}^{[\theta_0]} J_{u,p} \Sigma_p J_{u,p}^\top \dot{\ell}^{[\theta_0],\top} [\{\mathbb{E}_{\theta_0}(\dot{\ell}_{(U,V)}^{[\theta_0]} \dot{\ell}_{(U,V)}^{[\theta_0],\top})\}^{-1}]^\top,$$

where $\Sigma_p = \mathrm{diag}(\mathrm{vec}(p)) - \mathrm{vec}(p)\mathrm{vec}(p)^\top$.

Remark 4.5.12. The conditions on the hypothesized family $\mathcal{J}$ in (4.47) stated in the previous proposition are inspired by some of the least restrictive ones in the literature (see, e.g., van der Vaart, 1998, Chapter 5). Therefore, we anticipate that they will hold for many families $\mathcal{J}$. □

4.5.3 Monte Carlo Experiments

The asymptotic results stated in Corollary 4.5.6 and Proposition 4.5.11 do not offer any insights on the finite-sample behaviour of the three method-of-moments estimators in (4.50) and the maximum pseudo-likelihood estimator in (4.53). To provide a comparison of these four estimators, we carried out Monte Carlo simulations under the assumption that the parametric family $\mathcal{J}$ in (4.47) is constructed from a one-parameter family of classical copulas as in (4.48). For $(r,s) \in \{(3,3),(3,10),(10,10)\}$, the bias and the mean squared error (MSE) of the three method-of-moments estimators and

Table 4.5.1: For $(r,s) \in \{(3,3),(3,10),(10,10)\}$, bias and mean squared error (MSE) of the three method-of-moment estimators in (4.50) and the maximum pseudo-likelihood (PL) estimator in (4.53) estimated from 1000 random samples of size $n \in \{100,500,1000\}$ generated, as explained in Section 4.5.3, from pmfs whose copula pmf is of the form (4.48) with C_θ the Clayton copula with a Kendall's tau of 0.33. The column 'm' gives the marginal scenario. The column '$\mathcal{U}$' reports the number of times the IPFP did not numerically converge in 1000 steps. The column 'ni' report the number of numerical issues related to fitting.

					Bias				MSE			
(r,s)	m	n	$\mathcal{U}$	ni	ρ	γ	τ	PL	ρ	γ	τ	PL
(3,3)	1	100	0	0	0.02	0.01	0.02	0.02	0.11	0.10	0.10	0.10
		500	0	0	0.01	0.01	0.01	0.01	0.02	0.02	0.02	0.02
		1000	0	0	0.00	0.00	0.00	0.00	0.01	0.01	0.01	0.01
	2	100	0	0	0.03	0.02	0.03	0.03	0.14	0.13	0.14	0.15
		500	0	0	0.00	-0.00	0.00	0.00	0.02	0.02	0.02	0.03
		1000	0	0	0.00	0.00	0.00	0.00	0.01	0.01	0.01	0.01
	3	100	0	0	0.06	0.05	0.06	0.06	0.19	0.19	0.19	0.18
		500	0	0	0.01	0.01	0.01	0.01	0.03	0.03	0.03	0.03
		1000	0	0	0.01	0.01	0.01	0.01	0.02	0.02	0.02	0.02
(3,10)	1	100	0	0	0.03	0.01	0.03	0.03	0.09	0.09	0.09	0.08
		500	0	0	-0.00	-0.00	-0.00	0.00	0.02	0.02	0.02	0.01
		1000	0	0	0.01	0.00	0.01	0.01	0.01	0.01	0.01	0.01
	2	100	0	0	-0.07	-0.09	-0.07	-0.08	0.16	0.16	0.16	0.20
		500	0	0	-0.00	-0.01	-0.01	-0.01	0.03	0.03	0.03	0.03
		1000	0	0	0.00	0.00	0.00	0.01	0.01	0.01	0.01	0.02
	3	100	0	1	-0.58	-0.58	-0.58	-0.64	0.45	0.45	0.45	0.52
		500	0	0	-0.25	-0.26	-0.25	-0.28	0.23	0.23	0.23	0.27
		1000	0	0	-0.08	-0.09	-0.09	-0.10	0.16	0.17	0.17	0.18
(10,10)	1	100	0	0	0.02	0.00	0.02	0.02	0.09	0.08	0.09	0.07
		500	0	0	0.01	0.00	0.01	0.01	0.01	0.01	0.01	0.01
		1000	0	0	0.01	0.00	0.01	0.01	0.01	0.01	0.01	0.01
	2	100	0	0	-0.14	-0.17	-0.15	-0.15	0.14	0.14	0.14	0.17
		500	0	0	-0.03	-0.03	-0.03	-0.03	0.03	0.03	0.03	0.04
		1000	0	0	-0.01	-0.01	-0.01	-0.01	0.01	0.01	0.01	0.02
	3	100	0	0	-0.85	-0.82	-0.81	-0.81	0.74	0.70	0.70	0.67
		500	0	0	-0.69	-0.66	-0.65	-0.68	0.52	0.49	0.49	0.51
		1000	0	0	-0.57	-0.55	-0.54	-0.60	0.39	0.38	0.37	0.44

the maximum pseudo-likelihood estimator were estimated from 1000 random samples of size $n \in \{100,500,1000\}$ generated from a pmf with copula pmf of the form (4.48) with C_θ either the Clayton or the Gumbel–Hougaard copula with a Kendall's tau in $\{0.33,0.66\}$, or the Frank copula with a Kendall's tau in $\{-0.5,0,0.5\}$ (see, e.g., Hofert et al., 2018, Chapter 2 for the definitions of these copula families and the definition of Kendall's tau). For each of the above seven copula pmfs, three marginal pmf scenarios were considered:

1. $(1/r,\ldots,1/r)$ (resp. $(1/s,\ldots,1/s)$) as values of the first (resp. second)

Table 4.5.2: For $(r,s) \in \{(3,3),(3,10),(10,10)\}$, bias and mean squared error (MSE) of the three method-of-moment estimators in (4.50) and the maximum pseudo-likelihood (PL) estimator in (4.53) estimated from 1000 random samples of size $n \in \{100,500,1000\}$ generated, as explained in Section 4.5.3, from p.m.f.s whose copula pmf is of the form (4.48) with C_θ the Clayton copula with a Kendall's tau of 0.66. The column 'm' indicates the marginal scenario. The column '$\mathcal{U}$' presents the number of times the IPFP did not numerically converge in 1000 steps. The column 'ni' reports the number of numerical issues related to fitting.

					Bias				MSE			
(r,s)	m	n	$\mathcal{U}$	ni	ρ	γ	τ	PL	ρ	γ	τ	PL
(3,3)	1	100	0	0	0.09	0.02	0.10	0.09	0.79	0.75	0.83	0.76
		500	0	0	0.01	-0.00	0.01	0.01	0.13	0.14	0.14	0.13
		1000	0	0	0.01	0.00	0.01	0.01	0.07	0.07	0.07	0.07
	2	100	0	0	0.09	-0.01	0.09	0.09	0.77	0.65	0.79	0.76
		500	0	0	-0.00	-0.02	0.00	-0.00	0.14	0.14	0.14	0.14
		1000	0	0	-0.00	-0.01	-0.00	0.00	0.07	0.07	0.07	0.07
	3	100	0	0	0.05	-0.03	0.06	0.05	0.77	0.74	0.81	0.72
		500	0	0	0.02	0.01	0.01	0.02	0.13	0.14	0.13	0.13
		1000	0	0	0.00	-0.00	0.00	0.00	0.08	0.09	0.08	0.08
(3,10)	1	100	0	0	0.02	-0.05	0.04	0.03	0.59	0.59	0.63	0.53
		500	0	0	0.01	-0.01	0.01	0.01	0.09	0.10	0.10	0.09
		1000	0	0	-0.00	-0.01	-0.00	-0.00	0.05	0.05	0.05	0.04
	2	100	0	0	-0.48	-0.52	-0.43	-0.65	0.96	0.97	0.94	1.34
		500	0	0	-0.04	-0.05	-0.03	-0.06	0.11	0.11	0.11	0.13
		1000	0	0	-0.02	-0.03	-0.02	-0.03	0.05	0.05	0.05	0.06
	3	100	0	0	-2.63	-2.59	-2.55	-2.82	7.28	7.11	6.92	8.40
		500	0	0	-1.39	-1.34	-1.30	-1.54	2.93	2.86	2.74	3.50
		1000	0	0	-0.73	-0.70	-0.67	-0.84	1.35	1.36	1.29	1.54
(10,10)	1	100	19	0	-0.14	-0.22	-0.10	-0.09	0.39	0.40	0.39	0.30
		500	0	0	-0.02	-0.03	-0.01	-0.01	0.07	0.07	0.06	0.05
		1000	0	0	0.01	-0.00	0.01	0.01	0.04	0.03	0.03	0.03
	2	100	5	0	-0.54	-0.56	-0.44	-0.54	0.79	0.77	0.71	1.05
		500	18	0	-0.04	-0.05	-0.02	-0.01	0.06	0.06	0.06	0.08
		1000	2	0	-0.02	-0.02	-0.01	-0.00	0.03	0.03	0.03	0.04
	3	100	0	0	-3.30	-3.11	-3.09	-3.24	11.03	9.88	9.75	10.76
		500	0	0	-2.30	-2.03	-1.97	-2.01	5.95	4.80	4.64	4.96
		1000	0	0	-1.68	-1.46	-1.39	-1.36	3.47	2.74	2.60	2.63

marginal pmf,

2. $(1,\ldots,r)/(r(r+1)/2)$ (resp. $(1,\ldots,s)/(s(s+1)/2)$) as values of the first (resp. second) marginal pmf,

3. the values of the pmf of the binomial distribution with parameters $r-1$ and $1/2$ (resp. $s-1$ and $1/2$) as values of the first (resp. second) marginal pmf

For the second and third marginal scenarios, the resulting pmfs of the

form (4.31) were computed using the IPFP. For each generated sample from one of the 21 data generating pmfs, we first computed the nonparametric estimate $p^{[n]}$ in (4.39), then, using the IPFP, the corresponding empirical copula pmf $u^{[n]}$ in (4.40) and, finally, the four estimates of θ_0 using (4.50) and (4.53). All the computations were carried out using the R statistical environment (R Core Team, 2024) and its packages `mipfp` (Barthélemy and Suesse, 2018) and `copula` (Hofert et al., 2022). In particular, the IPFP was computed using the function `Ipfp()` of the package `mipfp` with its default parameter values (at most 1000 iterations and ε in (4.22) equal to 10^{-10}), the inverses of the (one-to-one) functions g_ρ, g_γ and g_τ in (4.49) were computed by numerical root finding using the `uniroot()` function while the maximization of the log-pseudo-likelihood in (4.52) was carried out using the `optim()` function with $\theta_\gamma^{[n]}$ in (4.50) as starting value.

The results when the data-generating pmf is based on the copula pmf in (4.48) with C_θ as a Clayton copula are presented in Tables 4.5.1 and 4.5.2. Note that the column labeled '$\mathcal{U}$' indicates, for each data-generating scenario, the number of times the IPFP did not numerically converge within 1000 steps, as per the criterion in (4.22). The subsequent column, 'ni', reports the number of numerical issues encountered during fitting (either related to numerical root finding for the method-of-moments estimators or numerical optimization for the maximum pseudo-likelihood). An examination of all the fitting simulation results revealed that the overall number of numerical issues was minimal. These issues primarily concerned the Clayton model with the strongest dependence (see Table 4.5.2) when $r = s = 10$. This can be attributed to the higher likelihood of zero counts in the bivariate contingency tables derived from the generated samples in this scenario. Regarding estimation precision, it is reassuring to observe that for each data-generating scenario, an increase in n leads to a decrease in the absolute value of the bias and the Mean Squared Error (MSE). Unsurprisingly, the poorest results are observed for the third marginal scenario, as it leads to the largest number of very small probabilities for some cells of the data-generating pmf For instance, there are notable large negative biases when r or s are equal to 10, as the smallest marginal probability is approximately 0.002, often resulting in zero counts in several cells of the contingency tables derived from the generated samples. In essence, the very non-uniform margins in this scenario obscure certain features of u, making the dependence appear significantly weaker on average than it actually is.

From a numerical standpoint, it is also worth noting that for such data-generating scenarios, without the smoothing considered in (4.38), the IPFP would converge significantly less frequently within 1000 steps.

The results when the data generating pmf is based on the copula pmf (4.48) with C_θ a Gumbel–Hougaard or a Frank copula are not qualitatively different and are not reported. In terms of MSE, the experiments did not reveal a uniformly better estimator.

4.6 Goodness-Of-Fit Testing

Recall the definition of the parametric family of copula pmfs $\mathcal{J}$ in (4.47). The three method-of-moments estimators in (4.50) and the maximum pseudo-likelihood estimator in (4.53) were both theoretically and empirically examined in Sections 4.5.2 and 4.5.3 under the hypothesis

$$H_0 : u \in \mathcal{J}, \text{ that is, there exists } \theta_0 \in \Theta \text{ such that } u = u^{[\theta_0]}. \tag{4.61}$$

Clearly, parameter estimates obtained for a given family $\mathcal{J}$ will be meaningful only if H_0 actually holds. The goal of goodness-of-fit testing is formally evaluate this hypothesis. To derive goodness-of-fit tests, we consider, as is commonly done in the literature, approaches based on comparing a nonparametric estimator of u with a parametric one under H_0.

In the remainder of this section, we first define a relevant class of chi-square statistics and provide their asymptotic null distributions. Next, we empirically investigare the finite-sample performance of the resulting asymptotic goodness-of-fit tests. Finally, we propose a semi-parametric bootstrap procedure as an alternative method for computing p-values.

4.6.1 Asymptotic Chi-Square-Type Goodness-Of-Fit Tests

One approach to test $H_0 : u \in \mathcal{J}$ versus $H_1 : u \notin \mathcal{J}$ consists of constructing test statistics as norms of the *goodness-of-fit process*

$$\sqrt{n}(u^{[n]} - u^{[\hat{\theta}^{[n]}]}), \tag{4.62}$$

where $u^{[n]}$ is defined in (4.40) and $\hat{\theta}^{[n]}$ is an estimator of θ_0 computed under H_0. As is typical in the goodness-of-fit testing literature, the process in (4.62) compares a nonparametric estimator of u which is consistent whether H_0 is true or not, with a parametric estimator that is only consistent under H_0.

The rationale behind this construction is that any norm of (4.62) should be typically smaller under H_0 than it is under H_1.

The following result, proven in Appendix 4.F, describes the asymptotic null behavior of the goodness-of-fit process in (4.62).

Proposition 4.6.1. *Assume that H_0 in (4.61) holds and that*

1. *$\sqrt{n}(\hat{p}^{[n]} - p)$ converges weakly in $\mathbb{R}^{r\times s}$, where $\hat{p}^{[n]}$ is defined in (4.37),*
2. *the map from $\Theta \subset \mathbb{R}^m$ to $\mathbb{R}^{r\times s}$ defined by $\theta \mapsto u^{[\theta]}$ is differentiable at θ_0 and let $\dot{u}^{[\theta_0]}$ be the $rs \times m$ matrix whose row $i + r(j-1)$, $(i,j) \in I_{r,s}$, is $\dot{u}_{ij}^{[\theta]}$ in (4.55),*
3. *there exists a $m \times rs$ matrix $V_{u,p}^{[\theta_0]}$ such that*

$$\sqrt{n}(\hat{\theta}^{[n]} - \theta_0) = V_{u,p}^{[\theta_0]} \sqrt{n}\,\mathrm{vec}(\hat{p}^{[n]} - p) + o_P(1),$$

where vec *is the operator defined as in (4.23) with $S = I_{r,s}$.*

Then,

$$\sqrt{n}\,\mathrm{vec}(u^{[n]} - u^{[\hat{\theta}^{[n]}]}) = (J_{u,p} - \dot{u}^{[\theta_0]} V_{u,p}^{[\theta_0]}) \sqrt{n}\,\mathrm{vec}(\hat{p}^{[n]} - p) + o_P(1),$$

where $J_{u,p}$ is defined in (4.41).

Note that Corollary 4.5.6 and Proposition 4.5.11 give conditions under which the third assumption in the previous proposition holds for the three method-of-moments estimators in (4.50) and the maximum pseudo-likelihood estimator in (4.53), respectively.

One natural test statistic that can be constructed from the goodness-of-fit process in (4.62) is the chi-square statistic

$$S^{[n]} = n \sum_{(i,j)\in I_{r,s}} \frac{(u_{ij}^{[n]} - u_{ij}^{[\hat{\theta}^{[n]}]})^2}{u_{ij}^{[\hat{\theta}^{[n]}]}}. \tag{4.63}$$

As is well-known in the statistical literature, a widely accepted rule of thumb to protect the level of a classical chi-square test is to regroup low observed counts such that eventually all observed counts are above 5 (see, e.g., van der Vaart, 1998, Chapter 17). To perform similar groupings in our context, we consider the following generalization of $S^{[n]}$ in (4.63):

$$S_G^{[n]} = n \| \mathrm{diag}(G\,\mathrm{vec}(u^{[\hat{\theta}^{[n]}]}))^{-1/2} G\,\mathrm{vec}(u^{[n]} - u^{[\hat{\theta}^{[n]}]}) \|_2^2, \tag{4.64}$$

where G is a chosen "grouping matrix", that is, a $q \times rs$ matrix with $q \in \{1,\ldots,rs\}$ whose elements are in $\{0,1\}$ with exactly rs of them equal to 1 and a unique 1 per column. Some thought reveals that, in (4.64), each row of G sums one or more copula pmf values. Clearly, $S_G^{[n]}$ coincides with $S^{[n]}$ in (4.63) when G is the $rs \times rs$ identity matrix.

The following result, proven in Appendix 4.F, provides the asymptotic null distribution of $S_G^{[n]}$ in (4.64) when computed from a random sample.

Proposition 4.6.2. *Under the conditions of Proposition 4.6.1, when* (X_1,Y_1), $\ldots$, (X_n,Y_n) *are independent copies of* (X,Y), $S_G^{[n]} \rightsquigarrow \mathcal{L}_G$ *in* $\mathbb{R}$, *where* $\mathcal{L}_G$ *is distributed as* $\sum_{k=1}^{q} \lambda_k Z_k^2$ *for i.i.d.* $N(0,1)$*-distributed random variables* $Z_1,\ldots,Z_q$ *and* $\lambda_1,\ldots,\lambda_q$ *the eigenvalues of*

$$\Sigma_{G,u,p}^{[\theta_0]} = \operatorname{diag}(G\operatorname{vec}(u^{[\theta_0]}))^{-1/2} G(J_{u,p} - \dot{u}^{[\theta_0]} V_{u,p}^{[\theta_0]}) \times \Sigma_p (J_{u,p} - \dot{u}^{[\theta_0]} V_{u,p}^{[\theta_0]})^\top G^\top \operatorname{diag}(G\operatorname{vec}(u^{[\theta_0]}))^{-1/2}, \quad (4.65)$$

with $J_{u,p}$ *defined in* (4.41) *and* $\Sigma_p = \operatorname{diag}(\operatorname{vec}(p)) - \operatorname{vec}(p)\operatorname{vec}(p)^\top$. *Furthermore, provided that* $(\theta,u',p') \mapsto \Sigma_{G,u',p'}^{[\theta]}$ *is continuous at* (θ_0,u,p), *we have that* $\Sigma_{G,u^{[n]},p^{[n]}}^{[\hat{\theta}^{[n]}]} \xrightarrow{\mathrm{P}} \Sigma_{G,u,p}^{[\theta_0]}$ *in* $\mathbb{R}^{q\times q}$, *where* $p^{[n]}$ *is defined in* (4.38).

Note that the eigenvalues $\lambda_1,\ldots,\lambda_q$ of $\Sigma_{G,u,p}^{[\theta_0]}$ are not in general equal to 0 or 1 so that $\mathcal{L}_G$ does not have a chi-square distribution in general. The previous proposition however suggests that a goodness-of-fit test based on $S_G^{[n]}$ in (4.64) could be carried out in practice as follows:

1. For the hypothesized parametric family of copula pmfs $\mathcal{J}$ in (4.47), estimate θ_0 by $\hat{\theta}^{[n]}$, where $\hat{\theta}^{[n]}$ is one the three method-of-moments estimators in (4.50) or the maximum pseudo-likelihood estimator in (4.53), and compute $S_G^{[n]}$ in (4.64).

2. Compute the eigenvalues $\hat{\lambda}_1,\ldots,\hat{\lambda}_q$ of $\Sigma_{G,u^{[n]},p^{[n]}}^{[\hat{\theta}^{[n]}]}$ which estimate the eigenvalues $\lambda_1,\ldots,\lambda_q$ of $\Sigma_{G,u,p}^{[\theta_0]}$ in (4.65).

3. For some large integer M, compute $S_G^{[n],l} = \sum_{k=1}^{q} \hat{\lambda}_k Z_{kl}^2$, $l \in \{1,\ldots,M\}$, for i.i.d. $N(0,1)$-distributed random variables Z_{kl}, $k \in \{1,\ldots,q\}$, $l \in \{1,\ldots,M\}$, and estimate the p-value of the test as

$$\frac{1}{M} \sum_{l=1}^{M} \mathbf{1}(S_G^{[n],l} \geq S_G^{[n]}).$$

As a proof of concept, we shall focus on the case when $\hat{\theta}^{[n]}$ is $\theta_{\rho}^{[n]}$ in (4.50), the estimator of θ_0 based on the inversion of Yule's coefficient. From Corollary 4.5.6, we then know that, in this case,

$$V_{u,p}^{[\theta_0]} = \frac{1}{g'_{\rho}(\theta_0)} \dot{\Upsilon}(u)^{\top} J_{u,p}.$$

Standard calculations show that

$$\dot{\Upsilon}(u) = \frac{12}{\sqrt{(r+1)(s+1)(r-1)(s-1)}} \big((i-1)(j-1)\big)_{(i,j)\in I_{r,s}}$$

and that

$$g'_{\rho}(\theta_0) = \frac{12}{\sqrt{(r+1)(s+1)(r-1)(s-1)}} \sum_{(i,j)\in I_{r,s}} (i-1)(j-1)\, \dot{u}_{ij}^{[\theta_0]}.$$

The previous formulas can be used to derive the expression of the covariance matrix $\Sigma_{G,u,p}^{[\theta_0]}$ in (4.65) in terms of $\dot{u}^{[\theta_0]}$. When the hypothesised family $\mathcal{J}$ in (4.47) is defined from a parametric copula family as in (4.48), $\dot{u}^{[\theta_0]}$ can be obtained by differentiating (4.48). This involves standard calculations for parametric copula families with explicit expressions for C_{θ}. Additionally, this differentiation can also be performed for implicit copula families, such as the normal or the t copula, using the expressions provided in Kojadinovic and Yan (2011).

4.6.2 Monte Carlo Experiments

To evaluate the finite-sample performance of the asymptotic chi-square goodness-of-fit tests described in the previous section, we consider data-generating scenarios similar to those in Section 4.5.3. Table 4.6.1 (respectively, Table 4.6.2 and Table 4.6.3) presents rejection percentages of the test based on $S^{[n]}$ in (4.63) and Yule's coefficient, computed from 1000 random samples of size $n \in \{100,500,1000\}$ generated from pmfs whose copula pmf is of the form (4.48) with C_{θ} being the Clayton (respectively, Gumbel-Hougaard, Frank) copula with a Kendall's tau in $\{0.33,0.66\}$ and whose margins align with the marginal scenarios listed in Section 4.5.3.

In all tables, the columns 'Cl' (respectively, 'GH', 'F') report rejection percentages when $\mathcal{J}$ in H_0 in (4.61) is constructed from a Clayton (respectively, Gumbel-Hougaard, Frank) copula. The integer in parentheses next to each rejection percentage represents the number of numerical issues (either related

Table 4.6.1: For $(r,s) \in \{(3,3),(3,5),(5,5)\}$, rejection percentages of the goodness-of-fit test based on $S^{[n]}$ in (4.63) and Yule's coefficient calculated from 1000 random samples of size $n \in \{100,500,1000\}$ generated, as explained in Section 4.6.2, from pmfs whose copula pmf is of the form (4.48) with C_θ the Clayton copula with a Kendall's tau in $\{0.33,0.66\}$. The column 'm' gives the marginal scenario. The integer between parentheses is the number of numerical issues encountered out of 1000 executions.

				$\tau = 0.33$			$\tau = 0.66$		
r	s	m	n	Cl	GH	F	Cl	GH	F
3	3	1	100	3.6 (0)	30.5 (0)	11 (0)	1.7 (0)	21.4 (0)	7.1 (0)
			500	4.9 (0)	93.8 (0)	59.3 (0)	2.9 (0)	97.7 (0)	78.7 (0)
			1000	5.7 (0)	100 (0)	89.8 (0)	3.3 (0)	100 (0)	99 (0)
		2	100	3.4 (0)	17.5 (0)	9.4 (0)	0.5 (0)	13.9 (0)	3.5 (0)
			500	3.5 (0)	85.6 (0)	46.7 (0)	2.8 (0)	96 (0)	69.9 (0)
			1000	4 (0)	99.7 (0)	77.4 (0)	3.8 (0)	99.9 (0)	97.7 (0)
		3	100	2.3 (0)	21.1 (0)	8.9 (0)	0.8 (0)	15.7 (0)	5.2 (0)
			500	4.8 (0)	90.7 (0)	52.1 (0)	1.7 (0)	96.4 (0)	69.6 (0)
			1000	5.7 (0)	99.7 (0)	87 (0)	2.8 (0)	100 (0)	97.4 (0)
3	5	1	100	1.9 (0)	18.2 (0)	8.4 (0)	1.3 (0)	7.6 (0)	0.3 (0)
			500	5.4 (0)	98.1 (0)	69.7 (0)	3.3 (0)	94.5 (0)	80.1 (0)
			1000	4.5 (0)	100 (0)	96.4 (0)	2.8 (0)	99.8 (0)	99.5 (0)
		2	100	0.4 (0)	6.8 (0)	2 (0)	0.8 (0)	1.8 (0)	0 (0)
			500	4 (0)	87.1 (0)	50.2 (0)	3 (0)	89.4 (0)	32.3 (0)
			1000	5.1 (0)	99.8 (0)	85.2 (0)	3.9 (0)	99.1 (0)	90.3 (0)
		3	100	0.3 (0)	1.9 (0)	0.3 (0)	0.2 (0)	0.1 (0)	0.1 (0)
			500	2.1 (0)	65.5 (0)	25.5 (0)	2.8 (0)	64.2 (0)	10.5 (0)
			1000	3.8 (0)	97.3 (0)	67.4 (0)	3.7 (0)	93.5 (0)	72.6 (0)
5	5	1	100	0.4 (0)	7.5 (0)	1 (0)	1.2 (2)	8.1 (2)	0.5 (1)
			500	4.1 (0)	98.7 (0)	78 (0)	2.5 (0)	88.2 (0)	56.9 (0)
			1000	5.2 (0)	100 (0)	99.3 (0)	3.9 (0)	98.9 (0)	91.2 (0)
		2	100	0.1 (0)	0.8 (0)	0.2 (0)	0.5 (0)	4.2 (0)	0.1 (1)
			500	4.1 (0)	88 (0)	46.3 (0)	3.8 (0)	94.7 (0)	48.5 (0)
			1000	4.5 (0)	99.6 (0)	87.7 (0)	5.5 (0)	99.5 (0)	89.2 (0)
		3	100	0 (0)	0 (0)	0 (0)	0.2 (0)	0 (0)	0 (0)
			500	0.3 (0)	29.1 (0)	3.9 (0)	0.6 (0)	38.3 (0)	6.7 (0)
			1000	0.8 (0)	77.5 (0)	26.1 (0)	3.5 (0)	83.6 (0)	68.1 (0)

to the convergence of the IPFP, numerical root finding, or the necessary eigenvalue decomposition) out of 1000 executions.

An inspection of the tables reveals that such numerical issues were very rare. In terms of rejection percentages, the tests were never excessively liberal and tended to be conservative in scenarios where the probability of

Table 4.6.2: For $(r,s) \in \{(3,3),(3,5),(5,5)\}$, rejection percentages of the goodness-of-fit test based on $S^{[n]}$ in (4.63) and Yule's coefficient computed from 1000 random samples of size $n \in \{100,500,1000\}$ generated, as explained in Section 4.6.2, from pmfs whose copula pmf is of the form (4.48) with C_θ the Gumbel–Hougaard copula with a Kendall's tau in $\{0.33,0.66\}$. The column 'm' gives the marginal scenario. The integer between parentheses reports the number of numerical issues encountered out of 1000 executions.

				$\tau = 0.33$			$\tau = 0.66$		
r	s	m	n	Cl	GH	F	Cl	GH	F
3	3	1	100	25.3 (0)	3.3 (0)	6.7 (0)	30.5 (0)	0.6 (0)	2.9 (0)
			500	92.7 (0)	5 (0)	19.9 (0)	98.8 (0)	3.1 (0)	17.7 (0)
			1000	99.9 (0)	4.1 (0)	36.9 (0)	100 (0)	3.1 (0)	38.7 (0)
		2	100	20.8 (0)	3.1 (0)	5.6 (0)	20.8 (0)	0.8 (0)	3 (0)
			500	88.1 (0)	4.5 (0)	15.6 (0)	96.8 (0)	2.4 (0)	10.9 (0)
			1000	99.7 (0)	4.2 (0)	27.2 (0)	100 (0)	4.2 (0)	33.2 (0)
		3	100	20 (0)	2.6 (0)	4.7 (0)	23.1 (0)	0.8 (0)	1.6 (0)
			500	90.8 (0)	4.2 (0)	15.8 (0)	98.2 (0)	1.4 (0)	10.8 (0)
			1000	99.3 (0)	4.3 (0)	35.9 (0)	100 (0)	3.7 (0)	28 (0)
3	5	1	100	21.1 (0)	2 (0)	2.5 (0)	30.3 (0)	0.1 (0)	0.1 (0)
			500	97.4 (0)	5.4 (0)	24.5 (0)	99.7 (0)	2.1 (0)	19.6 (0)
			1000	99.9 (0)	5.8 (0)	52.1 (0)	100 (0)	2 (0)	50.3 (0)
		2	100	12.2 (0)	0.6 (0)	1.5 (0)	16.8 (0)	0 (0)	0.2 (0)
			500	87.9 (0)	5.4 (0)	12.5 (0)	98.2 (0)	1.1 (0)	9.6 (0)
			1000	99.7 (0)	4.8 (0)	33.6 (0)	100 (0)	2.5 (0)	28.7 (0)
		3	100	3.5 (0)	0.3 (0)	0.1 (0)	2.3 (0)	0 (0)	0 (0)
			500	68.6 (0)	1.7 (0)	9.7 (0)	91.6 (0)	1 (0)	1.5 (0)
			1000	96.2 (0)	3.6 (0)	21.6 (0)	99.8 (0)	0.7 (0)	15.7 (0)
5	5	1	100	10.1 (0)	0.1 (0)	0.2 (0)	36 (0)	0.1 (0)	0.3 (0)
			500	99.2 (0)	4.3 (0)	30.5 (0)	99.3 (0)	0.7 (0)	32.7 (0)
			1000	100 (0)	4.3 (0)	67.7 (0)	100 (0)	2.6 (0)	83.5 (0)
		2	100	3.2 (0)	0 (0)	0.1 (0)	15.8 (0)	0 (0)	0.1 (0)
			500	89.6 (0)	3.6 (0)	12.1 (0)	98.2 (0)	1.2 (0)	15.1 (0)
			1000	100 (0)	5.6 (0)	39.6 (0)	100 (0)	2.3 (0)	63.7 (0)
		3	100	0.2 (1)	0 (0)	0 (0)	0.2 (0)	0 (0)	0 (0)
			500	37.2 (0)	0.4 (0)	1.1 (0)	86.6 (0)	0.3 (0)	2.2 (0)
			1000	82.3 (0)	1.3 (0)	8.8 (0)	99.4 (0)	1.8 (0)	19.9 (0)

zero counts in the contingency tables of the generated samples was high. Correspondingly, the tests generally demonstrated good power when they were not overly conservative.

The lowest empirical levels and powers in Tables 4.6.1, 4.6.2 and 4.6.3 are frequently observed for the third marginal scenario. For example, for

Table 4.6.3: For $(r,s) \in \{(3,3),(3,5),(5,5)\}$, rejection percentages of the goodness-of-fit test based on $S^{[n]}$ in (4.63) and Yule's coefficient computed from 1000 random samples of size $n \in \{100,500,1000\}$ generated, as explained in Section 4.6.2, from pmfs whose copula pmf is of the form (4.48) with C_θ the Frank copula with a Kendall's tau in $\{0.33,0.66\}$. The column 'm' gives the marginal scenario. The integer between parentheses reports the number of numerical issues encountered out of 1000 executions.

				$\tau = 0.33$			$\tau = 0.66$		
r	s	m	n	Cl	GH	F	Cl	GH	F
3	3	1	100	15.1 (0)	6.5 (0)	2.4 (0)	14.3 (0)	2.4 (0)	2.3 (0)
			500	61.7 (0)	22.4 (0)	4.7 (0)	78.1 (0)	12.4 (0)	2.3 (0)
			1000	90.9 (0)	40 (0)	4.9 (0)	98.5 (0)	32.9 (0)	3.4 (0)
		2	100	10.4 (0)	2.9 (0)	1.9 (0)	9.2 (0)	1.2 (0)	1.7 (0)
			500	52.6 (0)	15.5 (0)	5.1 (0)	67.7 (0)	8.4 (0)	2.4 (0)
			1000	87.2 (0)	32.8 (0)	5.2 (0)	97.9 (0)	24.3 (0)	2.9 (0)
		3	100	9.4 (0)	3.8 (0)	1.9 (0)	8.4 (0)	1 (0)	0.6 (0)
			500	56 (0)	18.7 (0)	5.6 (0)	71.3 (0)	7.2 (0)	1.6 (0)
			1000	88.1 (0)	35.4 (0)	2.7 (0)	97.8 (0)	20.6 (0)	2.2 (0)
3	5	1	100	11.7 (0)	3.4 (0)	1.5 (0)	20.8 (0)	1.1 (0)	0 (0)
			500	77.9 (0)	28.3 (0)	3.9 (0)	96.9 (0)	28.9 (0)	1.8 (0)
			1000	98.1 (0)	58.4 (0)	4.3 (0)	99.9 (0)	62.9 (0)	2.3 (0)
		2	100	5.9 (0)	1.3 (0)	1 (0)	13.1 (0)	0 (0)	0.1 (0)
			500	55.2 (0)	16.2 (0)	4.3 (0)	81 (0)	12.6 (0)	0.9 (0)
			1000	88 (0)	39.5 (0)	4.1 (0)	99.2 (0)	46.8 (0)	2.9 (0)
		3	100	2 (0)	0.7 (0)	0.1 (0)	1.8 (0)	0 (0)	0 (0)
			500	35.9 (0)	10 (0)	1.7 (0)	70.4 (0)	11.3 (0)	1.8 (0)
			1000	76.7 (0)	30.8 (0)	4 (0)	95.9 (0)	30 (0)	2.2 (0)
5	5	1	100	3.5 (0)	0.4 (0)	0.3 (0)	20.8 (0)	0.3 (0)	0 (0)
			500	83.4 (0)	35 (0)	4.5 (0)	99.6 (0)	39.2 (0)	1.6 (0)
			1000	99.4 (0)	75.4 (0)	5.4 (0)	100 (0)	85.9 (0)	1.6 (0)
		2	100	1.9 (0)	0.1 (0)	0 (0)	10.5 (0)	0.1 (0)	0 (0)
			500	53.6 (0)	15.9 (0)	3.6 (0)	98 (0)	19 (0)	1.1 (0)
			1000	89.8 (0)	48.3 (0)	3.8 (0)	100 (0)	74.5 (0)	3 (0)
		3	100	0 (1)	0 (0)	0 (0)	0.4 (0)	0 (0)	0 (0)
			500	13.2 (0)	2.3 (0)	0.2 (0)	74.5 (0)	3 (0)	1 (0)
			1000	36.8 (0)	8.8 (0)	0.2 (0)	99.6 (0)	26.3 (0)	1.7 (0)

that marginal scenario, C_θ the Clayton copula with a Kendall's tau of 0.66, $(r,s) = (5,5)$ and $n = 500$, a realization of the empirical copula pmf $u^{[n]}$

in (4.40) multiplied by n and rounded to the nearest integer is:

$$\begin{bmatrix} 80 & 16 & 3 & 0 & 0 \\ 18 & 53 & 21 & 9 & 0 \\ 2 & 18 & 40 & 27 & 12 \\ 0 & 9 & 19 & 47 & 26 \\ 0 & 3 & 17 & 17 & 62 \end{bmatrix}.$$

Note that, as already mentioned in Section 4.5.2, it could be interpreted as observed counts from the unknown copula pmf u. The low counts in the lower-left and upper right-corners might be the cause of the conservative behavior of the goodness-of-fit tests based on $S^{[n]}$ in (4.63). For that reason, in the next experiment, we focused on the $(r,s)=(5,5)$ case and the third marginal scenario, and considered groupings according to the following matrix:

$$\begin{bmatrix} & & 1 & 2 & 2 \\ & & 1 & 2 & 2 \\ 3 & 3 & & & \\ 4 & 4 & & & \\ 4 & 4 & & & \end{bmatrix} \tag{4.66}$$

where elements with the same integer are to be regrouped and from which we can form the 25 by 25 grouping matrix G to be used in the statistic $S_G^{[n]}$ in (4.64). Notice that the resulting groupings may not always guarantee that all aggregated counts are above 5. This will certainly not be true in general for $n=100$. The rejection percentages of the goodness-of-fit test based on $S_G^{[n]}$ and Yule's coefficient are presented in Table 4.6.4. A comparison with the horizontal blocks of Tables 4.6.1, 4.6.2 and 4.6.3 corresponding to $(r,s)=(5,5)$ and the third marginal scenario shows a clear improvement of the empirical levels and an increase of the powers when $n \geq 500$.

In a last experiment, we focused on the $(r,s)=(10,10)$ case and second marginal scenario listed in Section 4.5.3, and decided to form the grouping

Table 4.6.4: For $(r,s)=(5,5)$, rejection percentages of the goodness-of-fit test based on $S_G^{[n]}$ in (4.64) and Yule's coefficient with G formed according to (4.66) computed from 1000 random samples of size $n \in \{100,500,1000\}$ generated from a pmf whose copula pmf is of the form (4.48) with C_θ the Clayton (Cl), Gumbel–Hougaard (GH) or Frank copula (F) with a Kendall's tau in $\{0.33,0.66\}$ and whose margins are binomial with parameters 4 and 0.5.

		$C_\theta = $ Cl			$C_\theta = $ GH			$C_\theta = $ F		
τ	n	Cl	GH	F	Cl	GH	F	Cl	GH	F
0.33	100	0.1	2.3	0.3	1.1	0.1	0.2	0.8	0.6	0.1
	500	3.2	70.4	28.5	65.4	4.0	18.0	37.8	22.0	3.9
	1000	4.7	97.0	67.1	93.4	4.6	34.6	63.5	44.9	4.9
0.66	100	0.2	18.1	2.7	1.6	0.3	0.1	1.6	0.9	0.0
	500	3.5	100.0	95.0	86.2	2.2	27.8	83.8	53.2	2.6
	1000	4.7	100.0	100.0	98.8	4.2	77.9	99.6	84.7	4.4

matrix G according to the following matrix:

$$\begin{bmatrix} 1 & 1 & 1 & 2 & 2 & 2 & 2 & 3 & 3 & 3 \\ 1 & 1 & 1 & 2 & 2 & 2 & 2 & 3 & 3 & 3 \\ 1 & 1 & 1 & 2 & 2 & 2 & 2 & 3 & 3 & 3 \\ 4 & 4 & 4 & & & & & 5 & 5 & 5 \\ 4 & 4 & 4 & & & & & 5 & 5 & 5 \\ 4 & 4 & 4 & & & & & 5 & 5 & 5 \\ 4 & 4 & 4 & & & & & 5 & 5 & 5 \\ 6 & 6 & 6 & 7 & 7 & 7 & 7 & 8 & 8 & 8 \\ 6 & 6 & 6 & 7 & 7 & 7 & 7 & 8 & 8 & 8 \\ 6 & 6 & 6 & 7 & 7 & 7 & 7 & 8 & 8 & 8 \end{bmatrix}. \tag{4.67}$$

The rejection percentages are reported in Table 4.6.5 for $n \in \{500,1000,2000,4000\}$. The results seem to confirm that the goodness-of-fit tests based on $S_G^{[n]}$ in (4.64) can be well-behaved in many scenarios provided groupings are performed to avoid "very low counts" in the matrix $nu^{[n]}$.

4.6.3 Chi-Square Tests Based on a Semi-Parametric Bootstrap

When $(X_1,Y_1),\ldots,(X_n,Y_n)$ is a random sample, the goodness-of-fit tests based $S_G^{[n]}$ in (4.64) can also be carried out using an appropriate adaptation of the so-called parametric bootstrap (see, e.g., Stute et al., 1993; Genest and Rémillard, 2008). Specifically, in our case, the latter could be called a semi-parametric bootstrap. The testing procedure that we propose is as follows:

Table 4.6.5: For $(r,s)=(10,10)$, rejection percentages of the goodness-of-fit test based on $S_G^{[n]}$ in (4.64) and Yule's coefficient with G formed according to (4.67) computed from 1000 random samples of size $n\in\{500,1000,2000\}$ generated from a pmf whose copula pmf is of the form (4.48) with C_θ the Clayton (Cl), Gumbel–Hougaard (GH) or Frank copula (F) with a Kendall's tau in $\{0.33,0.66\}$ and whose margins are as in the second marginal scenario of Section 4.5.3.

		$C_\theta=$ Cl			$C_\theta=$ GH			$C_\theta=$ F		
τ	n	Cl	GH	F	Cl	GH	F	Cl	GH	F
0.33	500	6.2	65.6	28.0	73.4	7.5	15.6	39.8	11.1	8.5
	1000	5.8	97.8	62.4	98.7	7.3	26.2	68.6	19.4	6.0
	2000	4.1	100.0	94.1	100.0	6.8	45.6	95.6	42.1	5.2
0.66	500	2.6	90.9	49.6	93.8	2.5	18.5	56.2	5.3	2.8
	1000	4.4	99.9	97.6	100.0	3.7	40.5	97.0	24.1	3.7
	2000	4.1	100.0	100.0	100.0	4.2	76.9	100.0	64.4	3.8

1. For the hypothesized parametric family of copula pmfs $\mathcal{J}$ in (4.47), estimate θ_0 by $\hat{\theta}^{[n]}$, where $\hat{\theta}^{[n]}$ is one the three method-of-moments estimators in (4.50) or the maximum pseudo-likelihood estimator in (4.53), and compute $S_G^{[n]}$ in (4.64).
2. Let $p^{[n,1]}$ and $p^{[n,2]}$ denote the margins of $p^{[n]}$ in (4.38), and, using (4.14) with $a=p^{[n,1]}$ and $b=p^{[n,2]}$, form $p^{[n,\hat{\theta}^{[n]}]}=\mathcal{I}_{p^{[n,1]},p^{[n,2]}}(u^{[\hat{\theta}^{[n]}]})$, which is a consistent semi-parametric estimate of the pmf of (X,Y) under H_0 in (4.61).
3. For some large integer M, repeat the following steps for $l\in\{1,\ldots,M\}$:
   (a) Generate a random sample $(X_1^{(l)},Y_1^{(l)}),\ldots,(X_n^{(l)},Y_n^{(l)})$ from $p^{[n,\hat{\theta}^{[n]}]}$.
   (b) From the generated sample, compute the version $p^{[n],l}$ of $p^{[n]}$ in (4.38), the version $u^{[n],l}=\mathcal{U}(p^{[n],l})$ of $u^{[n]}$, where $\mathcal{U}$ is defined in (4.29), and the version $\hat{\theta}^{[n],l}$ of $\hat{\theta}^{[n]}$.
   (c) Form an approximate realization under H_0 of $S_G^{[n]}$ in (4.64) as
   $$S_G^{[n],l}=n\|\mathrm{diag}(G\,\mathrm{vec}(u^{[\hat{\theta}^{[n]},l]}))^{-1/2}G\,\mathrm{vec}(u^{[n],l}-u^{[\hat{\theta}^{[n]},l]})\|_2^2.$$
4. Finally, estimate the p-value of the test as
   $$\frac{1}{M}\sum_{l=1}^{M}\mathbf{1}(S_G^{[n],l}\geq S_G^{[n]}).$$

Remark 4.6.3. From a theoretical perspective, to attempt to show the null asymptotic validity of the proposed procedure, one could try to start

from the work of Genest and Rémillard (2008). From an empirical perspective, studying its finite-sample behavior is unfortunately substantially more computationally costly than studying the finite-sample performance of the asymptotic test described in Section 4.6.1. □

4.7 Data Example

We provide a brief demonstration of the proposed methodology using a bivariate data set previously analyzed in Goodman (1979). The discrete random variables involved are X, which represents the occupational status of a British male subject, and Y, which represents the occupational status of the subject's father. Both variables can take on $r = s = 8$ distinct ordered values, represented by the first eight integers. The reasoning for considering these values as ordered is not explained in Goodman (1979), and the latter refers the reader to Miller (1960) for further details, among others. It appears that there are no reasons to assume that any values within $I_{r,s}$ cannot be realizations of (X,Y). Hence, it is reasonable to assume that the unknown pmf p of (X,Y) is supported on $I_{r,s}$, implying that the decomposition in (4.31) is valid. The random vector (X,Y) was observed for $n = 3498$ subjects. The resulting bivariate contingency table can be obtained in R by executing `data(occupationalStatus)`, and is shown below for convenience:

$$\begin{bmatrix} 50 & 19 & 26 & 8 & 7 & 11 & 6 & 2 \\ 16 & 40 & 34 & 18 & 11 & 20 & 8 & 3 \\ 12 & 35 & 65 & 66 & 35 & 88 & 23 & 21 \\ 11 & 20 & 58 & 110 & 40 & 183 & 64 & 32 \\ 2 & 8 & 12 & 23 & 25 & 46 & 28 & 12 \\ 12 & 28 & 102 & 162 & 90 & 554 & 230 & 177 \\ 0 & 6 & 19 & 40 & 21 & 158 & 143 & 71 \\ 0 & 3 & 14 & 32 & 15 & 126 & 91 & 106 \end{bmatrix}. \tag{4.68}$$

To estimate the unknown copula pmf $u = \mathcal{U}(p)$ of (X,Y), we first computed $p^{[n]}$ in (4.39) from the above contingency table and then the empirical copula pmf $u^{[n]}$ using (4.40). The latter is represented in Figure 4.7.1 under the form of a *ballon plot* created using the R package `ggpubr` (Kassambara, 2023). To complement Figure 4.7.1, we also provide the matrix $nu^{[n]}$ rounded to

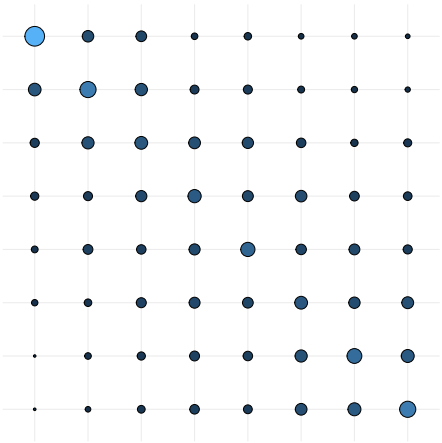

Figure 4.7.1: Ballon plot of the empirical copula pmf $u^{[n]}$ computed using (4.39) and (4.40) from the contingency table given in (4.68).

the nearest integer,

$$\begin{bmatrix} 253 & 70 & 58 & 14 & 21 & 8 & 8 & 4 \\ 88 & 160 & 82 & 35 & 36 & 17 & 12 & 7 \\ 38 & 80 & 90 & 74 & 66 & 42 & 20 & 26 \\ 28 & 37 & 65 & 99 & 61 & 71 & 44 & 32 \\ 16 & 46 & 42 & 64 & 118 & 55 & 60 & 37 \\ 13 & 22 & 48 & 62 & 58 & 91 & 68 & 76 \\ 0 & 15 & 28 & 47 & 42 & 80 & 130 & 94 \\ 0 & 8 & 24 & 43 & 34 & 73 & 95 & 160 \end{bmatrix}, \tag{4.69}$$

which, as already mentioned, could be interpreted as observed counts from the unknown copula pmf u. Estimates of Yule's coefficient ρ in (4.33) as well as the gamma coefficient γ and the tau coefficient τ in (4.36) can be computed using (4.46) and are $\rho^{[n]} \simeq 0.63$, $\gamma^{[n]} \simeq 0.56$ and $\tau^{[n]} \simeq 0.5$, respectively. The latter seems to indicate a moderately strong dependence between X and Y.

We now proceed to the parametric modelling of u. Initially, we fitted eight one-parameter families $\mathcal{J}$ in (4.47) constructed using (4.48), where C_θ is either a Clayton, Gumbel–Hougaard, Frank, Plackett, survival Clayton, survival Gumbel–Hougaard, or survival Joe copula (refer to, e.g., Hofert et al., 2018, Section 2.5 for the definition of a survival copula). Estimates obtained via the method of moments and maximum pseudo-likelihood are presented in Table 4.7.1. The last column of this table also includes the scaled maximised negative log-pseudo-likelihood $-\bar{L}^{[n]}(\theta^{[n]})/n$, where $\bar{L}$ is defined in (4.52). These results suggest that the model based on a survival

Table 4.7.1: Estimates of the parameter θ for families $\mathcal{J}$ in (4.47) constructed via (4.48), where C_θ is either a Clayton (Cl), Gumbel–Hougaard (GH), Frank (F), Plackett (P), survival Clayton (sCl), survival Gumbel–Hougaard (sGH) or survival Joe (sJ) copula. The first three columns give the three method-of-moment estimates defined in (4.50). The fourth and fifth colums report the maximum pseudo-likelihood estimate in (4.53) and the value of $-\bar{L}^{[n]}(\theta^{[n]})/n$, respectively, where $\bar{L}$ is defined in (4.52).

C_θ	$\theta_\rho^{[n]}$	$\theta_\gamma^{[n]}$	$\theta_\tau^{[n]}$	$\theta^{[n]}$	$-\bar{L}^{[n]}(\theta^{[n]})/n$
Cl	1.712	1.724	1.722	1.548	3.906
GH	1.865	1.861	1.863	1.789	3.940
J	2.591	2.596	2.592	2.070	3.998
F	4.886	4.945	4.968	5.001	3.914
P	9.147	8.915	8.961	8.717	3.909
sCl	1.712	1.724	1.722	1.208	3.987
sGH	1.865	1.861	1.863	1.861	3.896
sJ	2.591	2.596	2.592	2.393	3.909

Table 4.7.2: Results of the goodness-of-fit tests based on $S_G^{[n]}$ in (4.64) and Yule's coefficient with G formed such that the four values in the lower-left and upper-right corners of the copula p.m.f.s are grouped. The hypothesized family $\mathcal{J}$ in (4.61) is constructed via (4.48), where C_θ is either a Clayton (Cl), Gumbel–Hougaard (GH), Frank (F), Plackett (P), survival Clayton (sCl), survival Gumbel–Hougaard (sGH) or survival Joe (sJ) copula. The first row gives the values of $S_G^{[n]}$. The second (resp. third) row reports the p-values obtained via the asymptotic procedure (resp. semi-parametric bootstrap) described in Section 4.6.1 (resp. Section 4.6.3) with $M = 10^4$.

	Cl	GH	J	F	P	sCl	sGH	sJ
$S_G^{[n]}$	260.4	523.1	1374.9	279.9	245.4	1200.2	156.4	303.0
Asymptotic	0.000	0.000	0.000	0.000	0.000	0.000	0.004	0.000
Semi-p. boot.	0.000	0.000	0.000	0.000	0.000	0.000	0.011	0.000

Gumbel–Hougaard copula provides the best fit. This observation might not be unexpected since survival Gumbel–Hougaard copulas exhibit lower-tail dependence, and an examination of Figure 4.7.1 and the "observed counts" in (4.69) appears to indicate such "lower-tail" dependence in the upper left corner.

To further evaluate the fit of the previously mentioned eight models, we conducted goodness-of-fit tests based on $S_G^{[n]}$ in (4.64) and Yule's coefficient, with G structured so that, following (4.69), the four values in the lower-left and upper-right corners of the copula pmfs are grouped. Table 4.7.2 reports the p-values, calculated using both the asymptotic method from Section 4.6.1 and the semi-parametric bootstrap approach from Section 4.6.3 with $M = 10^4$. As shown, all models are rejected at the 2% significance level,

but the model based on the survival Gumbel–Hougaard copula family is the least rejected. A close examination of Figure 4.7.1 and (4.69) suggests asymmetry with respect to the diagonal in the unknown copula pmf u. This implies that models incorporating suitable non-exchangeable generalisations of the survival Gumbel–Hougaard copula family could offer a better fit. This will be explored in future research.

4.8 Concluding Remarks

Inspired by the seminal work of Geenens (2020), we explored a copula-like modelling approach for discrete bivariate distributions hinging upon I-projections on Fréchet classes and the IPFP. The proposed methodology find its roots in the copula-like decomposition of bivariate pmfs stated in Proposition 4.4.1. We focused on discrete bivariate distributions with rectangular supports, proposing both nonparametric and parametric estimation procedures, along with goodness-of-fit tests for the underlying copula pmf Asymptotic results were provided based on a differentiability result for I-projections on Fréchet classes, which may be of independent interest. Monte Carlo experiments were conducted to study the finite-sample performance of some inference procedures, and the methodology was illustrated with a data example. A discussion of the resulting future works is postponed to Chapter 7. Therein, among various theoretical extensions, a possible application to rainfall data shall also be proposed.

Throughout this chapter we have mentioned in numerous remarks the important connections of our work with the theory of ϕ-divergence estimators, which are based on the concept of ϕ-projections. The latter are a generalisation of the I-projections that proved to be a fundamental building block of the results in this chapter. It is then natural to investigate the extension of our differentiability result in Section 4.2.4 to this larger class of projections and explore its applications when studying the asymptotics of ϕ-divergence estimators. The next chapter is mainly theoretical and deals with this line of our research.

Appendix

4.A Proof of Proposition 4.2.6

Proof of Proposition 4.2.6. The proof is based on the implicit function theorem (see, e.g., Fitzpatrick, 2009, Theorem 17.6, p 450). Before applying this result, we need to define the underlying function and verify a certain number of related assumptions.

Since it is assumed that $\Gamma_{a,b,T} \neq \emptyset$, from Proposition 4.2.2, for any $x \in \Gamma_T$, $\mathcal{I}_{a,b}(x) = y_x^* = \arg\inf_{y \in \Gamma_{a,b}} D(y\|x)$ with $\operatorname{supp}(y_x^*) = T$. It follows that, for any $x \in \Gamma_T$, $\mathcal{I}_{a,b}(x) = \arg\inf_{y \in \Gamma_{a,b,T}} D(y\|x)$ and thus

$$\operatorname{vec}_A(\mathcal{I}_{a,b}(x)) = \operatorname*{arginf}_{z \in \Lambda_{a,b,A,T}} H(z\|\operatorname{vec}_B(x)), \tag{4.70}$$

where H is defined in (4.28) and $\Lambda_{a,b,A,T}$ in (4.27) is, by assumption, an open subset of $\mathbb{R}^{|A|}$. Notice that, from (4.13), for any $(z,w) \in \Lambda_{a,b,A,T} \times \Lambda_{B,T}$,

$$H(z\|w) = \sum_{(i,j) \in T} c(z)_{ij} \log \frac{c(z)_{ij}}{d(w)_{ij}} \tag{4.71}$$

and $\Lambda_{B,T}$ in (4.25) is an open subset of $\mathbb{R}^{|B|}$ under the considered assumptions on B. Next, let $F : \Lambda_{a,b,A,T} \times \Lambda_{B,T} \to \mathbb{R}^{|A|}$ be defined by

$$F(z,w) = \partial_1 H(z\|w), \qquad z \in \Lambda_{a,b,A,T} \subset (0,1)^{|A|}, w \in \Lambda_{B,T} \subset (0,1)^{|B|}.$$

From (4.71), F is differentiable at any $(z,w) \in \Lambda_{a,b,A,T} \times \Lambda_{B,T}$. Furthermore, from the definition of F, (4.70) and first-order necessary optimality conditions, we have that

$$F(\operatorname{vec}_A(\mathcal{I}_{a,b}(x)), \operatorname{vec}_B(x)) = 0_{\mathbb{R}^{|A|}}, \qquad \text{for all } x \in \Gamma_T. \tag{4.72}$$

Finally, to be able to apply the implicit function theorem, we shall verify that

$$\det[\partial_1 F(\operatorname{vec}_A(\mathcal{I}_{a,b}(x)), \operatorname{vec}_B(x))] \neq 0, \qquad \text{for all } x \in \Gamma_T. \tag{4.73}$$

Consider the set $\Lambda_{T,T} \subset (0,1)^{|T|}$ defined in (4.25) with $B = T$ and let $\tilde{D}$ be the map from $\Lambda_{T,T} \times \Lambda_{T,T}$ to $[0,\infty)$ defined by $\tilde{D}(s\|s') = D(\mathrm{vec}_T^{-1}(s)\|\mathrm{vec}_T^{-1}(s'))$, $s, s' \in \Lambda_{T,T}$. Then, from (4.13), for any $s, s' \in \Lambda_{T,T}$,

$$\tilde{D}(s\|s') = \sum_{i=1}^{|T|} s_i \log \frac{s_i}{s'_i}.$$

Lemma 4.A.1 below then implies that, for any $r \in \Lambda_{T,T}$, $\tilde{D}(\cdot\|r)$ is strongly convex with constant 1 on $\Lambda_{T,T}$ in the sense of Nesterov (2004, Section 2.1.3). From Theorem 2.1.9 in the previous reference, the latter is equivalent to the fact that, for any $r \in \Lambda_{T,T}$, $s, s' \in \Lambda_{T,T}$ and $t \in [0,1]$,

$$t\tilde{D}(s\|r) + (1-t)\tilde{D}(s'\|r) \geq \tilde{D}(ts + (1-t)s'\|r) + t(1-t)\frac{1}{2}\|s - s'\|_2^2. \quad (4.74)$$

From (4.28), we further have that, for any $z \in \Lambda_{a,b,A,T}$ and $w \in \Lambda_{B,T}$,

$$H(z\|w) = D(c(z)||d(w)) = \tilde{D}(\mathrm{vec}_T \circ c(z)||\mathrm{vec}_T \circ d(w)).$$

Fix $w \in \Lambda_{B,T}$, $z, z' \in \Lambda_{a,b,A,T}$ and $t \in [0,1]$. Using the fact that, by definition, $\mathrm{vec}_T \circ c$ is a linear map and (4.74), we obtain

$$\begin{aligned}
&tH(z\|w) + (1-t)H(z'\|w) - H(tz + (1-t)z'\|w) \\
&\quad = t\tilde{D}(\mathrm{vec}_T \circ c(z)\|\mathrm{vec}_T \circ d(w)) + (1-t)\tilde{D}(\mathrm{vec}_T \circ c(z')\|\mathrm{vec}_T \circ d(w)) \\
&\qquad - \tilde{D}(\mathrm{vec}_T \circ c(tz + (1-t)z')\|\mathrm{vec}_T \circ d(w)) \\
&\quad = t\tilde{D}(\mathrm{vec}_T \circ c(z)\|\mathrm{vec}_T \circ d(w)) + (1-t)\tilde{D}(\mathrm{vec}_T \circ c(z')\|\mathrm{vec}_T \circ d(w)) \\
&\qquad - \tilde{D}(t\mathrm{vec}_T \circ c(z) + (1-t)\mathrm{vec}_T \circ c(z')\|\mathrm{vec}_T \circ d(w)) \\
&\quad \geq t(1-t)\frac{1}{2}\|\mathrm{vec}_T \circ c(z') - \mathrm{vec}_T \circ c(z)\|_2^2 \geq t(1-t)\frac{1}{2}\|z' - z\|_2^2.
\end{aligned}$$

Hence, by Theorem 2.1.9 in Nesterov (2004), for any $w \in \Lambda_{B,T}$, $H(\cdot\|w)$ is strongly convex with constant 1 on $\Lambda_{a,b,A,T}$. From Theorem 2.1.11 in the same reference, the latter is equivalent to the fact that, for any $w \in \Lambda_{B,T}$ and $z \in \Lambda_{a,b,A,T}$, $\partial_1\partial_1 H(z\|w) - I_{|A|\times|A|}$ is positive semi-definite, where $I_{|A|\times|A|}$ is the $|A| \times |A|$ identity matrix. This implies that, for any $w \in \Lambda_{B,T}$ and $z \in \Lambda_{a,b,A,T}$, $\partial_1\partial_1 H(z\|w)$ is positive definite. Using the definition of F, the latter is equivalent to the fact that, for any $w \in \Lambda_{B,T}$ and $z \in \Lambda_{a,b,A,T}$, $\partial_1 F(z\|w)$ is positive definite, which implies (4.73).

Fix $x \in \Gamma_T$. We can now apply the implicit function theorem (see, e.g., Fitzpatrick, 2009, Theorem 17.6, p 450) to conclude that there exists a

scalar $r>0$ and a differentiable function $G:\mathcal{B}_r(x)\to(0,1)^{|A|}$, where $\mathcal{B}_r(x)$ is an open ball of radius r centered at $\mathrm{vec}_B(x)$, such that, whenever $\|z-\mathrm{vec}_A(\mathcal{I}_{a,b}(x))\|_2<r$, $\|w-\mathrm{vec}_B(x)\|_2<r$ and $F(z,w)=0$, then $G(w)=z$. The latter and (4.72) imply that, for any $x'\in\Gamma_T$ such that $\|\mathrm{vec}_A(\mathcal{I}_{a,b}(x'))-\mathrm{vec}_A(\mathcal{I}_{a,b}(x))\|_2<r$ and $\|\mathrm{vec}_B(x')-\mathrm{vec}_B(x)\|_2<r$, we have $G(\mathrm{vec}_B(x'))=\mathrm{vec}_A(\mathcal{I}_{a,b}(x'))$. From the continuity of $\mathcal{I}_{a,b}$ stated in Lemma 4.A.2 below, we thus obtain that $G(\mathrm{vec}_B(x'))=\mathrm{vec}_A(\mathcal{I}_{a,b}(x'))$ for x' in a neighborhood of x, or, equivalently, that $G(w)=\mathrm{vec}_A\circ\mathcal{I}_{a,b}\circ d(w)$ for w in a neighborhood of $\mathrm{vec}_B(x)$.

Another consequence of the implicit function theorem is that

$$\partial_1 F(G(w),w)J_G(w)+\partial_2 F(G(w),w))=0_{\mathbb{R}^{|A|\times|B|}}\text{ for all }w\in\mathcal{B}_r(x),$$

where $J_G(w)$ is the Jacobian matrix of G evaluated at w. From (4.73), the previous centered display implies that

$$J_G(w)=-[\partial_1 F(G(w),w)]^{-1}\partial_2 F(G(w),w))$$

for all w in a neighborhood of $\mathrm{vec}_B(x)$. □

Lemma 4.A.1. *For $k\in\mathbb{N}$, let $\Delta_k=\{u\in(0,1)^k:\sum_{i=1}^k u_i=1\}$. Furthermore, for any $u,v\in(0,1)^k$, let*

$$\bar{D}(v\|u)=\sum_{i=1}^k v_i\log\frac{v_i}{u_i}.$$

Then, for any $u\in\Delta_k$, the function $\bar{D}(\cdot\|u)$ is strongly convex with constant 1 on Δ_k in the sense of Nesterov (2004, Section 2.1.3).

Proof. Fix $u\in\Delta_k$. According to Nesterov (2004, Definition 2.1.2, p 63), we need to prove that, for any $v,w\in\Delta_k$,

$$\bar{D}(w\|u)\geq\bar{D}(v\|u)+\partial_1\bar{D}(v\|u)^T(w-v)+\frac{1}{2}\|w-v\|_2^2.$$

Note that, for any $v\in(0,1)^k$, $\partial_1\bar{D}(v\|u)=(\log(v_1/u_1)+1,\ldots,\log(v_k/u_k)+1)$ and thus that, for any $v,w\in\Delta_k$,

$$\begin{aligned}\partial_1\bar{D}(v\|u)^T(w-v)&=\sum_{i=1}^k(\log(v_i/u_i)+1)(w_i-v_i)=\sum_{i=1}^k w_i\log(v_i/u_i)-\bar{D}(v\|u)\\&=\sum_{i=1}^k w_i\log\left(\frac{w_iv_i}{u_iw_i}\right)-\bar{D}(v\|u)=\bar{D}(w\|u)-\bar{D}(w\|v)-\bar{D}(v\|u).\end{aligned}$$

Hence,

$$\begin{aligned}\bar{D}(w\|u) &= \bar{D}(v\|u) + \partial_1 \bar{D}(v\|u)^T (w-v) + \bar{D}(w\|v) \\ &\geq \bar{D}(v\|u) + \partial_1 \bar{D}(v\|u)^T (w-v) + \frac{1}{2}\|w-v\|_1^2,\end{aligned}$$

as a consequence of Pinsker's inequality (see, e.g., Tsybakov, 2008, Lemma 2.5, p 88) and the fact that the total variation dinstance coincides with the L_1 dinstance in the considered finite setting. The proof is complete since $\|\cdot\|_1 \geq \|\cdot\|_2$. □

Lemma 4.A.2. *For any $T \subset I_{r,s}$, $T \neq \emptyset$, such that there exists $y \in \Gamma_{a,b}$ with $\mathrm{supp}(y) \subset T$, the function $\mathcal{I}_{a,b} : \Gamma_T \to \Gamma_{a,b}$ is continuous.*

Proof. The result is a direct consequence of Theorem 3.3 (iii) in Gietl and Reffel (2017). □

4.B Proof of Theorem 4.3.2

Proof of Theorem 4.3.2. By the definition of u in (4.30), it is necessary to check the conditions for existence and uniqueness of a I-projection on a Frechèt class of pmfs with fixed margins, given generally in Proposition 4.2.1, along with the particular case depicted in Proposition 4.2.2. Then, recall that conditions equivalent to the latter ones are given in Proposition 4.2.3. The result follows by carefully rewriting the conditions of Proposition 4.2.3 after noting that, with the notation of the same proposition, for any $A \subset [r]$ (resp. $B \subset [s]$) we have $P_{u^{[1]}}(A) = \frac{|A|}{r}$ (resp. $P_{u^{[2]}}(B) = \frac{|B|}{s}$). □

4.C Proofs of Propositions 4.4.1 and 4.4.3

Proof of Proposition 4.4.1. We first prove that Assertion 1 implies Assertion 2. Since there exists $v \in \Gamma_{\text{unif}}$ such that $\mathrm{supp}(v) = \mathrm{supp}(p)$, from Proposition 4.2.1, the I-projection of p on Γ_{unif} exists and is unique, that is, $u = \mathcal{U}(p)$ exists and is unique. Furthermore, from Proposition 4.2.2, u is diagonally equivalent to p. Since there exists $q \in \Gamma_{p^{[1]},p^{[2]}}$ such that $\mathrm{supp}(q) = \mathrm{supp}(u)$ (take $q = p$), $p' = \mathcal{I}_{p^{[1]},p^{[2]}}(u)$ exists and is unique, and p' is diagonally equivalent to u. Since, by transitivity via u, p' is diagonally equivalent to p, and since p' and p have the same margins, we can conclude from Property 1 of Pretzel (1980) (see also Brossard and Leuridan, 2018, Lemma 27) that $p' = p$.

We shall now prove that Assertion 2 implies Assertion 1. Since $u = \mathcal{U}(p) = \arg\inf_{y \in \Gamma_{\text{unif}}} D(y\|p)$, we have that $\text{supp}(u) \subset \text{supp}(p)$ by the definition of D in (4.13). Similarly, $p = \mathcal{I}_{p^{[1]},p^{[2]}}(u) = \arg\inf_{y \in \Gamma_{p^{[1]},p^{[2]}}} D(y\|u)$ which implies that $\text{supp}(p) \subset \text{supp}(u)$. Assertion 1 thus holds since $u \in \Gamma_{\text{unif}}$ and $\text{supp}(u) = \text{supp}(p)$. □

Proof of Proposition 4.4.3. Let (U,V) have pmf $v \in \Gamma_{\text{unif}}$. Then, Goodman's and Kruskal's gamma (Goodman and Kruskal, 1954) and Kendall's tau b (see, e.g., Kendall and Gibbons, 1990) of (U,V) are respectively defined by

$$G(v) = \frac{\kappa(v) - \delta(v)}{\kappa(v) + \delta(v)} \qquad \text{and} \qquad T(v) = \frac{\kappa(v) - \delta(v)}{\sqrt{\mathrm{P}(U \neq U')\mathrm{P}(V \neq V')}},$$

where (U',V') is an independent copy of (U,V) and

$$\kappa(v) = \mathrm{P}\{(U-U')(V-V') > 0\} \qquad \text{and} \qquad \delta(v) = \mathrm{P}\{(U-U')(V-V') < 0\}.$$

For $\kappa(v)$, we have that

$$\begin{aligned}\kappa(v) &= \sum_{(i,j) \in I_{r,s}} \sum_{(i',j') \in I_{r,s}} \mathbf{1}\{(i-i')(j-j') > 0\} \mathrm{P}(U = i, U' = i', V = j, V' = j') \\ &= 2 \sum_{\substack{i \in \{1,\ldots,r-1\} \\ j \in \{1,\ldots,s-1\}}} \sum_{\substack{i' \in \{i+1,\ldots,r\} \\ j' \in \{j+1,\ldots,s\}}} v_{ij} v_{i'j'}.\end{aligned}$$

For $\delta(v)$, we can use the fact that $\delta(v) = \mathrm{P}\{(U-U')(V-V') \neq 0\} - \kappa(v)$ and the fact that $\mathrm{P}\{(U-U')(V-V') \neq 0\}$ can be expressed as

$$\begin{aligned}&\sum_{(i,j) \in I_{r,s}} \sum_{(i',j') \in I_{r,s}} \mathbf{1}\{(i-i')(j-j') \neq 0\} \mathrm{P}(U = i, U' = i', V = j, V' = j') \\ &\quad = \sum_{(i,j) \in I_{r,s}} \sum_{(i',j') \in I_{r,s}} \mathbf{1}(i \neq i') \mathbf{1}(j \neq j') v_{ij} v_{i'j'} = \sum_{(i,j) \in I_{r,s}} v_{ij} \sum_{\substack{i'=1 \\ i' \neq i}}^{r} \sum_{\substack{j'=1 \\ j' \neq j}}^{s} v_{i'j'} \\ &\quad = \sum_{(i,j) \in I_{r,s}} v_{ij} \sum_{\substack{i'=1 \\ i' \neq i}}^{r} \left(\frac{1}{r} - v_{i'j} \right) = \sum_{(i,j) \in I_{r,s}} v_{ij} \left(\frac{r-1}{r} - \sum_{\substack{i'=1 \\ i' \neq i}}^{r} v_{i'j} \right) \\ &\quad = \sum_{(i,j) \in I_{r,s}} v_{ij} \left(\frac{r-1}{r} - \frac{1}{s} + v_{ij} \right) = 1 - \frac{1}{r} - \frac{1}{s} + \sum_{(i,j) \in I_{r,s}} v_{ij}^2.\end{aligned}$$

To obtain the expression $T(v)$, it remains to obtain the expressions of $\mathrm{P}(U \neq U')$ and $\mathrm{P}(V \neq V')$. We have

$$\mathrm{P}(U \neq U') = \sum_{i,i'=1}^{r} \mathbf{1}(i \neq i')\mathrm{P}(U = i, U' = i') = \frac{1}{r^2} \sum_{i=1}^{r} \sum_{\substack{i'=1 \\ i' \neq i}}^{r} 1 = \frac{r-1}{r}$$

and, similarly, $\mathrm{P}(V \neq V') = (s-1)/s$. □

4.D Proof of Proposition 4.5.3

Proof of Proposition 4.5.3. We only prove the first claim as the other claims are immediate consequences of well-known results. We shall first apply Proposition 4.2.6 with $a = u^{[1]}$, $b = u^{[2]}$, $T = I_{r,s}$, $A = I_{r-1} \times I_{s-1}$ and $B = I_{r,s} \setminus \{(r,s)\}$. Note that, with some abuse of notation, the map d from $\Lambda_{B,T}$ in (4.25) to Γ_T in (4.24) mentioned in its statement can be defined, for any $w \in \Lambda_{B,T}$, by

$$d(w_{11}, w_{21}, \ldots, w_{r-2,s}, w_{r-1,s}) = \begin{bmatrix} w_{11} & \ldots & w_{1s} \\ \vdots & & \vdots \\ w_{r1} & \ldots & 1 - \sum_{(i,j) \in B} w_{ij} \end{bmatrix} \tag{4.75}$$

and that, with some abuse of notation, the map c from $\Lambda_{a,b,A,T}$ in (4.27) to $\Gamma_{a,b,T}$ in (4.26) can be defined, for any $z \in \Lambda_{a,b,A,T}$, by

$$\begin{aligned} &c(z_{11}, z_{21}, \ldots, z_{r-1,1}, \ldots, z_{1,s-1}, z_{2,s-1}, \ldots, z_{r-1,s-1}) \\ &= \begin{bmatrix} z_{11} & \cdots & z_{1,s-1} & 1/r - \sum_{j=1}^{s-1} z_{1j} \\ \vdots & & \vdots & \vdots \\ z_{r-1,1} & \cdots & z_{r-1,s-1} & 1/r - \sum_{j=1}^{s-1} z_{r-1,j} \\ 1/s - \sum_{i=1}^{r-1} z_{i1} & \ldots & 1/s - \sum_{i=1}^{r-1} z_{i,s-1} & 1/r + 1/s - 1 + \sum_{(i,j) \in A} z_{ij} \end{bmatrix}. \end{aligned} \tag{4.76}$$

Recall the definition of $\mathcal{U}$ in (4.29) and let $u = \mathcal{U}(p)$. Proposition 4.2.6 then implies that the map $\mathrm{vec}_A \circ \mathcal{U} \circ d$ is differentiable at $\mathrm{vec}_B(p)$ with Jacobian matrix at $\mathrm{vec}_B(p)$ equal to $-L_u^{-1} M_p$, where

$$L_u = \partial_1 \partial_1 H(\mathrm{vec}_A(u) \| \mathrm{vec}_B(p)) \text{ and } M_p = \partial_2 \partial_1 H(\mathrm{vec}_A(u) \| \mathrm{vec}_B(p)),$$

the map H is defined in (4.28) and, as we shall see below, the first (resp. second) matrix only depends on u (resp. p). With some abuse of notation, let us denote the $(r-1)(s-1)$ components of $\partial_1 H(\mathrm{vec}_A(u) \| \mathrm{vec}_B(p))$ by

$$(H'_{11}, H'_{21}, \ldots, H'_{r-1,1}, \ldots, H'_{1,s-1}, H'_{2,s-1}, \ldots, H'_{r-1,s-1}).$$

Standard calculations give

$$H'_{kl} = \log\left(\frac{c_{kl}(u)}{d_{kl}(p)}\right) - \log\left(\frac{c_{rl}(u)}{d_{rl}(p)}\right) - \log\left(\frac{c_{ks}(u)}{d_{ks}(p)}\right) + \log\left(\frac{c_{rs}(u)}{d_{rs}(p)}\right), \qquad (k,l) \in A,$$

where d_{kl} and c_{kl} are the component maps of the maps d and c defined in (4.75) and (4.76), respectively. Additional differentiation then leads to the expressions of the elements of the matrices L_u and M_p given in (4.42) and (4.43), respectively.

Next, from the assumption that $\sqrt{n}(\hat{p}^{[n]} - p) \rightsquigarrow Z_p$ in $\mathbb{R}^{r\times s}$ and since

$$\sqrt{n}(p^{[n]} - \hat{p}^{[n]}) \xrightarrow{\mathrm{P}} 0 \text{ in } \mathbb{R}^{r\times s}, \qquad (4.77)$$

where $p^{[n]}$ is defined in (4.38), we immediately obtain that $\sqrt{n}(\mathrm{vec}_B(p^{[n]}) - \mathrm{vec}_B(p)) \rightsquigarrow \mathrm{vec}_B(Z_p)$ in $\mathbb{R}^{|B|}$, which, combined with the delta method (see, e.g., van der Vaart, 1998, Theorem 3.1) for the map $\mathrm{vec}_A \circ \mathcal{U} \circ d$ and, again, (4.77), implies that

$$\sqrt{n}(\mathrm{vec}_A \circ \mathcal{U} \circ d \circ \mathrm{vec}_B(p^{[n]}) - \mathrm{vec}_A \circ \mathcal{U} \circ d \circ \mathrm{vec}_B(p)) + L_u^{-1} M_p \sqrt{n}(\mathrm{vec}_B(\hat{p}^{[n]}) - \mathrm{vec}_B(p)) \xrightarrow{\mathrm{P}} 0 \text{ in } \mathbb{R}^{|A|}.$$

Since, for the considered choices of T and B, for any $y \in \Gamma_T$, $d \circ \mathrm{vec}_B(y) = y$, the latter is equivalent to

$$\sqrt{n}\,\mathrm{vec}_A(u^{[n]} - u) + L_u^{-1} M_p \sqrt{n}\,\mathrm{vec}_B(\hat{p}^{[n]} - p) \xrightarrow{\mathrm{P}} 0 \text{ in } \mathbb{R}^{|A|}.$$

Notice from (4.76) that, for any $z \in \Lambda_{a,b,A,T}$, $c(z) = \mathrm{vec}^{-1}(Kz) + C$, where C is a constant $r \times s$ matrix. Hence, for any $y, y' \in \Gamma_{a,b,T}$, $K\mathrm{vec}_A(y - y') = \mathrm{vec}(y - y')$ and, by the continuous mapping theorem, we obtain that

$$\sqrt{n}\,\mathrm{vec}(u^{[n]} - u) + K L_u^{-1} M_p \sqrt{n}\,\mathrm{vec}_B(\hat{p}^{[n]} - p) \xrightarrow{\mathrm{P}} 0 \text{ in } \mathbb{R}^{rs}.$$

The desired result finally follows from the fact that $N\mathrm{vec}(y) = \mathrm{vec}_B(y)$, $y \in \mathbb{R}^{r\times s}$. □

4.E Proofs of the results of Section 4.5.2

Proof of Proposition 4.5.9. To prove the consistency of $\theta^{[n]}$, we shall use Theorem 5.7 in van der Vaart (1998). For any $\theta \in \Theta$, let $\ell_{ij}^{[\theta]} = \log u_{ij}^{[\theta]}$, $(i,j) \in I_{r,s}$, and let

$$M^{[n]}(\theta) = \frac{\bar{L}^{[n]}(\theta)}{n} = \sum_{(i,j)\in I_{r,s}} u_{ij}^{[n]} \ell_{ij}^{[\theta]} \quad \text{and} \quad M(\theta) = \sum_{(i,j)\in I_{r,s}} u_{ij} \ell_{ij}^{[\theta]}. \qquad (4.78)$$

From the definition of $\theta^{[n]}$ in (4.53), we have that $M^{[n]}(\theta^{[n]}) = \sup_{\theta\in\Theta} M^{[n]}(\theta)$, which implies that $M^{[n]}(\theta^{[n]}) \geq M^{[n]}(\theta_0)$. Furthermore, from the triangle inequality,

$$\begin{aligned}\sup_{\theta\in\Theta}|M^{[n]}(\theta) - M(\theta)| &\leq \sup_{\theta\in\Theta}\sum_{(i,j)\in I_{r,s}}|\ell_{ij}^{[\theta]}(u_{ij}^{[n]} - u_{ij})| \\ &= \sup_{\theta\in\Theta}\sum_{(i,j)\in I_{r,s}}|\ell_{ij}^{[\theta]}||u_{ij}^{[n]} - u_{ij}| \\ &\leq |\log\lambda|\sum_{(i,j)\in I_{r,s}}|u_{ij}^{[n]} - u_{ij}| \xrightarrow{\mathrm{P}} 0 \text{ in } \mathbb{R}\end{aligned}$$

by Proposition 4.5.2 and the continuous mapping theorem. Moreover, from the identifiability of the family $\mathcal{J}$ and Lemma 5.35 of van der Vaart (1998), we have that θ_0 is the unique maximizer of M. The latter implies that, for every $\varepsilon > 0$, $\sup_{\theta\in\Theta:\|\theta-\theta_0\|_2>\varepsilon} M(\theta) < M(\theta_0)$. The consistency of $\theta^{[n]}$ finally follows from Theorem 5.7 in van der Vaart (1998). □

Proof of Proposition 4.5.11. We proceed along the lines of the proof of Theorem 5.21 in van der Vaart (1998). First, for any $\theta\in\Theta$, let

$$\Psi_k^{[n]}(\theta) = \frac{\partial M^{[n]}}{\partial\theta_k}(\theta), k\in I_m,$$

where $M^{[n]}$ is defined in (4.78), and

$$\Psi^{[n]}(\theta) = \left(\Psi_1^{[n]}(\theta),\ldots,\Psi_m^{[n]}(\theta)\right) = \sum_{(i,j)\in I_{r,s}} u_{ij}^{[n]}\dot{\ell}_{ij}^{[\theta]},$$

where $\dot{\ell}_{ij}^{[\theta]}$ is defined in (4.57). Since Θ is assumed to be open, the estimator $\theta^{[n]}$ in (4.53) is a zero of $\Psi^{[n]}$. Also, let

$$\Psi(\theta) = \mathbb{E}_{\theta_0}(\dot{\ell}_{(U,V)}^{[\theta]}) = \sum_{(i,j)\in I_{r,s}} u_{ij}\dot{\ell}_{ij}^{[\theta]}, \tag{4.79}$$

where (U,V) has pmf u and $\dot{\ell}_{ij}^{[\theta]}$ is defined in (4.57). Then, from the discussion preceding the statement of Proposition 4.5.11, we have that θ_0 is a zero of Ψ.

Since, for any $(i,j)\in I_{r,s}$, $\theta\mapsto\ell_{ij}^{[\theta]}$ is twice differentiable at any $\theta\in\Theta$, by the mean value theorem (see, e.g., Fitzpatrick, 2009, Theorem 15.29, p 408), there exists $\eta > 0$ and a positive matrix $q\in\mathbb{R}^{r\times s}$ such that, whenever $\|\theta-\theta_0\|_2 < \eta$, for any $(i,j)\in I_{r,s}$,

$$\|\dot{\ell}_{ij}^{[\theta]} - \dot{\ell}_{ij}^{[\theta_0]}\|_2 \leq q_{ij}\|\theta-\theta_0\|_2. \tag{4.80}$$

Furthermore, from the triangle inequality and the inequality of Cauchy-Schwarz, we have

$$\begin{aligned}\Delta^{[n]} &= \left\| \sum_{(i,j)\in I_{r,s}} \dot{\ell}_{ij}^{[\theta^{[n]}]} \sqrt{n}(u_{ij}^{[n]} - u_{ij}) - \sum_{(i,j)\in I_{r,s}} \dot{\ell}_{ij}^{[\theta_0]} \sqrt{n}(u_{ij}^{[n]} - u_{ij}) \right\|_2^2 \\ &\leq \left\{ \sum_{(i,j)\in I_{r,s}} \|\dot{\ell}_{ij}^{[\theta^{[n]}]} - \dot{\ell}_{ij}^{[\theta_0]}\|_2 \times |\sqrt{n}(u_{ij}^{[n]} - u_{ij})| \right\}^2 \leq A^{[n]} \times B^{[n]},\end{aligned} \tag{4.81}$$

where

$$A^{[n]} = \sum_{(i,j)\in I_{r,s}} \|\dot{\ell}_{ij}^{[\theta^{[n]}]} - \dot{\ell}_{ij}^{[\theta_0]}\|_2^2 \quad \text{and} \quad B^{[n]} = \sum_{(i,j)\in I_{r,s}} |\sqrt{n}(u_{ij}^{[n]} - u_{ij})|^2. \tag{4.82}$$

To show that $\Delta^{[n]} \xrightarrow{\mathrm{P}} 0$, let us first prove that $A^{[n]} \xrightarrow{\mathrm{P}} 0$. Fix $\varepsilon > 0$. Then, $\mathrm{P}(A^{[n]} > \varepsilon) \leq I^{[n]} + J^{[n]}$, where

$$\begin{aligned} I^{[n]} &= \mathrm{P}\left(\sum_{(i,j)\in I_{r,s}} \|\dot{\ell}_{ij}^{[\theta^{[n]}]} - \dot{\ell}_{ij}^{[\theta_0]}\|_2^2 > \varepsilon, \|\theta^{[n]} - \theta_0\|_2 < \eta \right), \\ J^{[n]} &= \mathrm{P}\left(\|\theta^{[n]} - \theta_0\|_2 \geq \eta \right). \end{aligned}$$

The fact that that $J^{[n]}$ converges to zero is a consequence of the consistency of $\theta^{[n]}$ and the continuous mapping theorem. For $I^{[n]}$, using (4.80), we obtain

$$\begin{aligned} I^{[n]} &\leq \mathrm{P}\left(\|\theta^{[n]} - \theta_0\|_2^2 \sum_{(i,j)\in I_{r,s}} q_{ij}^2 > \varepsilon, \|\theta^{[n]} - \theta_0\|_2 < \eta \right) \\ &\leq \mathrm{P}\left(\|\theta^{[n]} - \theta_0\|_2^2 \sum_{(i,j)\in I_{r,s}} q_{ij}^2 > \varepsilon \right) \to 0, \end{aligned}$$

again as a consequence of the consistency of $\theta^{[n]}$. Hence, $A^{[n]} \xrightarrow{\mathrm{P}} 0$.

Let us next verify that $\Delta^{[n]}$ in (4.81) converges in probability to zero. From Proposition 4.5.3, we know that the weak convergence of $\sqrt{n}(\hat{p}^{[n]} - p)$ in $\mathbb{R}^{r\times s}$ implies the weak convergence of $\sqrt{n}(u^{[n]} - u)$ in $\mathbb{R}^{r\times s}$. Hence, from the continuous mapping theorem, $B^{[n]}$ in (4.82) is bounded in probability, which implies that $\Delta^{[n]}$ in (4.81) converges to zero in probability.

Moreover, since $\Psi^{[n]}(\theta^{[n]}) = 0$ and $\Psi(\theta_0) = 0$, we have that

$$\begin{aligned} \sum_{(i,j)\in I_{r,s}} \dot{\ell}_{ij}^{[\theta^{[n]}]} \sqrt{n}(u_{ij}^{[n]} - u_{ij}) &= -\sum_{(i,j)\in I_{r,s}} \dot{\ell}_{ij}^{[\theta^{[n]}]} \sqrt{n} u_{ij} \\ &= \sum_{(i,j)\in I_{r,s}} \sqrt{n}(\dot{\ell}_{ij}^{[\theta_0]} - \dot{\ell}_{ij}^{[\theta^{[n]}]}) u_{ij}. \end{aligned} \tag{4.83}$$

Since, for any $(i,j) \in I_{r,s}$, $\theta \mapsto \ell_{ij}^{[\theta]}$ is twice differentiable at any $\theta \in \Theta$, the map Ψ in (4.79) is differentiable at its zero θ_0 with Jacobian matrix at θ_0 given by

$$\sum_{(i,j)\in I_{r,s}} u_{ij} \ddot{\ell}_{ij}^{[\theta_0]} = \mathbb{E}_{\theta_0}(\ddot{\ell}_{(U,V)}^{[\theta_0]}),$$

where $\ddot{\ell}_{ij}^{[\theta]}$ is defined in (4.59). We then obtain from the delta method (see, e.g., van der Vaart, 1998, Theorem 3.1) that

$$\sqrt{n}\left(\sum_{(i,j)\in I_{r,s}} \dot{\ell}_{ij}^{[\theta^{[n]}]} u_{ij} - \sum_{(i,j)\in I_{r,s}} \dot{\ell}_{ij}^{[\theta_0]} u_{ij}\right) = \mathbb{E}_{\theta_0}(\ddot{\ell}_{(U,V)}^{[\theta_0]})\sqrt{n}(\theta^{[n]} - \theta_0) + o_P(1).$$

Combining the latter display with (4.83) and the fact that $\Delta^{[n]}$ in (4.81) converges to zero in probability, we obtain that

$$\mathbb{E}_{\theta_0}(\ddot{\ell}_{(U,V)}^{[\theta_0]})\sqrt{n}(\theta^{[n]} - \theta_0) = -\sum_{(i,j)\in I_{r,s}} \dot{\ell}_{ij}^{[\theta_0]} \sqrt{n}(u_{ij}^{[n]} - u_{ij}) + o_P(1).$$

The continuous mapping theorem, (4.60) and Proposition 4.5.3 finally imply that

$$\begin{aligned}
\sqrt{n}(\theta^{[n]} - \theta_0) &= -\{\mathbb{E}_{\theta_0}(\ddot{\ell}_{(U,V)}^{[\theta_0]})\}^{-1} \sum_{(i,j)\in I_{r,s}} \dot{\ell}_{ij}^{[\theta_0]} \sqrt{n}(u_{ij}^{[n]} - u_{ij}) + o_P(1), \\
&= \{\mathbb{E}_{\theta_0}(\dot{\ell}_{(U,V)}^{[\theta_0]} \dot{\ell}_{(U,V)}^{[\theta_0],\top})\}^{-1} \dot{\ell}^{[\theta_0]} \sqrt{n}\,\mathrm{vec}(u^{[n]} - u) + o_P(1), \\
&= \{\mathbb{E}_{\theta_0}(\dot{\ell}_{(U,V)}^{[\theta_0]} \dot{\ell}_{(U,V)}^{[\theta_0],\top})\}^{-1} \dot{\ell}^{[\theta_0]} J_{u,p} \sqrt{n}\,\mathrm{vec}(\hat{p}^{[n]} - p) + o_P(1).
\end{aligned}$$

□

4.F Proofs of the results of Section 4.6

Proof of Proposition 4.6.1. Under H_0 in (4.61), we can decompose the goodness-of-fit process as

$$\sqrt{n}(u^{[n]} - u^{[\hat{\theta}^{[n]}]}) = \sqrt{n}(u^{[n]} - u) - \sqrt{n}(u^{[\hat{\theta}^{[n]}]} - u).$$

From the delta method (see, e.g., van der Vaart, 1998, Theorem 3.1), we then obtain that, under H_0,

$$\sqrt{n}\,\mathrm{vec}(u^{[\hat{\theta}^{[n]}]} - u) = \sqrt{n}\,\mathrm{vec}(u^{[\hat{\theta}^{[n]}]} - u^{[\theta_0]}) = \dot{u}^{[\theta_0]} \sqrt{n}(\hat{\theta}^{[n]} - \theta_0) + o_P(1).$$

From the assumptions, it follows that, under H_0, $\sqrt{n}\,\mathrm{vec}(u^{[n]} - u^{[\hat{\theta}^{[n]}]})$ has the same limiting distribution as

$$\sqrt{n}\,\mathrm{vec}(u^{[n]} - u) - \dot{u}^{[\theta_0]} V_{u,p}^{[\theta_0]} \sqrt{n}\,\mathrm{vec}(\hat{p}^{[n]} - p).$$

Finally, from Proposition 4.5.3, we obtain that

$$\sqrt{n}\,\mathrm{vec}(u^{[n]} - u^{[\hat{\theta}^{[n]}]}) = J_{u,p}\sqrt{n}\,\mathrm{vec}(\hat{p}^{[n]} - p) - \dot{u}^{[\theta_0]} V_{u,p}^{[\theta_0]} \sqrt{n}\,\mathrm{vec}(\hat{p}^{[n]} - p) + o_P(1).$$

□

Proof of Proposition 4.6.2. From Proposition 4.6.1, the continuous mapping theorem and the fact that $u^{[\hat{\theta}^{[n]}]} \xrightarrow{\mathrm{P}} u^{[\theta_0]}$ in $\mathbb{R}^{r\times s}$, we obtain that

$$\begin{aligned}
&\mathrm{diag}(G\,\mathrm{vec}(u^{[\hat{\theta}^{[n]}]}))^{-1/2}\sqrt{n}\,G\,\mathrm{vec}(u^{[n]} - u^{[\hat{\theta}^{[n]}]}) \\
&\quad = \mathrm{diag}(G\,\mathrm{vec}(u^{[\theta_0]}))^{-1/2} G(J_{u,p} - \dot{u}^{[\theta_0]} V_{u,p}^{[\theta_0]})\sqrt{n}\,\mathrm{vec}(\hat{p}^{[n]} - p) + o_P(1).
\end{aligned}$$

Consequently, when $(X_1, Y_1), \ldots, (X_n, Y_n)$ are independent copies of (X, Y), the sequence

$$\mathrm{diag}(G\,\mathrm{vec}(u^{[\hat{\theta}^{[n]}]}))^{-1/2}\sqrt{n}\,G\,\mathrm{vec}(u^{[n]} - u^{[\hat{\theta}^{[n]}]})$$

is asymptotically centered normal with covariance matrix given by (4.65). The first claim is finally a consequence of Lemma 17.1 in van der Vaart (1998). The second claim immediately follows from the continuous mapping theorem. □

Chapter 5

On the Differentiability of ϕ-projections in the Discrete Finite Case

This chapter is based on the pre-print Geenens et al. (2024)

Geenens, G., I. Kojadinovic, and T. Martini (2024). On the differentiability of ϕ-projections in the discrete finite case, https://arxiv.org/abs/2407.05997,

submitted to the *Annals of the Institute of Statistical Mathematics*.

5.1 Introduction

The concept of a I-projection due to Csiszár (see, e.g., Csiszár, 1975; Csiszár and Shields, 2004) played a crucial role in the developments of the previous chapter. Actually, as known, it has applications in many areas of probability and statistics. Informally, given a probability distribution q_0 of interest and a set $\mathcal{M}$ of probability distributions, it consists of finding an element of $\mathcal{M}$, if it exists, that is the "closest" to q_0 in the sense of the Kullback–Leibler (or information) divergence. A well-known class of alternatives to the Kullback–Leibler divergence (containing the latter) are the so-called ϕ-divergences (see, e.g., Ali and Silvey, 1966; Csiszár, 1967; Liese and Vajda, 1987; Csiszár and Shields, 2004, and the references therein) and ϕ-projections are merely the analogs of I-projections based on ϕ-divergences.

Recall that, in the considered discrete finite setting, in the previous chapter

we obtained a first differentiability result for I-projections on Fréchet classes of bivariate probability arrays by exploiting, among other things, continuity results for ϕ-projections on convex sets obtained by Gietl and Reffel (2013, 2017).

Using the latter as a starting point, focusing on (probability) measures on finite spaces, the primary goal of this chapter is to establish conditions under which ϕ-projections are continuously differentiable with respect to the input distribution. From a statistical inference perspective, such findings enable the immediate determination of the consistency and the asymptotic distribution of a ϕ-projection estimator (also known as the minimum ϕ-divergence estimator in the context of parametric inference). Consider that q_n is a consistent estimator of a target probability vector q_0, which can be uniquely ϕ-projected onto a set $\mathcal{M}$ of probability vectors of interest. Then, under conditions ensuring the continuity of ϕ-projections, the consistency of the ϕ-projection of q_n onto $\mathcal{M}$ follows directly from the continuous mapping theorem. Additionally, under conditions guaranteeing the continuous differentiability of ϕ-projections, the limiting distribution of a properly scaled version of the ϕ-projection of q_n onto $\mathcal{M}$ can be derived using the delta method.

When the set $\mathcal{M}$ of probability vectors to which one aims to ϕ-project is not necessarily convex, related findings (though not differentiability results *per se*) were presented by Jiménez-Gamero et al. (2011, Section 2). When $\mathcal{M}$ is a set of probability vectors derived from a given parametric distribution and q_0 is assumed to belong to $\mathcal{M}$, the aforementioned results encompass the well-known asymptotic properties of minimum ϕ-divergence estimators (see, e.g., Read and Cressie, 1988; Morales et al., 1995; Basu et al., 2011, and references therein). As the results of Jiménez-Gamero et al. (2011) were actually obtained regardless of whether q_0 belongs to $\mathcal{M}$ (thus also considering potential model misspecification), they pave the way for developing various inference procedures related to goodness-of-fit testing and model selection, as discussed in Sections 3 and 4 of Jiménez-Gamero et al. (2011) (see also references therein).

The specific objective of this work is to attempt to unify and extend the previous results. Building on the approaches of Gietl and Reffel (2013, 2017), we consider finite measures on finite spaces, not just probability measures on finite spaces. With a slight strengthening of the main condition from Jiménez-Gamero et al. (2011) and a crucial additional assumption missing from that reference (see Remark 5.3.8 in Section 5.3.2), we first

present a general result on the differentiability of ϕ-projections. Under an additional assumption of convexity of the set $\mathcal{M}$, we demonstrate that the aforementioned conditions are implied by simpler ones that are particularly easy to verify. For instance, we find that for many common choices of ϕ-divergences, ϕ-projections are automatically continuously differentiable when $\mathcal{M}$ is a subset defined by linear equalities. Our proofs draw on results from Gietl and Reffel (2017) and Rüschendorf (1987), as well as on the fact that ϕ-divergences constructed from strongly convex functions ϕ (in a specific sense) are strongly convex in their first argument.

This chapter is structured as follows. In the first section, we define ϕ-divergences for finite measures on finite spaces, recall their main properties as stated in Gietl and Reffel (2017), provide conditions under which they are strongly convex in their first argument, and define ϕ-projections. Subsequently, we present conditions under which ϕ-projections are continuous and continuously differentiable, and demonstrate that these can be replaced by considerably simpler conditions when the set $\mathcal{M}$ onto which one wishes to ϕ-project is convex. Finally, we illustrate how the derived results can be utilised to determine the asymptotics of minimum ϕ-divergence estimators when projecting onto parametric sets of probability vectors, sets of probability vectors generated from distributions with specific moments fixed, and Fréchet classes of bivariate probability arrays. The latter example concerns essentially the same result that was given in Section 4.2.4 of the previous chapter. Nevertheless, as shall be explained, it sheds more light on the different technique employed therein and eventually opens up the path to an useful alternative point of view. For a smoother reading, all the proofs are found in a series of appendices.

5.2 Preliminaries on ϕ-divergences and ϕ-projections

5.2.1 Notation

Let $m \in \mathbb{N}^+$ be fixed and let $\mathcal{X} = \{x_1, \ldots, x_m\}$ be a set of interest. We will examine ϕ-divergences and ϕ-projections for finite measures on $(\mathcal{X}, 2^{\mathcal{X}})$. Obviously, any finite measure μ on $(\mathcal{X}, 2^{\mathcal{X}})$ is discrete and can be represented by a collection of m point masses since

$$\mu(A) = \sum_{x_i \in A} \mu(\{x_i\}), \qquad A \subset \{x_1, \ldots, x_m\}.$$

As will become evident in the following subsections, the forthcoming developments depend only on $m=|\mathcal{X}|$ and the values that finite measures on $\mathcal{X}$ take on subsets of $\mathcal{X}$, that is, on the singletons $\{x_1\}$, ..., $\{x_m\}$. Therefore, in the remainder of this chapter, a finite measure μ on $\mathcal{X}$ will be simply represented by a vector $s\in[0,\infty)^m$ defined by $s_i=\mu(\{x_i\})$, $i\in\{1,\ldots,m\}$. We will use the letters s,t for denoting such point mass vectors in the case of general finite measures, and reserve the letters p,q for point mass vectors of probability measures, that is, whose components sum up to one – these will be simply referred to as probability vectors in the rest of this chapter. We define the support of any vector $s\in[0,\infty)^m$ as $\mathrm{supp}(s)=\{i\in\{1,\ldots,m\}:s_i>0\}$. Finally, for any set $S\subset[0,\infty)^m$ and any $\mathcal{I}\subset\{1,\ldots,m\}$, $S_{\mathcal{I}}=\{s\in S:\mathrm{supp}(s)\subset\mathcal{I}\}$. For example, given $t\in[0,\infty)^m$, $[0,\infty)^m_{\mathrm{supp}(t)}$ denotes the set $\{s\in[0,\infty)^m:\mathrm{supp}(s)\subset\mathrm{supp}(t)\}$, which is the set of nonnegative vectors whose support is included in that of t.

5.2.2 ϕ-divergences, for Finite Measures on Finite Spaces

Let $\phi:[0,\infty)\to\mathbb{R}$ be a convex function that is continuous at 0. Then, let $f:[0,\infty)^2\to\mathbb{R}\cup\{\infty\}$ be the function defined by

$$f(v,w)=\begin{cases} w\phi\left(\dfrac{v}{w}\right) & \text{if } v\in[0,\infty),\, w\in(0,\infty),\\[2ex] v\lim\limits_{x\to\infty}\dfrac{\phi(x)}{x} & \text{if } v\in[0,\infty),\, w=0,\\[2ex] 0 & \text{if } v=0,\, w=0. \end{cases} \tag{5.1}$$

Note that, by Proposition A.24 of Liese and Vajda (1987), $\lim_{x\to\infty}\phi(x)/x$ exists in $\mathbb{R}\cup\{\infty\}$ and that, by Proposition A.35 in the same reference, the function f is lower semicontinuous on $[0,\infty)^2$. Following Gietl and Reffel (2017, Section 2), it can be confirmed that, for any fixed $v\geq 0$, $f(v,\cdot)$ is continuous.

We next state a trivial lemma which shows that, for any $M\in(0,\infty)$, the restriction of f to $[0,M]^2$ can be immediately recovered from its restriction to $[0,1]^2$.

Lemma 5.2.1. *Let $M\in(0,\infty)$. Then, for any $v,w\in[0,M]^2$, $f(v,w)=Mf(v/M,w/M)$.*

The following additional condition on ϕ is commonly found in the literature.

Condition 5.2.2 (Strict convexity of ϕ)**.** *The function ϕ is strictly convex.*

It is simple to show that, for any fixed $w>0$, $f(\cdot,w)$ is strictly convex if and only if Condition 5.2.2 holds.

Definition 5.2.3 (ϕ-divergence)**.** *For any $s,t\in[0,\infty)^m$, the ϕ-divergence of s relative to t is defined by*

$$D_\phi(s\,|\,t)=\sum_{i=1}^{m} f(s_i,t_i). \tag{5.2}$$

It can verified from (5.1) that, if $\operatorname{supp}(s)\subset\operatorname{supp}(t)$, $D_\phi(s\,|\,t)<\infty$. Note that since s and t are not necessarily probability vectors, $D_\phi(s\,|\,t)$ may be negative (take for instance $\phi(x)=x\log(x)$, $x\in(0,\infty)$, $\phi(0)=0$ and $s_i\leq t_i$, $i\in\{1,\ldots,m\}$, with at least one $s_i<t_i$). Moreover, the following properties (see Gietl and Reffel, 2017, Theorem 2.1) hold as a consequence of the properties of the function f in (5.1) discussed above.

Proposition 5.2.4. *The following assertions hold:*

(i) The function D_ϕ is lower semicontinuous on $[0,\infty)^m\times[0,\infty)^m$.

(ii) For fixed $s\in[0,\infty)^m$, $D_\phi(s\,|\,\cdot)$ is continuous on $[0,\infty)^m$.

(iii) For fixed $t\in[0,\infty)^m$, $t\neq 0_{\mathbb{R}^m}$, $D_\phi(\cdot\,|\,t)$ is strictly convex on $[0,\infty)^m_{\operatorname{supp}(t)}$ if and only if Condition 5.2.2 holds.

We additionally state an immediate corollary of Lemma 5.2.1 which shows that, for any $M\in(0,\infty)$, the restriction of D_ϕ to $[0,M]^m\times[0,M]^m$ can be immediately recovered from its restriction to $[0,1]^m\times[0,1]^m$.

Lemma 5.2.5. *Let $M\in(0,\infty)$. Then, for any $s,t\in[0,M]^m$, $D_\phi(s\,|\,t)=MD_\phi(s/M\,|\,t/M)$.*

5.2.3 Existence of ϕ-projections on Compact Subsets

Definition 5.2.6 (ϕ-projection)**.** *Let $t\in[0,\infty)^m$ and let $\mathcal{M}\subset[0,\infty)^m$, $\mathcal{M}\neq\emptyset$. A ϕ-projection of t on $\mathcal{M}$, if it exists, is an element $s^*\in\mathcal{M}$ satisfying*

$$D_\phi(s^*\,|\,t)=\inf_{s\in\mathcal{M}} D_\phi(s\,|\,t).$$

The next result (see Gietl and Reffel, 2017, Lemma 3.1 (i)) shows that ϕ-projections exist as soon as the set $\mathcal{M}$ on which one wishes to ϕ-project is compact.

Proposition 5.2.7 (Existence of ϕ-projection)**.** *Let $t \in [0,\infty)^m$ and let $\mathcal{M}$ be a nonempty compact subset of $[0,\infty)^m$. Then, there exists a ϕ-projection of t on $\mathcal{M}$.*

We end this subsection by noting that, if one wishes to ϕ-project only on compact sets of nonnegative vectors, without loss of generality, it suffices to study ϕ-projections of vectors of $[0,1]^m$ on compact subsets of $[0,1]^m$. We briefly explain why. Let $t \in [0,\infty)^m$ and let $\mathcal{M}$ be a nonempty compact subset of $[0,\infty)^m$. Hence, there exists $M \in (0,\infty)$ such $t \in [0,M]^m$ and $\mathcal{M} \subset [0,M]^m$. Let $t' = t/M \in [0,1]^m$ and let $\mathcal{M}' = \{s/M : s \in \mathcal{M}\} \subset [0,1]^m$. Clearly, $\mathcal{M}'$ is also compact. Moreover,

$$\operatorname*{arg\,inf}_{s' \in \mathcal{M}'} D_\phi(s' \,|\, t') = \frac{1}{M} \operatorname*{arg\,inf}_{s \in \mathcal{M}} D_\phi(s/M \,|\, t') = \frac{1}{M} \operatorname*{arg\,inf}_{s \in \mathcal{M}} D_\phi(s \,|\, t),$$

where the second equality holds from Lemma 5.2.5. Hence, as we shall only be interested in ϕ-projecting on compact sets of nonnegative vectors, without loss of generality, we shall only consider vectors in $[0,1]^m$ in the remainder of this work.

5.2.4 Unicity of ϕ-projections on Compact and Convex Sets

Another condition on ϕ commonly appearing in the literature is the following.

Condition 5.2.8. *The function ϕ satisfies $\lim_{x\to\infty} \frac{\phi(x)}{x} = \infty$.*

As a direct corollary of Lemma 3.1 of Gietl and Reffel (2017) (with a slight modification of statement (ii) coming from its proof), we have:

Proposition 5.2.9 (Unicity of ϕ-projection)**.** *Let $t \in [0,1]^m$ and $\mathcal{M}$ be a nonempty closed and convex subset of $[0,1]^m$. Then, under Condition 5.2.2:*

(i) If $\mathcal{M} \subset [0,1]^m_{\mathrm{supp}(t)}$, the ϕ-projection of t on $\mathcal{M}$ is unique.

(ii) If $\mathcal{M} \cap [0,1]^m_{\mathrm{supp}(t)} \neq \emptyset$ and Condition 5.2.8 holds, the ϕ-projection of t on $\mathcal{M}$ is unique and coincides with the ϕ-projection of t on $\mathcal{M} \cap [0,1]^m_{\mathrm{supp}(t)}$.

In the previous proposition, Condition 5.2.2 is crucial, as the next example shows (see Example 3.2 of Gietl and Reffel, 2017).

Example 5.2.10. *If ϕ is convex but not strictly convex, unicity of ϕ-projections on compact and convex sets may not be guaranteed. Indeed,*

consider ϕ *defined by* $\phi(x)=|x-1|$ *for* x *in* $[0,\infty)$. *The corresponding* ϕ*-divergence is commonly known as the total variation, since for any* s,t *in* $[0,1]^m$ *we have*

$$D_\phi(s\,|\,t)=\sum_{i=1}^{m}|s_i-t_i|.$$

Let $t=(1,1)\in[0,1]^2$ *and let the convex and compact set* $\mathcal{M}$ *be defined by*

$$\mathcal{M}=\{s\in[0,1]^2 : s_1+s_2=1\}.$$

For all $s\in\mathcal{M}$, *we have*

$$D_\phi(s\,|\,t)=|s_1-t_1|+|s_2-t_2|=t_1-s_1+t_2-s_2=1.$$

Hence all elements of $\mathcal{M}$ *are* ϕ*-projections of* t *on* $\mathcal{M}$.

5.2.5 ϕ-divergences Strongly Convex in Their First Argument

We conclude this section by presenting conditions under which ϕ-divergences are strongly convex in their first argument. Let $\mathcal{Y}$ be a convex subset of $\mathbb{R}^d$ ($d\in\mathbb{N}^+$) and recall that a function H from $\mathcal{Y}$ to $\mathbb{R}$ is strongly convex if and only if there exists a constant $\kappa_H\in(0,\infty)$ such that, for any $x,y\in\mathcal{Y}$ and $\alpha\in[0,1]$,

$$H(\alpha x+(1-\alpha)y)\leq\alpha H(x)+(1-\alpha)H(y)-\frac{\kappa_H}{2}\alpha(1-\alpha)\|x-y\|_2^2.$$

When H is twice continuously differentiable on $\mathring{\mathcal{Y}}$, the interior of $\mathcal{Y}$, it is known (see, e.g., Nesterov, 2004, Theorem 2.1.11) that H is strongly convex on $\mathcal{Y}$ if and only if, for any $x\in\mathring{\mathcal{Y}}$, $\mathcal{H}(x)-\kappa_H I_d$ is positive semi-definite, where $\mathcal{H}(x)$ is the Hessian matrix of H at x and I_d is the $d\times d$ identity matrix. We will utilise this equivalent characterisation in Section 5.3.3.

The following condition on ϕ will be essential in the subsequent developments.

Condition 5.2.11 (Strong convexity of ϕ)**.** *For any* $w\in(0,\infty)$, *the restriction of* ϕ *to* $[0,1/w]$ *is strongly convex.*

Note that many functions ϕ involved in the definition of classical ϕ-divergences satisfy this condition. The equivalent characterization of strong convexity mentioned above allows for the straightforward verification of Condition 5.2.11 for those functions ϕ satisfying:

Table 5.2.1: The function ϕ of several classical ϕ-divergences satisfying Conditions 5.2.11 and 5.2.12 and the associated strong convexity constant.

Divergence	$\phi(x)$	$\kappa_\phi(w)$
Kullback–Leibler	$x\log x$	w
Pearson's χ^2	$(x-1)^2$	2
Squared Hellinger	$2(1-\sqrt{x})$	$\frac{w^{\frac{3}{2}}}{2}$
Reverse relative entropy	$-\log x$	w^2
Vincze–Le Cam	$\frac{(x-1)^2}{x+1}$	$8\left(\frac{1}{w}+1\right)^{-3}$
Jensen–Shannon	$(x+1)\log\frac{2}{x+1}+x\log x$	$\frac{w^2}{w+1}$
Neyman's χ^2	$\frac{1}{x}-1$	$2w^3$
α-divergence	$\frac{4\left(1-x^{\frac{1+\alpha}{2}}\right)}{1-\alpha^2}, \alpha<3, \alpha\neq\pm1$	$w^{\frac{3-\alpha}{2}}$

Condition 5.2.12 (Differentiability of ϕ)**.** *The function ϕ is twice continuously differentiable on $(0,\infty)$.*

Indeed, Condition 5.2.11 is then equivalent to $\phi''(x)\geq\kappa_\phi(w)$ for all $x\in(0,1/w)$, where $\kappa_\phi(w)$ is the strong convexity constant of the restriction of ϕ to $[0,1/w]$. Note that Melbourne (2020) considered ϕ-divergences constructed from such functions ϕ, but did not investigate their strong convexity in their first argument. Table 5.2.1 shows the function ϕ and the associated strong convexity constant for several classical ϕ-divergences satisfying Condition 5.2.12.

The following result is proven in Section 5.A.

Proposition 5.2.13. *Under Condition 5.2.11, the function $D_\phi(\cdot\,|\,t)$ is strongly convex on $[0,1]^m_{\mathrm{supp}(t)}$ for all $t\in[0,1]^m$, $t\neq 0_{\mathbb{R}^m}$.*

We end this subsection with the statement of a short lemma, proven in Section 5.A, that we shall use in Section 5.3.3.

Lemma 5.2.14. *Condition 5.2.11 implies Condition 5.2.2.*

5.3 On the Continuous Differentiability of ϕ-projections

5.3.1 Framework

Let $t_0 \in [0,1]^m$ be the nonnegative vector to be ϕ-projected on a nonempty subset $\mathcal{M}$ of $[0,1]^m$. Since we are interested in stating differentiability results at t_0, it is natural impose that $t_0 \in (0,1)^m$. The goal of this section is to present conditions under which the function

$$\mathcal{S}^*(t) = \arg\inf_{s \in \mathcal{M}} D_\phi(s \,|\, t) \tag{5.3}$$

is well-defined over an open neighborhood $\mathcal{N}(t_0) \subset (0,1)^m$ of t_0 and continuously differentiable at t_0.

As we continue, we shall assume that the set $\mathcal{M}$ is constructed as follows.

Condition 5.3.1 (Construction of $\mathcal{M}$)**.** *There exists a bounded subset Θ of $\mathbb{R}^k$ for some strictly positive integer $k \leq m$ with $\mathring{\Theta} \neq \emptyset$ and a continuous injective function $\mathcal{S}$ from $\bar{\Theta}$ to $[0,1]^m$ that is twice continuously differentiable on $\mathring{\Theta}$ such that $\mathcal{M} = \mathcal{S}(\bar{\Theta})$ and $\mathcal{S}(\mathring{\Theta}) \subset (0,1)^m$.*

Remark 5.3.2. Assuming Θ is bounded in Condition 5.3.1 is not restrictive, as an unbounded parameter space can always be reparametrised using a suitable bijection into a bounded subset of $\mathbb{R}^k$. Additionally, we do not define Θ as a closed set as we find it more explicit to write $\bar{\Theta}$ when necessary. Assuming $\mathcal{S}$ injective is particularly natural in a statistical framework as it ensures the identifiability of the "parametric" model $\mathcal{S}(\bar{\Theta}) = \{\mathcal{S}(\theta) : \theta \in \bar{\Theta}\}$. Furthermore, note that Condition 5.3.1 implies that $\mathcal{M} = \mathcal{S}(\bar{\Theta})$ is compact (since $\bar{\Theta}$ is compact and $\mathcal{S}$ is continuous) and nonempty (because $\mathring{\Theta} \neq \emptyset$). Hence, by Proposition 5.2.7, for any $t \in [0,1]^m$, a ϕ-projection of t on $\mathcal{M}$ exists (but may not be unique). Finally, the assumption that $\mathcal{S}(\mathring{\Theta}) \subset (0,1)^m$ is necessary, on the one hand, so that, for any $t \in (0,1)^m$, the function $D_\phi(\mathcal{S}(\cdot) \,|\, t)$ is twice continuously differentiable on $\mathring{\Theta}$, and, on the other hand, to guarantee that $\mathcal{M} \cap (0,1)^m \neq \emptyset$ (which are both needed to obtain the desired differentiability results as we shall see in the next subsections). □

The previous setting is very general. To see this, consider for a moment that the vectors of interest in $[0,1]^m$ are probability vectors. Then, for some $k < m$, $\mathcal{M} = \mathcal{S}(\bar{\Theta})$ could represent a parametric family of probability vectors. For example, the first application in Section 5.4 considers the class of probability vectors obtained from the binomial distributions with parameters $m-1$ and $\theta \in \bar{\Theta} = [0,1]$. Upon reflection, it becomes clear that

most families of discrete distributions with finite support could indeed be considered. Note, however, that $\mathcal{M}$ is typically not convex in such parametric situations. As shall be illustrated in the second and third application in Section 5.4, Condition 5.3.1 can also be tailored to scenarios in which one wishes to ϕ-project on a set $\mathcal{M}$ of probability vectors defined by linear equalities, a common situation in applications implying the convexity of $\mathcal{M}$. In such convex cases, the key conditions necessary for the differentiability of ϕ-projections are substantially simpler and automatically verified for a rather large class of ϕ-divergences. This will be established in Section 5.3.3. Before that, Section 5.3.2 will first state general conditions under which $\mathcal{S}^*$ in (5.3) is well-defined and continuously differentiable at t_0. In what follows, Conditions 5.2.2, 5.2.12 and 5.3.1 will always be assumed to hold.

5.3.2 The General Case of $\mathcal{M}$ Compact

We will first examine the differentiability of the function $\mathcal{S}^*$ in (5.3) at $t_0 \in (0,1)^m$ without assuming that the set $\mathcal{M} = \mathcal{S}(\bar{\Theta})$ is convex. More in detail, we consider the following (nested) conditions.

Condition 5.3.3 (ϕ-projections in a neighborhood of t_0 I)**.** *There exists an open neighborhood $\mathcal{N}(t_0) \subset (0,1)^m$ of t_0 such that, for any $t \in \mathcal{N}(t_0)$, the ϕ-projection s^* of t on $\mathcal{M} = \mathcal{S}(\bar{\Theta})$ (exists and) is unique.*

Condition 5.3.4 (ϕ-projections in a neighborhood of t_0 II)**.** *Condition 5.3.3 holds and, for any $t \in \mathcal{N}(t_0)$, the unique ϕ-projection $s^* = \mathcal{S}^*(t)$ satisfies $s^* = \mathcal{S}(\theta^*)$ for some (unique) $\theta^* \in \mathring{\Theta}$.*

Clearly, Condition 5.3.4 is a strengthening of Condition 5.3.3, as it requires θ^* to belong to the interior of Θ. Also, when restricted to probability vectors, Condition 5.3.4 is a slight strengthening of Assumption 3 of Jiménez-Gamero et al. (2011) which is recovered by setting $\mathcal{N}(t_0) = \{t_0\}$. The cited reference notes that this type of condition is common in the literature (see, e.g., White, 1982; Lindsay, 1994; Broniatowski and Keziou, 2009).

We shall first determine additional conditions under which the function $\vartheta^* : \mathcal{N}(t_0) \to \bar{\Theta}$ defined by

$$\vartheta^*(t) = \operatorname*{arg\,inf}_{\theta \in \bar{\Theta}} D_\phi(\mathcal{S}(\theta) \,|\, t) \tag{5.4}$$

is continuously differentiable at t_0. Note that, under Condition 5.3.4, (5.4) can be equivalently expressed as

$$\vartheta^*(t) = \operatorname*{arg\,inf}_{\theta \in \mathring{\Theta}} D_\phi(\mathcal{S}(\theta) \,|\, t), \qquad t \in \mathcal{N}(t_0), \tag{5.5}$$

implying that it is then a function from $\mathcal{N}(t_0)$ to $\mathring{\Theta}$. As $\mathcal{S}^*$ in (5.3) can be expressed as

$$\mathcal{S}^*(t) = \mathcal{S}(\vartheta^*(t)), \quad t \in \mathcal{N}(t_0), \tag{5.6}$$

we see that conditions under which $\mathcal{S}^*$ is continuously differentiable at t_0 will follow from conditions under which ϑ^* in (5.5) is continuously differentiable at t_0.

The following continuity result is proven in Section 5.B.

Proposition 5.3.5 (Continuity of ϑ^* at t_0)**.** *Under Condition 5.3.3, the function ϑ^* in (5.4) is continuous at t_0.*

The next corollary is an immediate consequence of the previous proposition, the continuity of $\mathcal{S}$ on $\bar{\Theta}$ and (5.6).

Corollary 5.3.6 (Continuity of $\mathcal{S}^*$ at t_0)**.** *Under Condition 5.3.3, the function $\mathcal{S}^*$ in (5.3) is continuous at t_0.*

As shall be verified in the proof of Lemma 5.3.9 (given in Section 5.B), the differentiability assumptions made on ϕ and $\mathcal{S}$ from Section 5.3 onwards imply that, for any $t \in (0,1)^m$, the function $D_\phi(\mathcal{S}(\cdot)\,|\,t)$ is twice continuously differentiable on $\mathring{\Theta}$. The continuous differentiability of the function ϑ^* in (5.4) will be shown under Condition 5.3.4 and the following additional condition.

Condition 5.3.7 (Invertibility condition)**.** *Under Condition 5.3.4, the $k \times k$ matrix whose elements are*

$$\left.\frac{\partial^2 D_\phi(\mathcal{S}(\theta)\,|\,t_0)}{\partial\theta_i \partial\theta_j}\right|_{\theta=\vartheta^*(t_0)}, \qquad i,j \in \{1,\ldots,k\},$$

is positive definite.

Note that the matrix appearing in the previous condition is well-defined as, under Condition 5.3.4, (5.5) implies that $\vartheta^*(t_0) \in \mathring{\Theta}$.

Remark 5.3.8. The matrix defined in Condition 5.3.7 corresponds to the matrix given in Eq. (3) of Jiménez-Gamero et al. (2011). In that reference, it is claimed that, when restricted to probability vectors, Condition 5.3.4 with $\mathcal{N}(t_0) = \{t_0\}$ (which corresponds to their Assumption 3) implies Condition 5.3.7. However, Condition 5.3.4 with $\mathcal{N}(t_0) = \{t_0\}$ is equivalent to saying that the function $D_\phi(\mathcal{S}(\cdot)\,|\,t_0)$ has a unique minimum at $\theta_0 = \vartheta^*(t_0) \in \mathring{\Theta}$ which, from second order necessary optimality conditions, only implies the positive semi-definiteness of the matrix in Condition 5.3.7.

We provide here a simple counterexample, illustrating that Condition 5.3.7 cannot be dispensed with. Let $\phi(x) = (x-1)^4$, $x \in [0,\infty)$. The function ϕ is strictly convex on its domain and twice continuously differentiable on $(0,\infty)$. We restrict our attention to probability vectors. Let $q_0 = (1/3,1/3,1/3)$ be the element to be ϕ-projected on the subset of probability vectors $\mathcal{M} = \mathcal{S}(\bar{\Theta})$ where $\Theta = (0,1/2)$ and $\mathcal{S}(\theta) = (\theta,\theta,1-2\theta)$, $\theta \in \bar{\Theta}$. The function $\mathcal{S}$ is clearly twice continuously differentiable on $\mathring{\Theta} = \Theta$. Let h be the function from $\bar{\Theta}$ to $\mathbb{R}$ defined by

$$h(\theta) = D_\phi(\mathcal{S}(\theta) \,|\, q_0), \qquad \theta \in \bar{\Theta}.$$

It is then easy to verify that

$$h(\theta) = \frac{1}{3}(3\theta-1)^4 + \frac{1}{3}(3\theta-1)^4 + \frac{1}{3}(3(1-2\theta)-1)^4, \qquad \theta \in \bar{\Theta},$$

and that h attains its unique minimum at $\theta_0 = 1/3 \in \mathring{\Theta}$. Standard calculations show that

$$h''(\theta) = 36(3\theta-1)^2 + 36(3\theta-1)^2 + 144[3(1-2\theta)-1]^2, \qquad \theta \in \mathring{\Theta},$$

and that $h''(\theta_0) = h''(\vartheta^*(q_0)) = 0$. □

Before stating one of the main results of this work, let us show how to express the matrix defined in Condition 5.3.7 under the elegant form given in Eq. (3) of Jiménez-Gamero et al. (2011). To achieve that, further definitions are needed to be able to properly deal with the differentiability of matrix-valued functions. Let $a,b,c,d \in \mathbb{N}^+$. For any function $g : U \to \mathbb{R}^{b\times c}$, where U is an open subset of $\mathbb{R}^a$, we shall denote by $g_{i,j}$, $i \in \{1,\ldots,b\}$, $j \in \{1,\ldots,c\}$, its bc component functions. Furthermore, we shall say that g is r-times continuously differentiable on U, $r \geq 1$, if, for any $i \in \{1,\ldots,b\}$, $j \in \{1,\ldots,c\}$, the r-order partial derivatives of $g_{i,j}$ exist and are continuous on U. Next, consider the (column-major) vectorization operator which maps any function $g : \mathbb{R}^a \to \mathbb{R}^{b\times c}$ to its vectorized version $\overline{g} : \mathbb{R}^a \to \mathbb{R}^{bc}$ whose bc components functions $\overline{g}_1,\ldots,\overline{g}_{bc}$ are defined by $\overline{g}_{i+(j-1)b} = g_{i,j}$, $(i,j) \in \{1,\ldots,b\} \times \{1,\ldots,c\}$. For any continuously differentiable function $g : U \to \mathbb{R}^{b\times c}$, where U is an open subset of $\mathbb{R}^a$, we then define its Jacobian matrix at $y \in U \subset \mathbb{R}^a$ as the Jacobian matrix $J[\overline{g}](y)$ of the vectorized version $\overline{g}$ of g at y. In other words,

$$J[g](y) = J[\overline{g}](y) = \begin{bmatrix} \left.\frac{\partial \overline{g}_1(x)}{\partial x_1}\right|_{x=y} & \cdots & \left.\frac{\partial \overline{g}_1(x)}{\partial x_a}\right|_{x=y} \\ \vdots & & \vdots \\ \left.\frac{\partial \overline{g}_{bc}(x)}{\partial x_1}\right|_{x=y} & \cdots & \left.\frac{\partial \overline{g}_{bc}(x)}{\partial x_a}\right|_{x=y} \end{bmatrix}, \qquad y \in U \subset \mathbb{R}^a.$$

Note that $J[g] = J[\overline{g}]$ is a function from U to $\mathbb{R}^{bc \times a}$. If g is twice continuously differentiable on U, then, for any $y \in U$, we define $J_2[g](y)$ to be the Jacobian of the function $\overline{J[g]}$ at y, that is, $J_2[g](y) = J[\overline{J[g]}](y)$, $y \in U$. A little reflection reveals that, given the previous definitions, it makes no difference in terms of Jacobian whether a given function is regarded as taking its values in $\mathbb{R}^{b\times 1}$, $\mathbb{R}^{1\times b}$ or $\mathbb{R}^b$. Similarly, we adhere to the common convention that vectors, when appearing in matrix expressions, are treated as 1-column matrices.

The following lemma, proven in Section 5.B, can be used to obtain an explicit expression of the matrix appearing in Condition 5.3.7.

Lemma 5.3.9. *For any $t \in (0,1)^m$ and $\theta \in \mathring{\Theta}$, we have that*

$$J_2[D_\phi(\mathcal{S}(\cdot)\,|\,t)](\theta) = J[\mathcal{S}](\theta)^\top \operatorname{diag}\left(\frac{1}{t_1}\phi''\left(\frac{\mathcal{S}_1(\theta)}{t_1}\right),\ldots,\frac{1}{t_m}\phi''\left(\frac{\mathcal{S}_m(\theta)}{t_m}\right)\right) J[\mathcal{S}](\theta) + \left(I_k \otimes \left(\phi'\left(\frac{\mathcal{S}_1(\theta)}{t_1}\right),\ldots,\phi'\left(\frac{\mathcal{S}_m(\theta)}{t_m}\right)\right)^\top\right) J_2[\mathcal{S}](\theta), \quad (5.7)$$

where $\mathcal{S}_1,\ldots,\mathcal{S}_m$ are the m component functions of $\mathcal{S}$ and the symbol $\otimes$ denotes the Kronecker product.

We are now ready to state one of our main results. Its proof is given in Section 5.B.

Theorem 5.3.1 (Differentiability of ϑ^* at t_0)**.** *Under Conditions 5.3.4 and 5.3.7, ϑ^* in (5.5) is continuously differentiable at $t_0 = (t_{0,1},\ldots,t_{0,m}) \in (0,1)^m$ with Jacobian matrix at t_0 given by*

$$J[\vartheta^*](t_0) = J_2[D_\phi(\mathcal{S}(\cdot)\,|\,t_0)](\vartheta^*(t_0))^{-1} J[\mathcal{S}](\vartheta^*(t_0))^\top \Delta(t_0), \quad (5.8)$$

where the matrix-valued function $J_2[D_\phi(\mathcal{S}(\cdot)\,|\,t_0)]$ is as in (5.7) with $t = t_0$ and

$$\Delta(t_0) = \operatorname{diag}\left(\frac{\mathcal{S}_1(\vartheta^*(t_0))}{t_{0,1}^2}\phi''\left(\frac{\mathcal{S}_1(\vartheta^*(t_0))}{t_{0,1}}\right),\ldots \ldots,\frac{\mathcal{S}_m(\vartheta^*(t_0))}{t_{0,m}^2}\phi''\left(\frac{\mathcal{S}_m(\vartheta^*(t_0))}{t_{0,m}}\right)\right). \quad (5.9)$$

The following corollary is a direct consequence of (5.6), Theorem 5.3.1 and the chain rule.

Corollary 5.3.10 (Differentiability of $\mathcal{S}^*$ at t_0)**.** *Assume that Conditions 5.3.4 and 5.3.7 hold. Then, the function $\mathcal{S}^*$ in (5.3) is continuously differentiable at $t_0 \in (0,1)^m$ with Jacobian matrix at t_0 given by*

$$J[\mathcal{S}^*](t_0) = J[\mathcal{S}](\vartheta^*(t_0))J[\vartheta^*](t_0), \tag{5.10}$$

where $J[\vartheta^*](t_0)$ is given in (5.8).

Remark 5.3.11. When t_0 is known, one can attempt to verify the conditions of Theorem 5.3.1 or Corollary 5.3.10 analytically, or at least empirically. For Condition 5.3.4, one could choose a few vectors t in a "neighborhood" of t_0 and attempt to verify (at least numerically) that the function $D_\phi(\mathcal{S}(\cdot)\,|\,t)$ has a unique minimum on $\mathring{\Theta} \subset \mathbb{R}^k$. When $k \in \{1,2\}$ in particular, the preceding verification could be graphical (see Section 5.4.1). As to Condition 5.3.7, one could simply compute the matrix $J_2[D_\phi(\mathcal{S}(\cdot)\,|\,t_0)](\vartheta^*(t_0))$ and attempt to invert it. In a statistical context, t_0 is typically unknown, though. Naturally, one can simply replace t_0 by a consistent estimator t_n in the previous strategies. However, when the set $\mathcal{M}$ is convex, there is typically no need for such approximate checks as Conditions 5.3.4 and 5.3.7 are automatically verified for a rather large class of ϕ-divergences and many situations of practical interest, as exposed in Section 5.3.3. □

Remark 5.3.12. Under Conditions 5.3.4 and 5.3.7, Theorem 5.3.1 and Corollary 5.3.10 imply that the functions $J[\vartheta^*]$ and $J[\mathcal{S}^*]$ are continuous at t_0. In statistical applications, given a consistent estimator t_n of (the unknown) t_0, we thus immediately obtain from the continuous mapping theorem that $J[\vartheta^*](t_0)$ (resp. $J[\mathcal{S}^*](t_0)$) is consistently estimated by $J[\vartheta^*](t_n)$ (resp. $J[\mathcal{S}^*](t_n)$). □

We conclude this subsection by stating two simple conditions which imply Condition 5.3.4 when combined with Condition 5.3.3.

Condition 5.3.13 (Interior)**.** *The function $\mathcal{S}$ in Condition 5.3.1 satisfies $\mathcal{S}(\mathring{\Theta}) = \mathcal{M} \cap (0,1)^m$.*

Condition 5.3.14 (Same support)**.** *For any $t \in (0,1)^m$, all ϕ-projections of t on $\mathcal{M} = \mathcal{S}(\bar{\Theta})$ belong to $(0,1)^m$.*

Condition 5.3.13 is satisfied in many applications (see Section 5.4) while the verification of Condition 5.3.14 is discussed in the forthcoming subsection. The proof of the next result is given in Section 5.B.

Lemma 5.3.15. *Conditions 5.3.3, 5.3.13 and 5.3.14 imply Condition 5.3.4.*

5.3.3 The Case of $\mathcal{M}$ Convex

Recall that Conditions 5.2.2, 5.2.12 and 5.3.1 are assumed to hold. Below, using results of Gietl and Reffel (2017), we first demonstrate (see Section 5.C) that Condition 5.3.3 is automatically satisfied under an additional assumption of convexity of the set $\mathcal{M}$ on which we wish to ϕ-project the vector $t_0 \in (0,1)^m$.

Proposition 5.3.16. *Assume that $\mathcal{M} = \mathcal{S}(\bar{\Theta}) \subset [0,1]^m$ is convex. Then, for any $t \in (0,1)^m$, the ϕ-projection s^* of t on $\mathcal{M} = \mathcal{S}(\bar{\Theta})$ (exists and) is unique, which implies that Condition 5.3.3 holds.*

Since Condition 5.3.3 holds when $\mathcal{M}$ is convex, from Lemma 5.3.15, it suffices to verify Conditions 5.3.13 and 5.3.14 to obtain Condition 5.3.4.

Remark 5.3.17. Restricting attention to probability vectors, Condition 5.3.14 holds for the Kullback–Leibler divergence (that is, for I-projections) since $\mathcal{M} = \mathcal{S}(\bar{\Theta})$ contains probability vectors in $(0,1)^m$ by construction, as $\mathcal{S}(\mathring{\Theta}) \subset \mathcal{M} \cap (0,1)^m$ and $\mathcal{S}(\mathring{\Theta}) \neq \emptyset$ by Condition 5.3.1. Indeed, when one wishes to I-project a probability vector q_0 on a closed convex subset containing at least one probability vector with the same support as q_0, it is known from Csiszár (1975, Theorem 2.1, Theorem 2.2 and Corollary 3.3) that the support of the unique I-projection will be equal to that of q_0. □

Let us now turn to Condition 5.3.7, the other key condition in Theorem 5.3.1. To verify it when t_0 is unknown (see Remark 5.3.11), one could rely on more general conditions that imply Condition 5.3.7. One such condition is as follows.

Condition 5.3.18 (Strong convexity of $D_\phi(\mathcal{S}(\cdot)\,|\,t)$)**.** *The set $\bar{\Theta}$ is convex and, for any $t \in (0,1)^m$, the function $D_\phi(\mathcal{S}(\cdot)\,|\,t)$ from $\bar{\Theta}$ to $\mathbb{R}$ is strongly convex.*

The next proposition is proven in Section 5.C.

Proposition 5.3.19. *Condition 5.3.18 implies that, for any $t \in (0,1)^m$ and $\theta \in \mathring{\Theta}$, the $k \times k$ matrix $J_2[D_\phi(\mathcal{S}(\cdot)\,|\,t)](\theta)$ is positive definite, and thus that Condition 5.3.7 holds.*

We have (from Proposition 5.2.13 and Table 5.2.1) that, for many ϕ-divergences and any $t \in (0,1)^m$, the function $D_\phi(\cdot\,|\,t)$ is strongly convex on $[0,1]^m$. Thus, Proposition 5.3.19 suggests to consider conditions on $\mathcal{S}$ and Θ under which Condition 5.3.18 holds. A simple such condition is as follows.

Condition 5.3.20 ($\mathcal{S}$ is affine)**.** *The set $\bar{\Theta}$ is convex and the function $\mathcal{S}$ from $\bar{\Theta}$ to $[0,1]^m$ is affine, that is, there exists an $m \times k$ matrix A and $\gamma \in \mathbb{R}^m$ such that $\mathcal{S}(\theta) = A\theta + \gamma$, $\theta \in \bar{\Theta}$.*

The following lemma is proven in Section 5.C.

Lemma 5.3.21 (Strong convexity of $D_\phi(\mathcal{S}(\cdot)\,|\,t)$)**.** *Conditions 5.2.11 and 5.3.20 imply that Condition 5.3.18 holds.*

Note that, from Lemma 5.2.14, Condition 5.2.11 can be viewed as a strengthening of Condition 5.2.2. The previous derivations lead to the following result, proven in Section 5.C.

Corollary 5.3.22 (Differentiability of ϑ^* and $\mathcal{S}^*$ at t_0)**.** *Assume that $\mathcal{M} = \mathcal{S}(\bar{\Theta})$ is convex and that Conditions 5.2.11, 5.3.13, 5.3.14 and 5.3.20 hold. Then, the function ϑ^* in (5.5) is continuously differentiable at $t_0 = (t_{0,1}, \ldots, t_{0,m}) \in (0,1)^m$ with Jacobian matrix at t_0 given by*

$$J[\vartheta^*](t_0) = \left[A^\top \operatorname{diag}\left(\frac{1}{t_{0,1}}\phi''\left(\frac{\mathcal{S}_1(\vartheta^*(t_0))}{t_{0,1}}\right), \ldots \right.\right. \\ \left.\left. \ldots, \frac{1}{t_{0,m}}\phi''\left(\frac{\mathcal{S}_m(\vartheta^*(t_0))}{t_{0,m}}\right)\right) A\right]^{-1} A^\top \Delta(t_0), \quad (5.11)$$

where $\Delta(t_0)$ is defined in (5.9), and the function $\mathcal{S}^*$ in (5.3) is continuously differentiable at t_0 with Jacobian matrix at t_0 given by $J[\mathcal{S}^*](t_0) = AJ[\vartheta^*](t_0)$.

Figure 5.3.1 summarises the various implications of conditions leading to Corollary 5.3.22. The only condition in that result that may not be easily verifiable is Condition 5.3.14. Its verification would obviously be immediate if t_0 was known and its ϕ-projection on $\mathcal{M}$ could be computed. However, as discussed in Remark 5.3.11, this is usually not the case in a statistical context. From Remark 5.3.17, though, we know that Condition 5.3.14 is automatically verified for the Kullback–Leibler divergence when we restrict attention to probability vectors. Using results of Rüschendorf (1987), we can prove (see Section 5.C) that this may also be the case for other ϕ-divergences provided that Condition 5.3.20 holds and Θ is defined by linear inequalities.

Proposition 5.3.23. *Assume that Condition 5.3.20 holds, $\lim_{x \to 0^+} \phi'(x) = -\infty$, $\Theta \subset \mathbb{R}^k$ is defined by linear inequalities and $\mathcal{M} = \mathcal{S}(\bar{\Theta}) \subset [0,1]^m$ is a convex subset of probability vectors. Then, for any probability vector $q \in (0,1)^m$, the (unique) ϕ-projection p^* of q on $\mathcal{M} = \mathcal{S}(\bar{\Theta})$ belongs to $(0,1)^m$.*

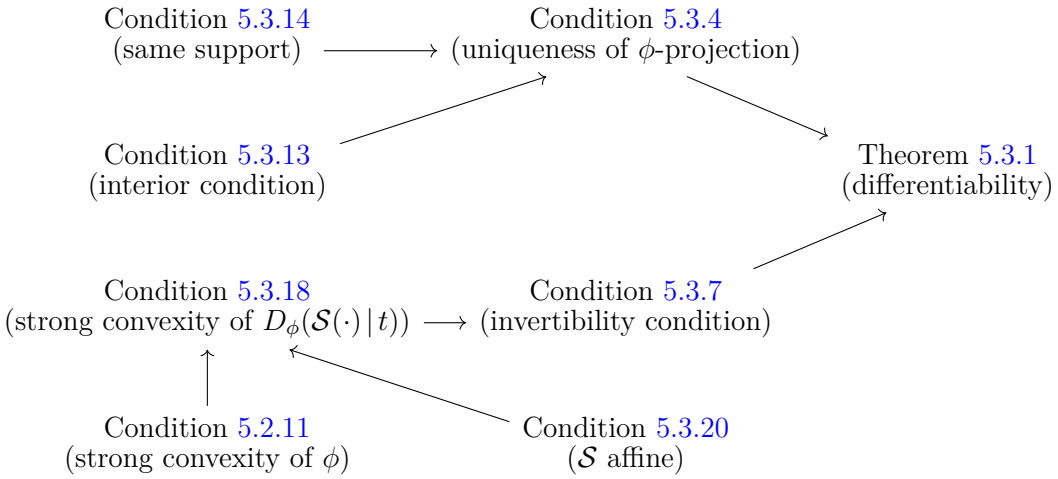

Figure 5.3.1: Diagram summarizing the various implications of conditions for differentiability when the set $\mathcal{M}$ of interest is convex.

Remark 5.3.24. For the ϕ-divergences listed in Table 5.2.1, $\lim_{x\to 0^+}\phi'(x) = -\infty$ holds for the Kullback–Leibler divergence, the squared Hellinger divergence, the reverse relative entropy, the Jensen–Shannon divergence, Neyman's χ^2 divergence and α-divergences with $\alpha < 1$.

□

The previous proposition leads to the following corollary, of importance for a large class of applications. It is proven in Section 5.C.

Corollary 5.3.25. *Assume that Conditions 5.2.11, 5.3.13 and 5.3.20 hold,* $\lim_{x\to 0^+}\phi'(x) = -\infty$, Θ *is defined by linear inequalities and* $\mathcal{M} = \mathcal{S}(\bar{\Theta}) \subset [0,1]^m$ *is a convex subset of probability vectors. Then, for any probability vector* $q \in (0,1)^m$, *the function* ϑ^* *in (5.5) is continuously differentiable at* q *with Jacobian matrix at* q *given by (5.11) with* $t_0 = q$, *and the function* $\mathcal{S}^*$ *in (5.3) is continuously differentiable at* q *with Jacobian matrix at* q *given by* $J[\mathcal{S}^*](q) = AJ[\vartheta^*](q)$.

5.3.4 The Case of $\mathcal{M}$ Defined By Linear Equalities

Many applications involve ϕ-projections on sets defined by linear equalities (see Sections 5.4.2 and 5.4.3). The proposition below, which states that in such cases Conditions 5.3.1, 5.3.13 and 5.3.20 are automatically verified, is therefore of great interest. It is proven in Section 5.C.

Proposition 5.3.26. *Assume that the set $\mathcal{M}$ is defined by linear equalities, that is, that there exist $d \in \mathbb{N}^+$, $\alpha = (\alpha_1, \ldots, \alpha_d) \in \mathbb{R}^d$ and d functions $f_1, f_2, \ldots, f_d$ on $\{1, \ldots, m\}$ such that*

$$\mathcal{M} = \left\{ s \in [0,1]^m : \sum_{i=1}^{m} s_i f_j(i) = \alpha_j, j \in \{1, \ldots, d\} \right\}.$$

Moreover, suppose that $\mathcal{M} \cap (0,1)^m$ is nonempty and $\mathcal{M}$ is not reduced to a singleton. Then Conditions 5.3.1, 5.3.13 and 5.3.20 hold.

Note that the previous result does not directly provide the expression of the function $\mathcal{S}$ in Condition 5.3.1 which is necessary to compute the Jacobians in Corollaries 5.3.22 and 5.3.25. It merely guarantees that, whenever $\mathcal{M}$ is defined by linear equalities, the framework for studying the differentiability of ϕ-projections set up in Section 5.3.1 applies with $\mathcal{S}$ affine and Condition 5.3.13 additionally satisfied.

5.4 Applications To Minimum ϕ-Divergence Estimators

This section illustrates the application of the results of Section 5.3 to the derivation of the limiting distributions of properly scaled ϕ-projection estimators in three different situations. To indicate that the element of $(0,1)^m$ to be ϕ-projected is a probability vector, it will be denoted by q_0. In all three applications, after defining the function $\mathcal{S}$ in Condition 5.3.1 and verifying the conditions of Corollary 5.3.10, Corollary 5.3.22 or Corollary 5.3.25, we shall compute the Jacobian matrix of the function $\mathcal{S}^*$ in (5.3) at q_0, that is, in the notation of Section 5.3, $J[\mathcal{S}^*](q_0)$. Given an estimator q_n of q_0, $\mathcal{S}^*(q_n)$ is the ϕ-projection and/or minimum ϕ-divergence estimator of q_0 (that is, the estimator of $\mathcal{S}^*(q_0)$) and, if $\sqrt{n}(q_n - q_0)$ converges weakly as n tends to infinity, it follows immediately with the delta method (see, e.g., van der Vaart, 1998, Theorem 3.1) that

$$\sqrt{n}(\mathcal{S}^*(q_n) - \mathcal{S}^*(q_0)) = J[\mathcal{S}^*](q_0)\sqrt{n}(q_n - q_0) + o_{\mathrm{P}}(1). \tag{5.12}$$

For illustration purposes, we shall specifically consider a random variable X taking its values in a set $\mathcal{X} = \{x_1, \ldots, x_m\}$ of interest such that $\mathrm{P}(\{X = x_i\}) = q_{0,i}$, $i \in \{1, \ldots, m\}$, and we shall assume that we have at hand n independent copies $X_1, \ldots, X_n$ of X. A natural estimator of q_0 is then the

probability vector q_n whose components are

$$q_{n,i} = \frac{1}{n}\sum_{j=1}^{n} \mathbf{1}(X_j = x_i), \qquad i \in \{1,\ldots,m\},$$

where $\mathbf{1}$ denotes the indicator function. The multivariate central limit theorem establishes that, as n tends to infinity, $\sqrt{n}(q_n - q_0)$ converges weakly to an m-dimensional centered normal distribution with covariance matrix $\Sigma_{q_0} = \mathrm{diag}(q_0) - q_0 q_0^\top$. As a result, we obtain from (5.12) that, as n tends to infinity, $\sqrt{n}(\mathcal{S}^*(q_n) - \mathcal{S}^*(q_0))$ converges weakly to an m-dimensional centered normal distribution with covariance matrix $\Sigma = J[\mathcal{S}^*](q_0)\Sigma_{q_0} J[\mathcal{S}^*](q_0)^\top$. To empirically verify the correctness of the asymptotic covariance matrix Σ, in all three applications, we shall finally compare it with $\Sigma_{n,N}$, the sample covariance matrix computed from $N = 5000$ independent replicates of $\sqrt{n}\mathcal{S}^*(q_n)$ for $n = 5000$.

The first application below is prototypical of the use of minimum ϕ-divergence estimators for parametric inference: it is about ϕ-projecting a probability vector q_0 on the set $\mathcal{M}$ of probability vectors generated from binomial distributions whose first parameter is fixed to $m-1$. As the set $\mathcal{M}$ is not convex in this case, we will have to rely on the general results of Section 5.3.2, and on Corollary 5.3.10 in particular. In the second and third scenarios, the set $\mathcal{M}$ is defined by linear equalities (and is thus convex) which allows using the results of Sections 5.3.3 and 5.3.4. Specifically, in the second (resp. third) application, $\mathcal{M}$ is the set of probability vectors generated from distributions with certain moments fixed (resp. the Fréchet class of all bivariate probability arrays with given univariate margins).

5.4.1 ϕ-projection on the Set of Binomial Probability Vectors

We first consider a parametric inference scenario in which one wishes to ϕ-project a probability vector $q_0 \in (0,1)^m$ of interest on a given parametric model. As an example, for $m > 1$, we consider the set of probability vectors generated from binomial distributions whose first parameter is fixed to $m-1$. Specifically, for any $\ell \in \mathbb{N}_0$ and $i \in \{0,\ldots,\ell\}$, let

$$p_{\ell,i}(\theta) = \binom{\ell}{i}\theta^i(1-\theta)^{\ell-i}, \qquad \theta \in [0,1],$$

and let $\mathcal{M} = \{(p_{m-1,0}(\theta),\ldots,p_{m-1,m-1}(\theta)) : \theta \in [0,1]\}$. The set Θ in Condition 5.3.1 is thus $[0,1]$ and the ith component function $\mathcal{S}_i$, $i \in \{1,\ldots,m\}$, of

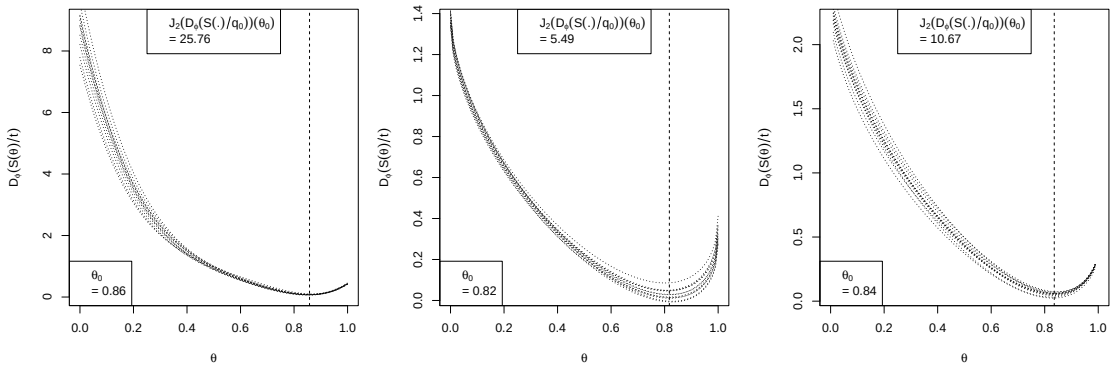

Figure 5.4.1: For Pearson's χ^2 divergence (left), the squared Hellinger divergence (middle) and the Kullback–Leibler divergence (right), graphs of the functions $D_\phi(\mathcal{S}(\cdot)\,|\,t)$ for 10 vectors $t\in(0,1)^3$ of the form $t=q_0+(z_1,\ldots,z_m)$, where $(z_1,\ldots,z_m)$ is drawn from the m-dimensional centered normal distribution with covariance matrix $0.01^2 I_m$. In each plot, the vertical dashed line marks the value θ_0 at which the function $D_\phi(\mathcal{S}(\cdot)\,|\,q_0)$ reaches its minimum and the top insert contains the (strictly positive) value of $J_2[D_\phi(\mathcal{S}(\cdot)\,|\,q_0)](\theta_0)$.

the function $\mathcal{S}$ is $p_{m-1,i-1}$. Furthermore, for any $\ell\in\mathbb{N}_0$, $i\in\{0,\ldots,\ell\}$ and $\theta\in(0,1)$,

$$\begin{aligned}p'_{\ell,i}(\theta)&=\binom{\ell}{i}\left[\mathbf{1}(i>0)i\theta^{i-1}(1-\theta)^{\ell-i}-\mathbf{1}(i<\ell)\theta^i(\ell-i)(1-\theta)^{\ell-i-1}\right]\\&=\ell\left[\mathbf{1}(i\geq 1)p_{\ell-1,i-1}(\theta)-\mathbf{1}(i\leq\ell-1)p_{\ell-1,i}(\theta)\right],\end{aligned}$$

and

$$p''_{\ell,i}(\theta)=\ell\left[\mathbf{1}(i\geq 1)p'_{\ell-1,i-1}(\theta)-\mathbf{1}(i\leq\ell-1)p'_{\ell-1,i}(\theta)\right].$$

From the above, we obtain that, for any $\theta\in(0,1)$,

$$J[\mathcal{S}](\theta)=(p'_{m-1,0}(\theta),\ldots,p'_{m-1,m-1}(\theta))$$

and

$$J_2[\mathcal{S}](\theta)=(p''_{m-1,0}(\theta),\ldots,p''_{m-1,m-1}(\theta)).$$

As an illustration, we take $m=3$ and $q_0=(0.1,0.2,0.7)\notin\mathcal{M}=\mathcal{S}(\bar{\Theta})$. As ϕ-divergence, we consider Pearson's χ^2 divergence, the squared Hellinger divergence and the Kullback–Leibler divergence (see Table 5.2.1). To be able to apply Corollary 5.3.10 and compute $J[\mathcal{S}^*](q_0)$ given in (5.10) (to eventually compute $\Sigma=J[\mathcal{S}^*](q_0)\Sigma_{q_0}J[\mathcal{S}^*](q_0)^\top$), we need to verify Conditions 5.3.4 and 5.3.7. We proceed as suggested in Remark 5.3.11. To empirically verify Condition 5.3.4, we plot the function $D_\phi(\mathcal{S}(\cdot)\,|\,t)$ for 10 vectors $t\in(0,1)^3$ in

a "neighborhood" of q_0. Specifically, we consider vectors t of the form $t = q_0 + (z_1, \ldots, z_m)$, where $(z_1, \ldots, z_m)$ is drawn from the m-dimensional centered normal distribution with covariance matrix $0.01^2 I_m$. Figure 5.4.1 reveals that Condition 5.3.4 seems to hold for all three divergences. In all three cases, the value θ_0 at which $D_\phi(\mathcal{S}(\cdot) \,|\, q_0)$ reaches its minimum is found numerically using the `optim()` function of the R statistical environment (R Core Team, 2024) and its rounded value is reported in the lower left corner of each plot. The rounded value of $J_2[D_\phi(\mathcal{S}(\cdot) \,|\, q_0)](\theta_0)$ is also provided in each plot of Figure 5.4.1 and confirms that Condition 5.3.7 is likely to hold as well. For all three ϕ-divergences, the asymptotic covariance matrix Σ of $\sqrt{n}(\mathcal{S}^*(q_n) - \mathcal{S}^*(q_0))$ and its approximation $\Sigma_{n,N}$ show very good agreement. We print their rounded versions below for Pearson's χ^2 divergence:

$$\Sigma = \begin{bmatrix} 0.010 & 0.049 & -0.059 \\ 0.049 & 0.246 & -0.295 \\ -0.059 & -0.295 & 0.353 \end{bmatrix}, \qquad \Sigma_{n,N} = \begin{bmatrix} 0.010 & 0.049 & -0.059 \\ 0.049 & 0.248 & -0.297 \\ -0.059 & -0.297 & 0.356 \end{bmatrix}.$$

5.4.2 ϕ-projection on the Set Of Probability Vectors Generated from Distributions with Certain Moments Fixed

Let X be a random variable whose support is $\mathcal{X} = \{x_1, x_2, \ldots, x_m\}$ with $x_1 < \cdots < x_m$ and let $q_0 \in (0,1)^m$ be defined by $q_{0,i} = \mathrm{P}(\{X = x_i\})$, $i \in \{1, \ldots, m\}$. Assume that the aim is to ϕ-project q_0 on the set $\mathcal{M}$ of probability vectors obtained from distributions on $\mathcal{X}$ whose $r < m-1$ first raw (non-centered) moments are fixed to some known values $\mu_1, \ldots, \mu_r$, that is,

$$\mathcal{M} = \left\{ p \in [0,1]^m : \sum_{i=1}^m p_i = 1, \sum_{i=1}^m x_i^j p_i = \mu_j, j \in \{1, \ldots, r\} \right\}. \tag{5.13}$$

For the previous problem to be well-defined, it is further necessary to assume that $\mathcal{M}$ is nonempty. Note that the above setting is strongly related to the so-called *moment problem*. In the current finite discrete setting, Tagliani (2000, 2001) proposed to address it using the maximum entropy principle (Jaynes, 1957). It appears that his approach is a particular case of the considered ϕ-projection approach when D_ϕ in (5.2) is the Kullback–Leibler divergence and q_0 is the uniform distribution on $\mathcal{X}$.

The set $\mathcal{M}$ in (5.13) is clearly convex. To illustrate the usefulness of the differentiability results stated in Section 5.3.3, we shall further assume

that $\mathcal{M} \cap (0,1)^m$ is nonempty. Since $\mathcal{M}$ is defined by linear equalities, Proposition 5.3.26 guarantees that Conditions 5.3.1 and 5.3.13 hold for some set Θ and some affine function $\mathcal{S}$ to be determined. Moreover, let $\mu = (1, \mu_1, \ldots, \mu_r) \in \mathbb{R}^{r+1}$ and, for any $\ell \in \{1, \ldots, m\}$, let U_ℓ be the $\ell \times m$ matrix whose ith row is $(x_1^{i-1}, \ldots, x_m^{i-1})$, $i \in \{1, \ldots, \ell\}$. With the above notation, $\mathcal{M} = \{p \in [0,1]^m : U_{r+1} p = \mu\}$. Next, take an element p of $\mathcal{M}$ and denote by $\theta_1, \ldots, \theta_{m-r-1}$ the corresponding raw (non-centered) moments of order $r+1, \ldots, m-1$, respectively. Then,

$$U_m p = (\mu, \theta),$$

where $\theta = (\theta_1, \ldots, \theta_{m-r-1}) \in \mathbb{R}^{m-r-1}$ and

$$(\mu, \theta) = (1, \mu_1, \ldots, \mu_r, \theta_1, \ldots, \theta_{m-r-1}) \in \mathbb{R}^m.$$

Since U_m is a Vandermonde matrix with $x_i \neq x_j$ for $i, j \in \{1, \ldots, m\}$, $i \neq j$, it is invertible (see, e.g., Golub and Loan, 2013, p 203) and we have

$$p = U_m^{-1}(\mu, \theta)$$

confirming that any p.m.f. on $\mathcal{X}$ is perfectly defined by its $m-1$ raw moments. As a consequence, any $p \in \mathcal{M}$ can be parametrised by a vector $\theta \in \mathbb{R}^{m-r-1}$ (of raw moments) in

$$\Theta = \left\{\theta \in \mathbb{R}^{m-r-1} : 0_{\mathbb{R}^m} \leq U_m^{-1}(\mu, \theta) \leq 1_{\mathbb{R}^m}\right\}$$

and the function $\mathcal{S}$ from $\Theta = \bar{\Theta}$ to $[0,1]^m$ such that $\mathcal{M} = \mathcal{S}(\bar{\Theta})$ is given by

$$\mathcal{S}(\theta) = U_m^{-1}(\mu, \theta), \qquad \theta \in \bar{\Theta}.$$

Some thought reveals that $\mathcal{S}(\mathring{\Theta}) = \mathcal{M} \cap (0,1)^m$, that is, Condition 5.3.13 holds as expected.

Let V_m (resp. W_m) be the $m \times (r+1)$ (resp. $m \times (m-r-1)$) matrix obtained from U_m^{-1} by keeping its first $r+1$ (resp. last $m-r-1$) columns. Then, note that Θ and $\mathcal{S}$ can be equivalently expressed as

$$\Theta = \left\{\theta \in \mathbb{R}^{m-r-1} : -V_m \mu \leq W_m \theta \leq 1_{\mathbb{R}^m} - V_m \mu\right\} \tag{5.14}$$

and

$$\mathcal{S}(\theta) = W_m \theta + V_m \mu, \qquad \theta \in \bar{\Theta}, \tag{5.15}$$

respectively. From the previous expressions, we have that $\bar{\Theta}$ is convex, $\mathcal{S}$ is affine (that is, Condition 5.3.20 holds) and the Jacobian matrix of $\mathcal{S}$ at any $\theta \in \mathring{\Theta}$ is constant and equal to W_m.

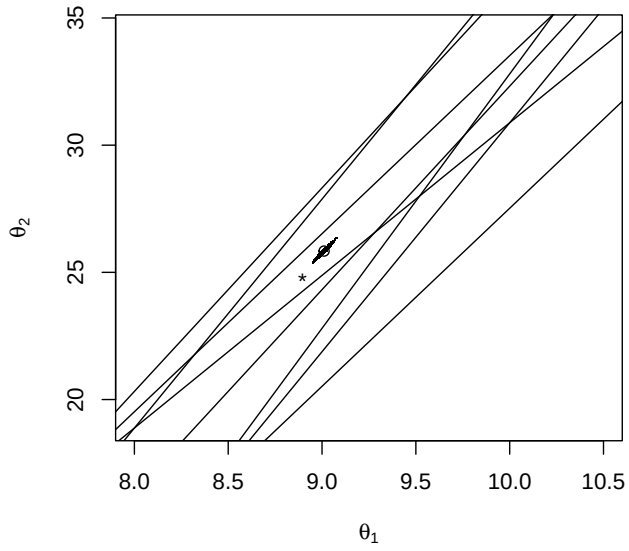

Figure 5.4.2: Lines representing the linear inequalities appearing in the definition of Θ in (5.14). The symbol '$\star$' represents the point $(8.896, 24.8704) \in \mathring{\Theta}$ corresponding to the probability vector of the binomial distribution with parameters $m-1$ and 0.4. For the squared Hellinger divergence, the vector θ_0 at which $D_\phi(\mathcal{S}(\cdot) \,|\, q_0)$ reaches its minimum is represented by the symbol 'o'. The small oblique cloud of points around 'o' consists of realizations of the value θ_n at which the function $D_\phi(\mathcal{S}(\cdot) \,|\, q_n)$ reaches its minimum.

As an illustration, we take $m = 5$, $\mathcal{X} = \{0, \ldots, m-1\}$, $r = 2$ and $\mu = (1, 1.6, 3.52)$. The linear inequalities appearing in the definition of Θ in (5.14) are then represented in Figure 5.4.2. Note that $\mathcal{M} = \mathcal{S}(\bar{\Theta})$ is nonempty as $\mathcal{S}(\mathring{\Theta})$ contains the probability vector of the binomial distribution with parameters $m-1$ and 0.4 (whose raw moments of order 1, 2, 3 and 4 are 1.6, 3.52, 8.896 and 24.8704, respectively). This also implies that $(8.896, 24.8704) \in \mathring{\Theta}$. The latter point is represented by the symbol '$\star$' in Figure 5.4.2. For q_0, we take $(0.35, 0.3, 0.15, 0.1, 0.1)$. Note that q_0 does not belong to $\mathcal{M}$ as the corresponding raw moments of order 1, 2, 3 and 4 are 1.3, 3.4, 10.6 and 36.4, respectively. As in the previous application, as ϕ-divergence, we consider Pearson's χ^2 divergence, the squared Hellinger divergence and the Kullback–Leibler divergence. In all three cases, we use the `constrOptim()` function of the R statistical environment to find the vector θ_0 at which $D_\phi(\mathcal{S}(\cdot) \,|\, q_0)$ reaches its minimum. For Pearson's χ^2 divergence (resp. the squared Hellinger divergence, the Kullback–Leibler divergence), after rounding to two decimal

places, we obtain (9.18,26.84) (resp. (9.01,25.84), (9.07,26.19)). For the squared Hellinger divergence, θ_0 is represented by the symbol 'o' in Figure 5.4.2. Additional calculations show that, in all three cases, $\mathcal{S}(\theta_0) \in (0,1)^m$, which implies that Condition 5.3.14 holds. Note that, from Table 5.2.1, Condition 5.2.11 holds for all three divergences under consideration and that, from (5.14) and (5.15), Condition 5.3.20 holds as well. Proposition 5.3.23 and Remark 5.3.24 then imply that Condition 5.3.14 automatically holds for the Kullback–Leibler and the squared Hellinger divergences. For Pearson's χ^2 divergence, however, the additional calculations mentioned above cannot be dispensed with.

For the χ^2 divergence (resp. the Kullback–Leibler divergence and the squared Hellinger divergence), we can apply Corollary 5.3.22 (resp. Corollary 5.3.25) and compute $J[\mathcal{S}^*](q_0)$ as given therein as well as $\Sigma = J[\mathcal{S}^*](q_0)\Sigma_{q_0}J[\mathcal{S}^*](q_0)^\top$. Again, for all three ϕ-divergences, the asymptotic covariance matrix Σ of $\sqrt{n}(\mathcal{S}^*(q_n) - \mathcal{S}^*(q_0))$ and its approximation $\Sigma_{n,N}$ showed very good agreement. We print their rounded versions below for the squared Hellinger divergence:

$$\Sigma = \begin{bmatrix} 0.021 & -0.058 & 0.048 & -0.005 & -0.005 \\ -0.058 & 0.168 & -0.154 & 0.037 & 0.007 \\ 0.048 & -0.154 & 0.176 & -0.080 & 0.011 \\ -0.005 & 0.037 & -0.080 & 0.070 & -0.022 \\ -0.005 & 0.007 & 0.011 & -0.022 & 0.009 \end{bmatrix},$$

$$\Sigma_{n,N} = \begin{bmatrix} 0.021 & -0.059 & 0.048 & -0.006 & -0.005 \\ -0.059 & 0.169 & -0.156 & 0.039 & 0.007 \\ 0.048 & -0.156 & 0.179 & -0.083 & 0.011 \\ -0.006 & 0.039 & -0.083 & 0.072 & -0.022 \\ -0.005 & 0.007 & 0.011 & -0.022 & 0.009 \end{bmatrix}.$$

5.4.3 ϕ-projection on a Fréchet Class of Bivariate Probability Arrays

In this last application we essentially generalise the setting of the previous chapter, albeit with a slightly different notation which is more congenial to the present framework. Let $r, s \in \mathbb{N}^+$. The initial form of the probability

vector to be ϕ-projected is a $r \times s$ matrix $\underline{q}_0$. We further assume that all its elements are strictly positive, that is, $\underline{q}_0 \in (0,1)^{r\times s}$. Let $a \in (0,1)^r$ and $b \in (0,1)^s$ be two fixed probability vectors. This last illustration is about ϕ-projecting the bivariate probability array $\underline{q}_0$ on the Fréchet class $\underline{\mathcal{M}} \subset [0,1]^{r\times s}$ of bivariate probability arrays with univariate margins a and b, which, after the strong suggestion in Geenens (2020), we have used as the main building block of the models presented in the previous chapter. The slight differences of the latter with the approach here delineated will be discussed in Remark 5.4.1. Furthermore, it is a problem which arises in many areas of probability and statistics (see, e.g., Csiszár, 1975; Little and Wu, 1991; Vajda and van der Meulen, 2005). To be able to exploit the results of Section 5.3.3, we will need to consider a vectorised version of this problem.

Let us first define the set $\underline{\mathcal{M}}$ in the spirit of Condition 5.3.1. Let

$$\underline{\Theta} = \left\{ \underline{\theta} \in [0,1]^{(r-1)\times(s-1)} : \sum_{i=1}^{r-1} \underline{\theta}_{ij} \leq b_j, j \in \{1,\dots,s-1\}, \right. \\ \left. \sum_{j=1}^{s-1} \underline{\theta}_{ij} \leq a_j, i \in \{1,\dots,r-1\} \right\}.$$

Note that $\underline{\Theta}$ is closed and convex. Some thought reveals that the function $\underline{\mathcal{S}}$ from $\underline{\Theta}$ to $[0,1]^{r\times s}$ defined by

$$\underline{\mathcal{S}}(\underline{\theta}) = \begin{bmatrix} \theta_{11} & \dots & \theta_{1,s-1} & a_1 - \sum_{j=1}^{s-1}\theta_{1j} \\ \vdots & & \vdots & \vdots \\ \theta_{r-1,1} & \dots & \theta_{r-1,s-1} & a_{r-1} - \sum_{j=1}^{s-1}\theta_{r-1,j} \\ b_1 - \sum_{i=1}^{r-1}\theta_{i1} & \dots & b_{s-1} - \sum_{i=1}^{r-1}\theta_{i,s-1} & a_r + b_s - 1 + \sum_{i=1}^{r-1}\sum_{j=1}^{s-1}\theta_{ij} \end{bmatrix}$$

then generates the Fréchet class $\underline{\mathcal{M}} \subset [0,1]^{r\times s}$ of bivariate probability arrays with univariate margins a and b, that is, $\underline{\mathcal{M}} = \underline{\mathcal{S}}(\bar{\underline{\Theta}})$. Furthermore, some thought reveals that $\underline{\mathcal{S}}(\mathring{\underline{\Theta}}) = \underline{\mathcal{M}} \cap (0,1)^{r\times s}$ and $\underline{\mathcal{M}}$ is not empty as $\underline{\mathcal{S}}(\mathring{\underline{\Theta}})$ contains the bivariate probability array $ab^\top$. Finally, $\underline{\mathcal{M}}$ is compact (since $\underline{S}$ is continuous and $\bar{\underline{\Theta}}$ is compact) and convex.

For $c,d \in \mathbb{N}^+$, let $\mathrm{vec}_{c,d}$ be the vectorization operator, which, given a matrix in $\mathbb{R}^{c\times d}$, returns its column-major vectorization in $\mathbb{R}^{cd}$ and let $\mathrm{vec}_{c,d}^{-1}$ denote the inverse operator. Next, let $\Theta = \mathrm{vec}_{r-1,s-1}(\underline{\Theta})$ and let $\mathcal{S}$ be the function from Θ to $[0,1]^{rs}$ defined by

$$\mathcal{S}(\theta) = \mathrm{vec}_{r,s} \circ \underline{\mathcal{S}} \circ \mathrm{vec}_{r-1,s-1}^{-1}(\theta), \qquad \theta \in \Theta.$$

Some thought then reveals that

$$\mathcal{S}(\theta) = A\theta + \mathcal{S}(0_{\mathbb{R}^{(r-1)(s-1)}}), \qquad \theta \in \Theta,$$

where A is the $rs \times (r-1)(s-1)$ matrix given by

$$\begin{bmatrix} Q & 0 & \dots & 0 \\ 0 & Q & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & Q \\ -Q & -Q & \dots & -Q \end{bmatrix} \text{ with } Q = \begin{bmatrix} 1 & 0 & \dots & 0 \\ 0 & 1 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 1 \\ -1 & -1 & \dots & -1 \end{bmatrix} \in \mathbb{R}^{r \times (r-1)}.$$

Hence, Condition 5.3.1 holds with $m = rs$ and $k = (r-1)(s-1)$, and the resulting set $\mathcal{M} = \mathcal{S}(\bar{\Theta})$ is simply $\mathrm{vec}_{r,s}(\underline{\mathcal{M}})$. Since $\underline{\mathcal{S}}(\mathring{\Theta}) = \underline{\mathcal{M}} \cap (0,1)^{r \times s}$, Condition 5.3.13 is also satisfied. Finally, Condition 5.3.20 holds as well since $\Theta = \bar{\Theta}$ is convex and $\mathcal{S}$ is affine. We shall thus be able to apply Corollary 5.3.22 provided Conditions 5.2.11 and 5.3.14 hold. Alternatively, since Θ is defined by linear inequalities, we could also rely on Corollary 5.3.25 provided Condition 5.2.11 and $\lim_{x \to 0^+} \phi'(x) = -\infty$ hold.

Let us now ϕ-project $\underline{q}_0$ on $\underline{\mathcal{M}}$ or, equivalently, $q_0 = \mathrm{vec}_{r,s}(\underline{q}_0)$ on $\mathcal{M}$. As an example, we take $r = s = 3$, $a = (0.2, 0.3, 0.5)$, $b = (0.5, 0.25, 0.25)$ and

$$\underline{q}_0 = \begin{bmatrix} 0.04 & 0.11 & 0.13 \\ 0.10 & 0.07 & 0.08 \\ 0.14 & 0.12 & 0.21 \end{bmatrix}.$$

To simplify computations, we shall only use the Kullback–Leibler divergence here, as the I-projection of a bivariate probability array on a Fréchet class can be conveniently carried out in practice using the iterative proportional fitting procedure (IPFP), also known as Sinkhorn's algorithm or matrix scaling in the literature (see, e.g., Pukelsheim, 2014; Idel, 2016; Brossard and Leuridan, 2018). Specifically, we use the function `Ipfp()` of the R package `mipfp` (Barthélemy and Suesse, 2018). The I-projection of $\underline{q}_0$ on $\underline{\mathcal{M}}$, rounded to two decimal places, is found to be

$$\underline{p}_0 = \begin{bmatrix} 0.06 & 0.08 & 0.06 \\ 0.18 & 0.07 & 0.06 \\ 0.27 & 0.10 & 0.13 \end{bmatrix},$$

so that the vector θ_0 (rounded to two decimal places) at which $D_\phi(\mathcal{S}(\cdot)\,|\,q_0)$ reaches its minimum on $\mathring{\Theta}$ is $(0.06, 0.18, 0.08, 0.07)$. Since Condition 5.2.11 and $\lim_{x\to 0^+}\phi'(x) = -\infty$ hold for the Kullback–Leibler divergence (see Table 5.2.1), we can directly apply Corollary 5.3.25 and compute $J[\mathcal{S}^*](q_0)$ as well as the asymptotic covariance matrix $\Sigma = J[\mathcal{S}^*](q_0)\Sigma_{q_0}J[\mathcal{S}^*](q_0)^\top$. The rounded version of the latter as well as of its approximation $\Sigma_{n,N}$ are:

$$\Sigma = \begin{bmatrix} 0.041 & -0.013 & -0.028 & -0.026 & 0.009 & 0.017 & -0.014 & 0.004 & 0.011 \\ -0.013 & 0.059 & -0.047 & 0.009 & -0.037 & 0.028 & 0.003 & -0.023 & 0.019 \\ -0.028 & -0.047 & 0.075 & 0.017 & 0.028 & -0.045 & 0.011 & 0.019 & -0.030 \\ -0.026 & 0.009 & 0.017 & 0.034 & -0.011 & -0.023 & -0.007 & 0.002 & 0.006 \\ 0.009 & -0.037 & 0.028 & -0.011 & 0.039 & -0.028 & 0.002 & -0.002 & 0.001 \\ 0.017 & 0.028 & -0.045 & -0.023 & -0.028 & 0.051 & 0.005 & 0.001 & -0.006 \\ -0.014 & 0.003 & 0.011 & -0.007 & 0.002 & 0.005 & 0.022 & -0.005 & -0.017 \\ 0.004 & -0.023 & 0.019 & 0.002 & -0.002 & 0.001 & -0.005 & 0.025 & -0.020 \\ 0.011 & 0.019 & -0.030 & 0.006 & 0.001 & -0.006 & -0.017 & -0.020 & 0.036 \end{bmatrix},$$

$$\Sigma_{n,N} = \begin{bmatrix} 0.040 & -0.013 & -0.027 & -0.025 & 0.009 & 0.016 & -0.015 & 0.004 & 0.011 \\ -0.013 & 0.059 & -0.046 & 0.010 & -0.037 & 0.027 & 0.003 & -0.022 & 0.019 \\ -0.027 & -0.046 & 0.073 & 0.016 & 0.028 & -0.043 & 0.012 & 0.018 & -0.030 \\ -0.025 & 0.010 & 0.016 & 0.033 & -0.011 & -0.022 & -0.007 & 0.001 & 0.006 \\ 0.009 & -0.037 & 0.028 & -0.011 & 0.039 & -0.028 & 0.002 & -0.002 & 0.000 \\ 0.016 & 0.027 & -0.043 & -0.022 & -0.028 & 0.050 & 0.006 & 0.001 & -0.006 \\ -0.015 & 0.003 & 0.012 & -0.007 & 0.002 & 0.006 & 0.022 & -0.005 & -0.017 \\ 0.004 & -0.022 & 0.018 & 0.001 & -0.002 & 0.001 & -0.005 & 0.024 & -0.019 \\ 0.011 & 0.019 & -0.030 & 0.006 & 0.000 & -0.006 & -0.017 & -0.019 & 0.036 \end{bmatrix},$$

showing again very good agreement.

Remark 5.4.1. Let us briefly explain the slight difference that exists between the above application of Corollary 5.3.22 to I-projections on Fréchet classes and what we have presented in the last chapter. Essentially, instead of obtaining differentiability results for the function $\mathcal{S}^*$ in (5.3) at q_0, in Section 4.2.4 of Chapter 4 we have provided differentiability results for the function

$$(q_1, \ldots, q_{rs-1}) \mapsto \mathcal{S}^*\left(q_1, \ldots, q_{rs-1}, 1 - \sum_{i=1}^{rs-1} q_i\right)$$

at q_0. By the chain rule, the Jacobian of the latter function at q_0 is simply $J[\mathcal{S}^*](q_0)R$, where $J[\mathcal{S}^*](q_0)$ is given in Corollary 5.3.22 and R is the $rs \times (rs-1)$ matrix obtained by adding a row of -1's below the identity matrix

I_{rs-1}. In other words, in Section 4.2.4 of Chapter 4 we have somehow bypassed the parametrisation aspect individuated by the definition of Θ and the continuous injection $\mathcal{S}$, as instead done in this chapter, and we have proceeded in a more direct way. However, the results of the previous chapter acted as the starting point for the developments we have presented here. Furthermore, the use of the vectorisation operator as a tool to employ the results of this chapter for the problem of the Frechèt class of pmfs, which is naturally usually stated using matrices, is particularly enlightening when aiming to generalise the methodology delineated in Chapter 4 to the multivariate case, as shall be explained in the future work section of this thesis. □

This chapter concludes the treatment of the line of our research that started by the development of copula-like models for finite discrete bivariate random vectors, of which some results prompted us to study the generalisations presented here. Clearly, the latter investigations have not been conducted in a context that directly connects with the hydrological underpinnings of this thesis, which, for instance, initially led us to consider the problem of copulas in the case of discrete margins. In the following chapter we shall briefly return to a hydrological motivation behind the study of another theoretical and applicable tool.

Appendix

5.A Proofs of the results of Section 5.2

Proof of Proposition 5.2.13. First, notice that the set $[0,1]^m_{\mathrm{supp}(t)}$ is indeed convex since any convex combination of elements of $[0,1]^m_{\mathrm{supp}(t)}$ is in $[0,1]^m$ and has support included in $\mathrm{supp}(t)$. Next, recall the definition of f in (5.1) and let us first show that Condition 5.2.11 implies that, for any $w \in (0,\infty)$, $f(\cdot,w)$ is strongly convex on $[0,1]$. Fix $w \in (0,\infty)$ and let $\kappa_\phi(w)$ be the strong convexity constant of the restriction of ϕ to $[0,1/w]$. Then, for any $v,v' \in [0,1]$ and $\alpha \in [0,1]$,

$$\begin{aligned} f(\alpha v + (1-\alpha)v', w) &= w\phi\left(\alpha\frac{v}{w} + (1-\alpha)\frac{v'}{w}\right) \\ &\leq w\alpha\phi\left(\frac{v}{w}\right) + w(1-\alpha)\phi\left(\frac{v'}{w}\right) - w\frac{\kappa_\phi(w)}{2}\alpha(1-\alpha)\left|\frac{v}{w} - \frac{v'}{w}\right|^2 \\ &= \alpha f(v,w) + (1-\alpha) f(v',w) - \frac{\kappa_\phi(w)}{2w}\alpha(1-\alpha)|v-v'|^2. \end{aligned}$$

Hence, for any $w \in (0,\infty)$, $f(\cdot,w)$ is strongly convex on $[0,1]$ with strong convexity constant $\kappa_\phi(w)/w$.

Now, let $t \in [0,1]^m$, $t \neq 0_{\mathbb{R}^m}$, and notice that, for any $s \in [0,1]^m_{\mathrm{supp}(t)}$, all the summands in the expression of $D_\phi(s\,|\,t)$ in (5.2) corresponding to the t_i's that are not strictly positive are zero because of the definition of f in (5.1). As a consequence, without loss of generality, we can assume that $t \in (0,1]^m$.

Hence, for any $t \in (0,1]^m$, $s, s' \in [0,1]^m$ and $\alpha \in [0,1]$,

$$\begin{aligned} D_\phi(\alpha s + (1-\alpha)s' \,|\, t) &= \sum_{i=1}^m f(\alpha s_i + (1-\alpha)s'_i, t_i) \\ &\leq \alpha \sum_{i=1}^m f(s_i, t_i) + (1-\alpha) \sum_{i=1}^m f(s'_i, t_i) - \alpha(1-\alpha) \sum_{i=1}^m \frac{\kappa_\phi(t_i)}{2t_i} |s_i - s'_i|^2 \\ &\leq \alpha D_\phi(s \,|\, t) + (1-\alpha) D_\phi(s' \,|\, t) - \alpha(1-\alpha) \|s - s'\|_2^2 \min_{i \in \{1,\dots,m\}} \frac{\kappa_\phi(t_i)}{2t_i}. \end{aligned}$$

Hence, for any $t \in (0,1]^m$, the function $D_\phi(\cdot \,|\, t)$ is strongly convex on $[0,1]^m$ with strong convexity constant $\min_{i \in \{1,\dots,m\}} \kappa_\phi(t_i)/t_i$. □

Proof of Lemma 5.2.14. Let x, y in $[0,\infty)$ such that $x \neq y$. Notice that any $w \leq 1/\max\{x,y\}$ is such that x and y are in $[0,1/w]$. Fix one such w and let $\kappa_\phi(w)$ be the strong convexity constant of the restriction of ϕ to $[0,1/w]$. Then, for any $\alpha \in (0,1)$,

$$\begin{aligned} \phi(\alpha x + (1-\alpha)y) &\leq \alpha\phi(x) + (1-\alpha)\phi(y) - \frac{1}{2}\kappa_\phi(w)\alpha(1-\alpha)|x-y|^2 \\ &< \alpha\phi(x) + (1-\alpha)\phi(y), \end{aligned}$$

since $\frac{1}{2}\kappa_\phi(w)\alpha(1-\alpha)|x-y|^2 > 0$. As x, y are arbitrary in $[0,\infty)$, this establishes the strict convexity of ϕ. □

5.B Proofs of the results of Section 5.3.2

Proof of Proposition 5.3.5. The proof is an adaption of the proof of Theorem 3.3 (iii) in Gietl and Reffel (2017). Let (t_n) be a sequence in $[0,1]^m$ such that $\lim_{n\to\infty} t_n = t_0$. For n large enough, $\theta_n = \vartheta^*(t_n)$ is well-defined via (5.4) and $s_n = \mathcal{S}(\theta_n)$ is the ϕ-projection of t_n on $\mathcal{M} = \mathcal{S}(\bar{\Theta})$. Since $\bar{\Theta}$ is a compact subset of $\mathbb{R}^k$, by the Bolzano–Weierstrass theorem, the sequence (θ_n) in $\bar{\Theta}$ has a convergent subsequence which converges to an element of $\bar{\Theta}$. Let (θ_{η_n}) be such a subsequence and let $\lim_{n\to\infty} \theta_{\eta_n} = \theta^{**}$. By continuity of $\mathcal{S}$, we immediately obtain that $\lim_{n\to\infty} \mathcal{S}(\theta_{\eta_n}) = \mathcal{S}(\theta^{**})$. Then, using the fact that every subsequence of (t_n) converges to t_0 and the lower semicontinuity of D_ϕ (see Proposition 5.2.4 (i)), we obtain that, for any $\theta \in \bar{\Theta}$,

$$\begin{aligned} D_\phi(\mathcal{S}(\theta^{**}) \,|\, t_0) &= D_\phi(\lim_{n\to\infty} \mathcal{S}(\theta_{\eta_n}) \,|\, \lim_{n\to\infty} t_n) = D_\phi(\lim_{n\to\infty} \mathcal{S}(\theta_{\eta_n}) \,|\, \lim_{n\to\infty} t_{\eta_n}) \\ &\leq \liminf_{n\to\infty} D_\phi(\mathcal{S}(\theta_{\eta_n}) \,|\, t_{\eta_n}) \leq \liminf_{n\to\infty} D_\phi(\mathcal{S}(\theta) \,|\, t_{\eta_n}), \end{aligned}$$

where we have used the fact that, for n large enough, $D_\phi(\mathcal{S}(\theta_{\eta_n})\,|\,t_{\eta_n}) \leq D_\phi(\mathcal{S}(\theta)\,|\,t_{\eta_n})$ since $\theta_{\eta_n} = \vartheta^*(t_{\eta_n})$. Finally, using the continuity of D_ϕ with respect to its second argument (see Proposition 5.2.4 (ii)), we obtain

$$\liminf_{n\to\infty} D_\phi(\mathcal{S}(\theta)\,|\,t_{\eta_n}) = \lim_{n\to\infty} D_\phi(\mathcal{S}(\theta)\,|\,t_{\eta_n}) = D_\phi(\mathcal{S}(\theta)\,|\,t_0).$$

In other words, we have shown that, for any $\theta \in \bar{\Theta}$, $D_\phi(\mathcal{S}(\theta^{**})\,|\,t_0) \leq D_\phi(\mathcal{S}(\theta)\,|\,t_0)$, which, since $\mathcal{S}(\vartheta^*(t_0))$ is the unique ϕ-projection of t_0 on $\mathcal{M} = \mathcal{S}(\bar{\Theta})$ by Condition 5.3.3, implies that $\mathcal{S}(\theta^{**}) = \mathcal{S}(\vartheta^*(t_0))$, and, by the injectivity of $\mathcal{S}$, that $\theta^{**} = \vartheta^*(t_0)$. Hence, $\lim_{n\to\infty} \theta_{\eta_n} = \lim_{n\to\infty} \vartheta^*(t_{\eta_n}) = \vartheta^*(t_0)$. Reasoning by contraposition, this implies that ϑ^* is continuous at t_0. □

Proof of Lemma 5.3.9. Let us first state a formula that we will need later in the proof. Having in mind the definitions given above Lemma 5.3.9, let $g : U \to \mathbb{R}^{b\times c}$ and $h : U \to \mathbb{R}^{c\times d}$, where U is an open subset of $\mathbb{R}^a$, be continuously differentiable on U. Then, the product function gh given by $(gh)(x) = g(x)h(x)$, $x \in U$, is well-defined and, using for instance Theorem T4.3 in Table IV of Brewer (1978), we have

$$J[gh](x) = \left(h(x)^\top \otimes I_b\right) J[g](x) + (I_d \otimes g(x))\, J[h](x), \qquad x \in U. \tag{5.16}$$

We shall now proceed with the proof. From (5.1) and (5.2), for any $t \in (0,1)^m$ and any $s \in [0,1]^m$,

$$D_\phi(s\,|\,t) = \sum_{i=1}^m t_i \phi\left(\frac{s_i}{t_i}\right),$$

implying that, under Condition 5.2.12, for any $t \in (0,1)^m$, $D_\phi(\cdot\,|\,t)$ is twice continuously differentiable on $(0,1)^m$. Then, since $\mathcal{S}$ is assumed to be twice continuously differentiable on $\mathring{\Theta}$ and $\mathcal{S}(\mathring{\Theta}) \subset (0,1)^m$ by Condition 5.3.1, for any $t \in (0,1)^m$, $D_\phi(\mathcal{S}(\cdot)\,|\,t)$ is twice continuously differentiable on $\mathring{\Theta}$. Next, fix $t \in (0,1)^m$ and consider the function from $(0,1)^m$ to $\mathbb{R}^m$ defined by

$$\Phi_t'(s) = J[D_\phi(\cdot\,|\,t)](s) = \left(\phi'\left(\frac{s_1}{t_1}\right), \ldots, \phi'\left(\frac{s_m}{t_m}\right)\right), \qquad s \in (0,1)^m.$$

Then, by the chain rule, for any $\theta \in \mathring{\Theta}$,

$$J[D_\phi(\mathcal{S}(\cdot)\,|\,t)](\theta) = J[D_\phi(\cdot\,|\,t)](\mathcal{S}(\theta))\, J[\mathcal{S}](\theta) = \Phi_t'(\mathcal{S}(\theta))^\top J[\mathcal{S}](\theta).$$

Note in passing that

$$J[\Phi_t'](s) = J_2[D_\phi(\cdot\,|\,t)](s) = \operatorname{diag}\left(\frac{1}{t_1}\phi''\left(\frac{s_1}{t_1}\right), \ldots, \frac{1}{t_m}\phi''\left(\frac{s_m}{t_m}\right)\right), \qquad s \in (0,1)^m.$$

Finally, from (5.16) and the chain rule, for any $\theta \in \mathring{\Theta}$,

$$\begin{aligned}
J_2[D_\phi(\mathcal{S}(\cdot)\,|\,t)](\theta) &= J\Big[J[D_\phi(\mathcal{S}(\cdot)\,|\,t)]\Big](\theta) = J\Big[\Phi_t'(\mathcal{S}(\cdot))^\top J[\mathcal{S}]\Big](\theta) \\
&= \Big(J[\mathcal{S}](\theta)^\top \otimes I_1\Big)\, J\left[\Phi_t'(\mathcal{S}(\cdot))\right](\theta) + \Big(I_k \otimes \Phi_t'(\mathcal{S}(\theta))^\top\Big)\, J_2[\mathcal{S}](\theta) \\
&= J[\mathcal{S}](\theta)^\top J\left[\Phi_t'\right](\mathcal{S}(\theta)) J[\mathcal{S}](\theta) + \Big(I_k \otimes \Phi_t'(\mathcal{S}(\theta))^\top\Big)\, J_2[\mathcal{S}](\theta) \\
&= J[\mathcal{S}](\theta)^\top \operatorname{diag}\left(\frac{1}{t_1}\phi''\left(\frac{\mathcal{S}_1(\theta)}{t_1}\right), \ldots, \frac{1}{t_m}\phi''\left(\frac{\mathcal{S}_m(\theta)}{t_m}\right)\right) J[\mathcal{S}](\theta) \\
&\quad + \left(I_k \otimes \left(\phi'\left(\frac{\mathcal{S}_1(\theta)}{t_1}\right), \ldots, \phi'\left(\frac{\mathcal{S}_m(\theta)}{t_m}\right)\right)^\top\right) J_2[\mathcal{S}](\theta).
\end{aligned}$$

□

Proof of Theorem 5.3.1. As verified in the proof of Lemma 5.3.9, under the conditions considered from Section 5.3 onwards, we have that, for any $t \in (0,1)^m$, the function $D_\phi(\mathcal{S}(\cdot)\,|\,t)$ is twice continuously differentiable on $\mathring{\Theta}$. Next, let F be the function from $(0,1)^m \times \mathring{\Theta}$ to $\mathbb{R}^k$ defined by $F(t,\theta) = J[D_\phi(\mathcal{S}(\cdot)\,|\,t)](\theta)$, $t \in (0,1)^m$, $\theta \in \mathring{\Theta}$. From the chain rule, we have that, for any $t \in (0,1)^m$ and $\theta \in \mathring{\Theta}$,

$$F(t,\theta) = J[D_\phi(\cdot\,|\,t)](\mathcal{S}(\theta)) J[\mathcal{S}](\theta) = \left(\phi'\left(\frac{\mathcal{S}_1(\theta)}{t_1}\right), \ldots, \phi'\left(\frac{\mathcal{S}_m(\theta)}{t_m}\right)\right)^\top J[\mathcal{S}](\theta), \tag{5.17}$$

implying that F is continuously differentiable on the open subset $(0,1)^m \times \mathring{\Theta}$ of $\mathbb{R}^{m+k}$.

From Condition 5.3.4, for any $t \in \mathcal{N}(t_0)$, the function $D_\phi(\mathcal{S}(\cdot)\,|\,t)$ from $\mathring{\Theta}$ to $\mathbb{R}$ reaches its unique minimum at $\vartheta^*(t) \in \mathring{\Theta}$, where the function ϑ^* is defined in (5.5). First order necessary optimality conditions then imply that

$$J[D_\phi(\mathcal{S}(\cdot)\,|\,t)](\vartheta^*(t)) = F(t,\vartheta^*(t)) = 0_{\mathbb{R}^k} \quad \text{for all } t \in \mathcal{N}(t_0). \tag{5.18}$$

Since $F(t_0,\vartheta^*(t_0)) = 0_{\mathbb{R}^k}$ by (5.18) and

$$J_2[D_\phi(\mathcal{S}(\cdot)\,|\,t_0)](\vartheta^*(t_0)) = J[F(t_0,\cdot)](\vartheta^*(t_0))$$

is invertible by Condition 5.3.7, we can apply the implicit function theorem (see, e.g., Fitzpatrick, 2009, Theorem 17.6, p 450) to obtain that there exists $r \in (0,\infty)$, an open ball $\mathcal{B} \subset (0,1)^m$ of radius r centered at t_0 and a continuously differentiable function $g : \mathcal{B} \to \mathbb{R}^k$ such that

$$F(t,g(t)) = 0_{\mathbb{R}^k} \qquad \text{for all } t \in \mathcal{B},$$

and, whenever $\|t-t_0\|<r$, $\|\theta-\vartheta^*(t_0)\|<r$ and $F(t,\theta)=0$, then $\theta=g(t)$. Moreover,

$$J[F(\cdot,g(t)](t)+J[F(t,\cdot)](g(t))J[g](t)=0_{\mathbb{R}^k} \qquad \text{for all } t\in\mathcal{B}.$$

Now, under Condition 5.3.4, Proposition 5.3.5 states that the function ϑ^* in (5.5) is continuous at t_0. Hence, there exists an open neighborhood $\mathcal{B}'\subset\mathcal{B}$ of t_0 such that, for any $t\in\mathcal{B}'$, $\|\vartheta^*(t)-\vartheta^*(t_0)\|<r$. Besides, by (5.18), $F(t,\vartheta^*(t))=0_{\mathbb{R}^k}$ for all $t\in\mathcal{B}'\cap\mathcal{N}(t_0)$. Therefore, by the implicit function theorem, for any $t\in\mathcal{B}'\cap\mathcal{N}(t_0)$, $\vartheta^*(t)=g(t)$ and thus ϑ^* is continuously differentiable at t_0 with Jacobian matrix at t_0 given by

$$J[\vartheta^*](t_0)=-J[F(t_0,\cdot)](\vartheta^*(t_0))^{-1}J[F(\cdot,\vartheta^*(t_0)](t_0).$$

Consider the continuously differentiable function from $(0,1)^m$ to $\mathbb{R}^m$ defined by

$$\Psi'(t)=\left(\phi'\left(\frac{\mathcal{S}_1(\vartheta^*(t_0))}{t_1}\right),\dots,\phi'\left(\frac{\mathcal{S}_m(\vartheta^*(t_0))}{t_m}\right)\right), \qquad t\in(0,1)^m,$$

and note that $J[\Psi'](t_0)$ is equal to $-\Delta(t_0)$, where $\Delta(t_0)$ is given in (5.9). The expression in (5.8) finally follows from the fact that $J[F(t_0,\cdot)](\vartheta^*(t_0))=J_2[D_\phi(\mathcal{S}(\cdot)\,|\,t_0)](\vartheta^*(t_0))$ and the fact that, by (5.16) and (5.17),

$$J[F(\cdot,\vartheta^*(t_0)](t_0)=J\Big[\Psi'(\cdot)^\top J[\mathcal{S}](\vartheta^*(t_0))\Big](t_0)=\Big(J[\mathcal{S}](\vartheta^*(t_0))^\top\otimes I_1\Big)J[\Psi'](t_0).$$

□

Proof of Lemma 5.3.15. From Conditions 5.3.3 and 5.3.14, there exists an open neighborhood $\mathcal{N}(t_0)\subset(0,1)^m$ of t_0 such that, for any $t\in\mathcal{N}(t_0)$, the ϕ-projection s^* of t on $\mathcal{M}$ exists, is unique and belongs to $(0,1)^m$. Furthermore, Condition 5.3.1 implies that the function $\mathcal{S}$ is a bijection between $\bar{\Theta}$ and $\mathcal{M}$. Under Condition 5.3.13, we immediately obtain that it is also a bijection between $\mathring{\Theta}$ and $\mathcal{M}\cap(0,1)^m$. Hence, for any $t\in\mathcal{N}(t_0)$, the unique ϕ-projection $s^*=\mathcal{S}^*(t)\in\mathcal{M}\cap(0,1)^m$ satisfies $s^*=\mathcal{S}(\theta^*)$ for some (unique) $\theta^*\in\mathring{\Theta}$. □

5.C Proofs of the results of Section 5.3.3

Proof of Proposition 5.3.16. First recall that, by continuity of $\mathcal{S}$ and compactness of $\bar{\Theta}$, $\mathcal{M}=\mathcal{S}(\bar{\Theta})$ is a compact subset of $[0,1]^m$. Note also that, for any $t\in(0,1)^m$, $[0,1]^m_{\mathrm{supp}(t)}=[0,1]^m$. Then, by Proposition 5.2.7 and, since Condition 5.2.2 holds, by Proposition 5.2.9 (i), for any $t\in(0,1)^m$, the ϕ-projection s^* of t on $\mathcal{M}=\mathcal{S}(\bar{\Theta})$ exists and is unique (which implies that Condition 5.3.3 holds). □

Proof of Proposition 5.3.19. Fix $t \in (0,1)^m$. As already verified in the proof of Lemma 5.3.9, the function $D_\phi(\mathcal{S}(\cdot)\,|\,t)$ is twice continuously differentiable on $\mathring{\Theta}$. From Theorem 2.1.11 in Nesterov (2004), we have that the strong convexity of the function $D_\phi(\mathcal{S}(\cdot)\,|\,t)$ from $\bar{\Theta}$ to $\mathbb{R}$ is equivalent to the fact that, for any $\theta \in \mathring{\Theta}$, $J_2[D_\phi(\mathcal{S}(\cdot)\,|\,t)](\theta) - \kappa(t) I_k$ is positive semi-definite, where $\kappa(t)$ is the strong convexity constant of the function $D_\phi(\mathcal{S}(\cdot)\,|\,t)$. This then immediately implies that, for any $t \in (0,1)^m$ and $\theta \in \mathring{\Theta}$, $J_2[D_\phi(\mathcal{S}(\cdot)\,|\,t)](\theta)$ is positive definite. Since, under Condition 5.3.4, ϑ^* can be expressed as in (5.5), we immediately obtain $J_2[D_\phi(\mathcal{S}(\cdot)\,|\,t_0)](\vartheta^*(t_0))$ is positive definite, that is, that Condition 5.3.7 holds. □

Proof of Lemma 5.3.21. From Condition 5.3.20, for any $\theta, \theta' \in \bar{\Theta}$ and $\alpha \in [0,1]$,

$$\mathcal{S}(\alpha\theta + (1-\alpha)\theta') = A(\alpha\theta + (1-\alpha)\theta') + \gamma = \alpha\mathcal{S}(\theta) + (1-\alpha)\mathcal{S}(\theta').$$

Next, since Condition 5.2.11 holds, from Proposition 5.2.13, for any $\theta, \theta' \in \bar{\Theta}$, $t \in (0,1)^m$ and $\alpha \in [0,1]$,

$$\begin{aligned}
D_\phi(\mathcal{S}(\alpha\theta + (1-\alpha)\theta')\,|\,t) &= D_\phi(\alpha\mathcal{S}(\theta) + (1-\alpha)\mathcal{S}(\theta')\,|\,t) \\
&\leq \alpha D_\phi(\mathcal{S}(\theta)\,|\,t) + (1-\alpha) D_\phi(\mathcal{S}(\theta')\,|\,t) - \frac{\kappa_{D_\phi(\cdot|t)}}{2}\alpha(1-\alpha)\|\mathcal{S}(\theta) - \mathcal{S}(\theta')\|_2^2 \\
&= \alpha D_\phi(\mathcal{S}(\theta)\,|\,t) + (1-\alpha) D_\phi(\mathcal{S}(\theta')\,|\,t) - \frac{\kappa_{D_\phi(\cdot|t)}}{2}\alpha(1-\alpha)\|A(\theta - \theta')\|_2^2 \\
&\leq \alpha D_\phi(\mathcal{S}(\theta)\,|\,t) + (1-\alpha) D_\phi(\mathcal{S}(\theta')\,|\,t) - \frac{\kappa_{D_\phi(\cdot|t)}}{2}\alpha(1-\alpha)c\|\theta - \theta'\|_2^2,
\end{aligned}$$

where $\kappa_{D_\phi(\cdot|t)}$ is the strong convexity constant of the function $D_\phi(\cdot\,|\,t)$ and c is a constant that depends on A such that, for any $x \in \mathbb{R}^k$, $\|Ax\|_2^2 \geq c\|x\|_2^2$. To prove that, for any $t \in (0,1)^m$, the function $D_\phi(\mathcal{S}(\cdot)\,|\,t)$ from Θ to $\mathbb{R}$ is strongly convex, it remains to show that the constant c can be chosen strictly positive. Let $S^{k-1} = \{x \in \mathbb{R}^k : \|x\|_2 = 1\}$ be the unit sphere in $\mathbb{R}^k$ and let $c = \min_{x \in S^{k-1}} \|Ax\|_2^2$. Then $c \geq 0$ and, as expected, for any $x \in \mathbb{R}^k$,

$$\|Ax\|_2^2 = \|x\|_2^2 \left\| A\frac{x}{\|x\|_2} \right\|_2^2 \geq c\|x\|_2^2.$$

To prove that $c > 0$, we proceed by contradiction. Assume that $c = 0$. Then, there would exist $x \in S^{k-1}$ such that $Ax = 0$. Since $\mathcal{S}$ is assumed to be affine and injective, $x \mapsto Ax$ is a linear injective function and $Ax = 0$ implies that $x = 0_{\mathbb{R}^k}$. This is a contradiction since $0_{\mathbb{R}^k} \notin S^{k-1}$. Hence, for any $t \in (0,1)^m$, the function $D_\phi(\mathcal{S}(\cdot)\,|\,t)$ from Θ to $\mathbb{R}$ is strongly convex. □

Proof of Corollary 5.3.22. To show the result, it suffices to verify that the conditions of Theorem 5.3.1 and Corollary 5.3.10, namely Conditions 5.3.4 and 5.3.7, are satisfied. From Proposition 5.3.16, we know that Condition 5.3.3 holds. Lemma 5.3.15 then implies that Condition 5.3.4 holds since Conditions 5.3.13 and 5.3.14 are assumed to be satisfied. From Lemma 5.3.21, we have that Conditions 5.2.11 and 5.3.20 imply Condition 5.3.18, which implies Condition 5.3.7 by Proposition 5.3.19. □

Proof of Proposition 5.3.23. Recall that the fact that, for any probability vector $q \in (0,1)^m$, the ϕ-projection p^* of q on $\mathcal{M} = \mathcal{S}(\bar{\Theta})$ exists and is unique follows from Proposition 5.3.16. Note also that the function $\mathcal{S}$ defined in Condition 5.3.1 is a bijection between $\bar{\Theta}$ and $\mathcal{M} = \mathcal{S}(\bar{\Theta})$. When it is affine (that is, under Condition 5.3.20), an expression of its inverse is given by Lemma 5.C.1 below. Specifically, we then have that

$$\mathcal{S}^{-1}(p) = (A^\top A)^{-1} A^\top (p - \gamma), \qquad p \in \mathcal{M}.$$

This implies that the set $\mathcal{M}$ can be equivalently expressed as

$$\mathcal{M} = \left\{ p \in [0,1]^m : \sum_{i=1}^m p_i = 1, (A^\top A)^{-1} A^\top (p - \gamma) \in \bar{\Theta} \right\},$$

and, since Θ is defined by linear inequalities, that $\mathcal{M}$ is defined by linear inequalities. Some thought then reveals that there exists a $d \in \mathbb{N}^+$, $\alpha_1, \ldots, \alpha_d \in \mathbb{R}$ and d functions $g_1, g_2, \ldots, g_d$ from $\{1, \ldots, m\}$ to $\mathbb{R}$ such that

$$\mathcal{M} = \left\{ p \in [0,1]^m : \sum_{i=1}^m p_i = 1, \sum_{i=1}^m p_i g_j(i) \geq \alpha_j, j \in \{1, \ldots, d\} \right\}.$$

Setting $f_j(i) = g_j(i) - \alpha_j$ for $i \in \{1, \ldots, m\}$ and $j \in \{1, \ldots, d\}$, it is easy to verify that $\mathcal{M}$ can be equivalently rewritten as

$$\mathcal{M} = \left\{ p \in [0,1]^m : \sum_{i=1}^m p_i = 1, \sum_{i=1}^m p_i f_j(i) \geq 0, j \in \{1, \ldots, d\} \right\}.$$

Theorem 2 (c) of Rüschendorf (1987) and the remark following it then imply that, for any probability vector $q \in (0,1)^m$, there exist $c = (c_0, \ldots, c_d) \in \mathbb{R}^{d+1}$ such that, for any $i \in \{1, \ldots, m\}$,

$$p_i^* = q_i \phi'^{-1} \left(c_0 + \sum_{j=1}^d c_j f_j(i) \right),$$

where $p^* = \mathcal{S}^*(q)$ is the ϕ-projection of q on $\mathcal{M}$. If $\lim_{x\to 0^+} \phi'(x) = -\infty$, $\phi'^{-1}(x) > 0$ for all $x \in (0,\infty)$ and the previous centered display immediately implies that $\mathcal{S}^*(q) \in (0,1)^m$ whenever $q \in (0,1)^m$. □

The following lemma is necessary for proving Proposition 5.3.23.

Lemma 5.C.1. *Under Condition 5.3.20, the inverse $\mathcal{S}^{-1}$ of $\mathcal{S} : \bar{\Theta} \to \mathcal{M}$ can be expressed as*

$$\mathcal{S}^{-1}(s) = (A^\top A)^{-1} A^\top (s - \gamma), \qquad s \in \mathcal{M}. \tag{5.19}$$

Proof. Since A is of full rank because of the injectivity of $\mathcal{S}$, we know, for instance from Ben-Israel and Greville (2003, Theorem 5, p 48), that $A^\top A$ is invertible. Moreover, for instance from expression (6.13) in Trefethen and Bau III (1997), we have that the orthogonal projection of $x \in \mathbb{R}^m$ on $\mathrm{Im}(A) = \{y \in \mathbb{R}^m : y = Az, z \in \mathbb{R}^k\}$ can be computed as $A(A^\top A)^{-1} A^\top x$. This implies that

$$A(A^\top A)^{-1} A^\top y = y, \qquad \text{for all } y \in \mathrm{Im}(A). \tag{5.20}$$

We shall use these facts to prove that (5.19) provides an expression of the inverse of the function $\mathcal{S} : \bar{\Theta} \to \mathcal{M}$. Let $\theta \in \bar{\Theta}$. From the invertibility of $A^\top A$, we have that

$$\mathcal{S}^{-1}(\mathcal{S}(\theta)) = \mathcal{S}^{-1}(A\theta + \gamma) = (A^\top A)^{-1} A^\top A\theta = \theta.$$

Next, let $s \in \mathcal{M}$. Then, using (5.20),

$$\mathcal{S}(\mathcal{S}^{-1}(s)) = \mathcal{S}((A^\top A)^{-1} A^\top (s - \gamma)) = A(A^\top A)^{-1} A^\top (s - \gamma) + \gamma = (s - \gamma) + \gamma = s,$$

since $s - \gamma$ belongs to $\mathrm{Im}(A)$. □

Proof of Corollary 5.3.25. To show the result, we shall check that the conditions of Corollary 5.3.22 are satisfied. It thus just remains to verify Condition 5.3.14. The latter follows from Proposition 5.3.23 since attention is restricted to probability vectors. □

Proof of Proposition 5.3.26. The fact that $\mathcal{M}$ is defined by linear constraints is equivalent to the existence of a $d \times m$ matrix B such that $\mathcal{M} = \{s \in [0,1]^m : Bs = \alpha\}$. Because $\mathcal{M}$ is convex and not reduced to a singleton, it contains an infinity of elements. Since $\mathcal{M} \cap (0,1)^m$ is nonempty, we can choose $s_0 \in \mathcal{M} \cap (0,1)^m$. Then, for any $s \in \mathcal{M}$, $s - s_0 \in \ker(B)$, where $\ker(B) = \{x \in \mathbb{R}^m : Bx = 0_{\mathbb{R}^d}\}$. Since $\ker(B) \neq \{0_{\mathbb{R}^m}\}$, there exists a strictly

positive integer $k<m$ and $\delta_1,\ldots,\delta_k\in\mathbb{R}^m$ that form a basis of $\ker(B)$. Let A be the matrix whose column vectors are $\delta_1,\ldots,\delta_k$ and let

$$\Theta=\left\{\theta\in\mathbb{R}^k : 0_{\mathbb{R}^m}\leq s_0+A\theta\leq 1_{\mathbb{R}^m}\right\}. \tag{5.21}$$

Since, for any $s\in\mathcal{M}$ (resp. $\mathcal{M}\cap(0,1)^m$), $s-s_0\in\ker(B)$, there exists a unique $\theta\in\mathbb{R}^k$ such that $s-s_0=A\theta$ and, as $s=s_0+A\theta$ belongs to $[0,1]^m$ (resp. $(0,1)^m$), θ necessarily belongs to $\bar{\Theta}$ (resp. $\mathring{\Theta}$). In other words,

$$\forall s\in\mathcal{M} \text{ (resp. } \mathcal{M}\cap(0,1)^m\text{), there exists a unique } \theta\in\bar{\Theta} \text{ (resp. } \mathring{\Theta}\text{)} \quad \text{s.t. } s=s_0+A\theta. \tag{5.22}$$

To complete the proof, we shall show that Conditions 5.3.1, 5.3.13 and 5.3.20 are satisfied for Θ in (5.21) and the function $\mathcal{S}$ defined by $\mathcal{S}(\theta)=A\theta+s_0$, $\theta\in\bar{\Theta}$. Condition 5.3.20 clearly holds since $\bar{\Theta}$ is convex and $\mathcal{S}$ has the right form with $\gamma=s_0$. Let us next verify that Condition 5.3.1 is satisfied. Note that Θ is clearly bounded. Furthermore, from (5.21), the fact that $s_0\in\mathcal{M}\cap(0,1)^m$ is equivalent to $0_{\mathbb{R}^m}<s_0+A0_{\mathbb{R}^k}<1_{\mathbb{R}^m}$ implies that $0_{\mathbb{R}^k}\in\mathring{\Theta}$ and thus that $\mathring{\Theta}\neq\emptyset$. The fact that $\mathcal{S}$ is a bijection between $\bar{\Theta}$ and $\mathcal{M}$ follows, on one hand, from the fact that, for any $\theta\in\bar{\Theta}$, $\mathcal{S}(\theta)\in\mathcal{M}$ (since $B(A\theta+s_0)=BA\theta+Bs_0=0_{\mathbb{R}^d}+\alpha$) and, on the other hand, from (5.22) (since, for any $s\in\mathcal{M}$, there exists a unique $\theta\in\bar{\Theta}$ such that $\mathcal{S}(\theta)=s$). In a similar way, Condition 5.3.13 is a consequence of the fact that, for any $\theta\in\mathring{\Theta}$, $\mathcal{S}(\theta)\in\mathcal{M}\cap(0,1)^m$ and (5.22). □

Chapter 6

First Passage Time Density Approximation

This chapter is based on Di Nardo et al. (2023)

Di Nardo, E., G. D'Onofrio, and T. Martini (2023), Approximating the first passage time density from data using generalised Laguerre polynomials, *Communications in Nonlinear Science and Numerical Simulation* 118, 106991,

and on Di Nardo et al. (2024)

Di Nardo, E., G. D'Onofrio, and T. Martini (2024), Orthogonal gamma-based expansion for the cir's first passage time distribution, *Applied Mathematics and Computation* 480, 128911.

In the context of stochastic modelling of hydrological processes, one may consider a stochastic process representing the concentration of a tracer or the presence of water particles in a hydrological system. Then, it could be of interest to investigate the first time at which a particle, or a set of particles, first reaches a certain threshold after being introduced into the system. This threshold could represent various hydrological events, such as the arrival of water at a river outlet, the appearance of a contaminant at a monitoring well, or the time when the water table reaches a critical level. For example, Stechmann and Neelin (2014) recalls that in models for precipitation and column vapour, precipitation events commence when water vapor attains a specific threshold value and conclude when it drops to a slightly lower threshold, as supported by recent observational and modelling studies. In

addition to the previously mentioned hydrological applications, this same problem also arises in a wide array of fields, including finance, engineering, computational neuroscience, mathematical biology, and reliability theory (see Redner (2001) for a thorough exposition). The dynamics of a noisy system in these contexts are often modelled by a stochastic process $\{Y(t)\}_{t\geq 0}$ evolving in the presence of a threshold $S(t)$ for $t \geq 0$. The mathematical study of the FPT problem consists in finding the probability density function (pdf) $g[S(t),t|y_\tau,\tau] = \frac{d}{dt}\mathbb{P}(T<t)$ of the random variable (rv) T, defined by

$$T = \begin{cases} \inf_{t\geq\tau}\{Y(t) > S(t)\}, & Y(\tau) = y_\tau < S(\tau), \\ \inf_{t\geq\tau}\{Y(t) < S(t)\}, & Y(\tau) = y_\tau > S(\tau), \end{cases}$$

representing the time it takes for the process $\{Y(t)\}_{t\geq 0}$ to cross the threshold $S(t)$ for the first time. As we continue, the pdf of interest $g[S(t),t|y_\tau,\tau]$ will be denoted by g for notational convenience. Despite the problem's classical nature and its seemingly straightforward formulation, a closed-form solution exists only in very limited cases, dependent on the characteristics of both $\{Y(t)\}_{t\geq 0}$ and $S(t)$. Various strategies have been developed to tackle this problem, as comprehensively reviewed in Ricciardi et al. (1999). For instance, Doob's representation formula (Doob, 1949) has been used, or, as another example, Siegert's equation (Siegert, 1951). It consists in a partial differential equation involving either the moments of T or its Laplace transform. Albeit rarely, closed-form expressions can emerge (see, e.g., Buonocore et al., 2015). Asymptotic expressions of g can be studied using the Volterra integral equations (Giorno et al., 1990; Nobile et al., 1985), Laplace transform techniques (Martin et al., 2019) or Large Deviation estimates (Baldi et al., 2020; D'Onofrio et al., 2018). However, since in general the latter techniques may not work, in practice numerical evaluations of g can come in handy. For instance, g can be approximated by exploiting the fact that it is the solution of a non-singular second kind Volterra integral equation (Buonocore et al., 1987; Ricciardi et al., 1984; Giorno et al., 1989; Gutiérrez-Jáimez et al., 1995; Di Nardo et al., 2001) or by using a Sturm-Liouville series expansion with (Linetsky, 2004; Alili et al., 2005; Kent, 1980). It is clear that all these techniques rely on the knowledge of the nature of the stochastic process $\{Y(t)\}_{t\geq 0}$. More precisely, they exploit the availability of some computable objects regarding $\{Y(t)\}_{t\geq 0}$. In this context, as already mentioned in the introduction, by starting from the FPT problem with constant boundary $S(t) \equiv S > 0$ for the CIR process, which is notoriously tricky, the authors of Di Nardo and D'Onofrio (2021) introduced a series

expansion of g as a tool to approximate it. In the particular case of the CIR, the feasibillity of the expansion was based on novel way to compute the cumulants of T. Then, more precisely, the FPT pdf has been expanded in series of generalised Laguerre polynomials, involving moments computed from cumulants and weighted by a gamma pdf. The idea of approximating a pdf by truncating a suitable series expansion is not new. Indeed, such an expansion when Hermite polynomials are used instead of the Laguerre, and a normal pdf instead of the gamma, is known as Gram-Charlier type-A series. We refer the reader to Provost and Ha (2016) and Asmussen et al. (2019) for two interesting papers where the more general methodology of approximating a pdf based on the knowledge of its moments is introduced, using the product of a weight function, usually called parent or reference distribution, and a family of associated orthogonal polynomials. Nevertheless, Section 6.1 of this chapter will provide the necessary preliminaries on this subject.

Before proceeding, we would like to discuss an issue, that, as far as our knowledge goes, is seldom treated in the literature. That is, this type of approximation is not guaranteed to be always positive. Two main approaches can be found scattered in the literature. The first one involves finding constrained regions on the values of the moments (or cumulants) such that the resulting approximation will provide a non-negative pdf, and, due to the complexity of this approach, it is applied when the approximation is built with a low truncation order. For instance, this technique was used when approximating the FPT pdf of an Ornstein-Uhlenbeck process in Smith (1991). Similarly, some restrictions on the first four moments that guarantee the non-negativity of the approximated density are discussed in Lung (1998). More recently, in the case of a normal reference density and for arbitrary even order, the valid region of cumulants has been found numerically through a semi-definite algorithm by Lin and Zhang (2022). It must be noted that, with the low truncation order employed in these references, the approximated density may fail to be close to the desired one, especially for objective densities that are not sufficiently close to the reference density. A second way of tackling this issue consists in replacing values of a suitable positive function to the negative ones assumed by the approximated pdf. For instance, Wilson and Wragg (1973) suggests a second-degree polynomial as an interpolating function for pdfs with support $(0, \infty)$, when the negativity happens in a right-handed neighborhood of the origin. In this thesis, this second approach is developed along a different direction and for the first time within the FPT framework, as shall be seen in Section

6.4.5. We mention that, as will become clearer in Section 6.4.4, this issue can also be tackled by producing a correct approximation of the FPT cumulative distribution function (cdf).

The issues highlighted thus far are not the sole concerns related to the application of this approximation. Additional factors, such as the selection of gamma pdf parameters and the series truncation order, may also influence the approximation's accuracy. These topics, which are only briefly mentioned in Provost and Ha (2016), are explored comprehensively in this chapter. For instance, the critical importance of the coefficient of variation in determining the gamma pdf parameters is discussed, and the truncation order is managed through suitable stopping criteria.

The efficiency of the proposed approximation will be tested on two diffusion processes: the geometric brownian motion and the First passage time, also known as Cox-Ingersoll-Ross (CIR). While the former has been selected for its mathematical tractability, which leads to the possibility of computing a closed form of the FPT pdf with almost all the techniques mentioned in the beginning, the latter was considered precisely because its FPT problem is historically difficult to address and in such a context our tool finds its most useful application.

Having a reliable approximation of the FPT pdf which can be conveniently corrected to a bonafide pdf opens up the path to three interesting applications.

The first one is the proposal of an acceptance-rejection-like method that hinges upon the particular form of the series expansion. This approach enables the generation of FPT data even when its distribution is unknown. Although acceptance-rejection methods have previously been employed in the FPT context, as seen in Herrmann and Zucca (2019) and Mijatović et al. (2015), they have not been applied to the CIR process. The novelty of our proposed approximation strategy lies in its innovative exploitation of a series representation of the unknown density. This method is particularly advantageous given the lack of exact simulation techniques for CIR sample paths. Existing methods, which rely on discretization or transition densities, incur substantial computational costs when the fixed time step is small.

The second one exploits an additional benefit of the Laguerre-Gamma series representation. Indeed, it enables the generation of density estimates derived from sample moments. In situations where the FPT moments/cumulants are unknown, this method can be used to construct an estimator of the FPT pdf from a sample of FPT data. This technique is

known in the literature. More in general, when dealing with a series expansion with respect to an arbitrary family of orthogonal polynomials, these estimators are referred to in the literature as orthogonal series estimators. The Laguerre-Gamma series itself can be highly competitive compared to traditional density estimators like the kernel density estimator (KDE) or the histogram when the distribution generating the data is supported on $(0, \infty)$, as noted by Hall (1980).

Finally, in the third application, we will set forth an example of how the Laguerre-Gamma series expansion could be used to perform parameter estimation of the underlying diffusion through an approximated maximum likelihood, in the case where a sample of FPT data is available, and, unlike in the previous scenario, the moments are known.

This chapter is organized as follows.

In Section 6.1 we describe in general a procedure for obtaining approximations of unknown pdfs by truncating series representations with respect to a family of orthogonal polynomials and a known reference density. The choice of the latter is then treated in the context of the FPT problem.

Next, in Section 6.2, after briefly recalling the framework in which we shall give our proposals, we provide the theoretical results fundamental for applying the procedure outlined in the previous section to a FPT density g and FPT cdf G, such as a study of the order of convergence.

Given the framework set up in Section 6.1 and 6.2, in Section 6.3 we define the approximant function, we delineate some of its properties and we introduce the issues discussed above.

Section 6.4 is then written with the aim of providing ways to solve these issues. More in general, it provides elucidations on how to actually compute the given approximant. For instance, corrections for positivity, an iterative algorithm and related stopping criteria are introduced. The method for approximating the cdf is developed as well. We stress that this approach, new in the FPT context, has some numerical advantages and allows an easier approximation of quantiles.

In Section 6.5 we recall the FPT problem for the GBM and the CIR process resuming the known results useful for carrying out the proposed approximation.

Numerical examples are given in Section 6.6 for both the GBM and the CIR process aiming to discuss the strengths and weaknesses of the proposed

approach. We set three different choices of the CIR process (resp. GBM) parameters and boundaries that corresponds to different forms and statistical properties of the FPT pdf and cdf.

The last section presents the three applications introduced above. Concluding remarks then end the chapter.

6.1 Moment Based Probability Density Function Expansions

Let X be a real valued absolutely continuous rv and let g be its density. Suppose g is not available in a closed form, but the moments μ_j (or cumulants κ_j) of X are known. We denote the support of g by $\text{supp}(g) = \{x \in \mathbb{R} : g(x) > 0\}$. Furthermore, suppose that additional information on g, such as knowledge of the support or of some dispersion measures, is available such that it is possible to identify a well known "reference" distribution f resembling g. It is clear that we must have $\text{supp}(f) \subset \text{supp}(g)$. For example, knowing that g is supported on the real line and that g is symmetric would suggest us to select f as the normal pdf. Let ν be a measure having density f and consider the weighted Hilbert space L^2_ν defined as

$$L^2_\nu = \left\{ h : \mathbb{R} \to \mathbb{R} \,|\, \left(\int h^2 \, d\nu \right)^{1/2} = \left(\int_{-\infty}^{\infty} h^2(x) f(x) \, dx \right)^{1/2} < \infty \right\}$$

and equipped with the usual inner product $< h, l >_\nu = \int hl \, d\nu$ and the corresponding norm $||h||^2_\nu = < h, h >_\nu$ for $h, l \in L^2_\nu$.

To have a tractable form of g, as it has been proposed in a variety of contexts in the literature, one could exploit a series expansion of g based on the moments (or cumulants) of X and on a set of polynomials $\{P_k\}_{k \geq 0}$, with $P_0 = 1$, which is orthonormal with respect to some measure having density f. That is

$$< P_n, P_m >_\nu = \delta_{mn}, \quad m, n \in \mathbb{N}.$$

If for all $j \in \mathbb{N}$ the moments $\int_{\text{supp}(f)} x^j f(x) \mathrm{d}x$ of f are finite, classical results (see, e.g., Szegő, 1975, Section 2.2) shows the existence of the desired family $\{P_k\}_{k \geq 0}$, which can then be computed with the classical Gram-Smidth orthogonalization procedure. That is why, in the rest of this chapter, when discussing a reference density whose moments are known to be finite we will often refer to its associated family of orthonormal polynomials. The aim is

being able to meaningfully write

$$\frac{g}{f} = \sum_{k=0}^{\infty} a_k(\mu_1, \ldots, \mu_k) P_k, \tag{6.1}$$

where $a_k(\mu_1, \ldots, \mu_k)$ are coefficients depending on the first k moments. Series representations of this type are also known as generalised Fourier series. Let $m \in \mathbb{N}$. Equation (6.1) remains a formal writing until a definition of the convergence as m goes to infinity of the partial sums $\sum_{j=1}^{m} a_j P_j$ to $\frac{g}{f}$ is precisely given. Usually, as shall become clearer in the forthcoming discussion, conditions on g, f and $\{P_k\}_{k\geq 0}$ are found such that the convergence in L^2_ν is guaranteed. Indeed, a common assumption is that the orthonormal sequence $\{P_k\}_{k\geq 0}$ is complete in L^2_ν, that is (see, e.g., Sansone, 1991).

Definition 6.1.1. *The orthonormal sequence $\{P_j\}_{j\geq 0}$ is said complete in L^2_ν if any of the following equivalent conditions hold:*

1. *for all $k \geq 0$, if $g \in L^2_\nu$ and $< g, P_k >_\nu = 0$ then $g = 0$;*
2. *$h = \sum_{j\geq 0} < f, P_k >_\nu P_k$, $\forall h \in L^2_\nu$ (Series Representation);*
3. *$||h||^2 = \sum_{k\geq 0} | < h, P_k >_\nu |^2$, $\forall h \in L^2_\nu$ (Parseval Identity).*

In other words, the sequence $\{P_j\}_{j\geq 0}$ is complete if it forms an orthonormal basis of L^2_ν. Thus, if $\{P_k\}_{k\geq 0}$ is proven to be complete and assuming furthermore that $\frac{g}{f} \in L^2_\nu$, from 2. in Definition 6.1.1, we have that

$$\begin{aligned}
\frac{g}{f} &= \sum_{k\geq 0} < f, P_k >_\nu P_k = \sum_{k\geq 0} < f, P_k >_\nu P_k = \sum_{k\geq 0} \int_{-\infty}^{\infty} \frac{g(x)}{f(x)} P_k(x) f(x)\, dx\, P_k \\
&= \sum_{k\geq 0} \mathbb{E}[P_k(X)] P_k = 1 + \sum_{k\geq 1} \mathbb{E}[P_k(X)] P_k = 1 + \sum_{k\geq 1} \left(\sum_{i=0}^{k} b_{i,k} \mathbb{E}[X^i] \right) P_k \\
&= 1 + \sum_{k\geq 1} \left(\sum_{i=0}^{k} b_{i,k} \mu_i \right) P_k
\end{aligned}$$

where $b_{i,k}$ for $i \in \{1, \ldots, k\}$ are the coefficients of P_k. Thus, by setting $a_k(\mu_1, \ldots, \mu_k) = \mathbb{E}[P_k(X)]$ we recover the desired expression (6.1). Note that the latter, without additional assumptions, holds only as a serie representation in L^2_ν in the following sense: we have $||\sum_{k=0}^{n} \mathbb{E}[P_k(X)] P_k - \frac{g}{f}||^2$ goes to zero as n tends to infinity. From now on, we will assume that $\{P_k\}_{k\geq 0}$ is complete and that $\frac{g}{f} \in L^2_\nu$ and, for simplicity of notation, we will write $a_k = \mathbb{E}[P_k(X)]$. Before proceeding, we stress out that pointwise and uniform

convergence of $\sum_{k=0}^{n} a_k P_k$ to $\frac{g}{f}$ are usually harder to prove and must be investigated case by case, especially in the case of an unbounded $\mathrm{supp}(g)$, such as $\mathrm{supp}(g) = (0,\infty)$ in our forthcoming application to the FPT problem.

After having obtained a series expansion of $\frac{g}{f}$, an approximation of g is readily computed by truncating the series to an order n. That is, we define our approximant $\hat{g}_n$ of g as

$$\hat{g}_n(x) = f(x)(1 + \sum_{k=1}^{n} a_k P_k(x)), \quad x \in \mathrm{supp}(g), \tag{6.2}$$

recalling that

$$a_k = \mathbb{E}[P_k(X)] = \left(\sum_{i=0}^{k} b_{i,k} \mu_i \right). \tag{6.3}$$

Throughout this chapter, n in $\hat{g}_n$ shall be referred to as the *truncation order* or *approximation order*. Usually, it is possible to rearrange the above expression by putting together elements multiplying the same power of x to obtain

$$\hat{g}_n(x) = f(x)\,(1 + p_n(x)), \tag{6.4}$$

where

$$p_n = \sum_{k=1}^{n} A_k x^k, \tag{6.5}$$

for some other coefficients A_k to be determined. In the following proposition we collect some well known and useful properties of the approximant $\hat{g}_n$. Recall that $\mu_k = \mathbb{E}[X^k]$ for $k \geq 0$, where $\mu_0 = 1$.

Proposition 6.1.2. *Let any $n > 0$ and let X_f be the rv having density f. We have that*

1. $\int_{\mathrm{supp}(g)} \hat{g}_n(x)\mathrm{d}x = 1$;

2. $\int_{\mathrm{supp}(g)} x^j \hat{g}_n(x)\mathrm{d}x = \mu_j, \quad j \in \{0,\ldots,n\}$;

3. *let $j \in \{1,\ldots,n\}$. If $\mathbb{E}[X_f^k] = \mu_k$ for $k \in \{1,\ldots,j\}$, then $a_k = 0$ in (6.3) for $k \in \{1,\ldots,j\}$.*

Proof. All the three statements follow from the orthonormality of the sequence $\{P_k\}_{k\geq 0}$. To prove 1., simply observe that

$$\begin{aligned}\int_{\mathrm{supp}(g)} \hat{g}_n(x)\mathrm{d}x &= \int_{\mathrm{supp}(g)} f(x)\mathrm{d}x + \int_{\mathrm{supp}(g)} f(x)\sum_{k=1}^{n} a_k P_k(x)\mathrm{d}x \\ &= 1 + \sum_{k=1}^{n} a_k \int_{\mathrm{supp}(g)} f(x)P_k(x)\mathrm{d}x \\ &= 1 + \sum_{k=1}^{n} a_k < 1, P_k >_\nu \\ &= 1 + \sum_{k=1}^{n} a_k < P_0, P_k >_\nu = 1.\end{aligned}$$

Now let us consider point 2. We shall proceed by induction. For $j=0$ the the statement has been proven in 1. above. Suppose the statement holds for $j=n-1$ and let us then prove it for n. For the remainder of the proof, let $b_{i,k}$ for $i\in\{1,\ldots,k\}$ be the coefficients of P_k. We have that

$$\begin{aligned}\int_{\mathrm{supp}(g)} P_n(x)\hat{g}_n(x)\mathrm{d}x &= \int_{\mathrm{supp}(g)} P_n(x)f(x)\mathrm{d}x + \int_{\mathrm{supp}(g)} f(x)\sum_{k=1}^{n} a_k P_n(x)P_k(x)\mathrm{d}x \\ &= < 1, P_n >_\nu + \sum_{k=1}^{n} a_k < P_n, P_k >_\nu \\ &= < P_0, P_n >_\nu + a_n \\ &= a_n = \sum_{i=0}^{n} b_{i,n}\mu_i.\end{aligned}$$

Hence, since by the induction hypothesis

$$\begin{aligned}\int_{\mathrm{supp}(g)} P_n(x)\hat{g}_n(x)\mathrm{d}x &= \sum_{i=0}^{n} b_{i,n}\int_{\mathrm{supp}(g)} x^i \hat{g}_n(x) \\ &= \sum_{i=0}^{n-1} b_{i,n}\mu_i + b_{n,n}\int_{\mathrm{supp}(g)} x^n \hat{g}_n(x),\end{aligned}$$

the statement follows. Finally, note that under the assumption of statement 3. for $k\in\{1,\ldots,j\}$ we have

$$\begin{aligned}a_k = \mathbb{E}[P_k(X)] &= \sum_{i=1}^{k} b_{i,k}\mathbb{E}[X^i] \\ &= \sum_{i=1}^{k} b_i \mathbb{E}[X_f^i] = \mathbb{E}_\nu[P_k(X_f)] = < 1, P_k >_\nu \\ &= < P_0, P_k >_\nu = 0,\end{aligned}$$

where $\mathbb{E}_\nu$ denotes the expected value with respect to ν.

□

Unfortunately, $\hat{g}_n$ is not guaranteed to be a bonafide pdf since it could be negative. This problem, along with the choice of the order n of approximation, will be more thoroughly discussed in Sections 6.4.2 and 6.4.5, when applying the above density approximation to the FPT problem. The next subsection is devoted to the choice of the reference density f for our problem.

6.1.1 Choice of a Reference Density

To be able to employ the method described in Section 6.1, a first step to be taken is the choice of a reference density f. In this context we can rely on the fact that the FPT density g is supported on $(0,\infty)$. Some well known probability densities whose domain is $(0,\infty)$ are

1. the log-normal pdf $\tilde{f}_{\mu,\sigma}$ with parameters μ and σ;
2. the inverse Gaussian pdf $\bar{f}_{\mu,\lambda}$ with parameters μ and λ;
3. the gamma pdf $f_{\alpha,\beta}$ with parameters α and β.

To be able to obtain the desired approximation, it is crucial to investigate the associated sequence of orthonormal polynomials. For the log-normal pdf $\tilde{f}_{\mu,\sigma}$, the associated sequence $\{P_k\}_{k\geq 0}$ seems not to be classically known. It has been computed for $\mu=0$ and $\sigma=1$ in Ernst et al. (2012) and for arbitrary μ and σ in Asmussen et al. (2019) and Zheng et al. (2012). However, $\{P_k\}_{k\geq 0}$ is not a complete sequence (see, e.g., Asmussen et al., 2019, Proposition 1.1). The latter result is deeply connected to the fact that the log-normal distribution is not determined by its moments (Heyde, 1963). Therefore, $\hat{g}_n$ might converge to a density different from g, but sharing the same moments as g (this result is well known, see for example (Ernst et al., 2012, Proposition 4.1) for a non trivial example of a family of densities for which the convergence fails).

As will be explained in Section 6.5.2, it is known that the GBM FPT rv has an inverse Gaussian distribution. Therefore, as a proof of concept, one could choose the inverse Gaussian pdf $\bar{f}_{\mu,\lambda}$ as a reference density in the GBM case. However, Nishii (1996) shows that a standard method of differentiating the density $\bar{f}_{\mu,\lambda}$ does not lead to an orthogonal polynomial system $\{P_k\}_{k\geq 0}$. At the same time, when using the Gram–Schmidt orthogonalisation procedure, the resulting polynomials are not easy to work with (see also Goffard and

Laub, 2020). In Hassairi and Zarai (2004) the authors propose another procedure which hinges upon the so-called bi-orthogonality property, but they do not discuss whether this construction leads to a basis. Therefore, with the aim of providing a simple and practical method, we decided not to opt for the inverse Gaussian pdf $\bar{f}_{\mu,\lambda}$ as a reference density.

As we continue, we consider the gamma density $f_{\alpha,\beta}$ defined as

$$f_{\alpha,\beta}(t) = \beta(\beta t)^{\alpha} \frac{e^{-\beta t}}{\Gamma(\alpha+1)}, \quad t>0, \tag{6.6}$$

with scale parameter $\alpha+1>0$ and shape parameter $\beta>0$. The gamma pdf $f_{\alpha,\beta}$ for $\beta=1$ has already been suggested in Wilson and Wragg (1973) as reference density to approximate a pdf over $(0,\infty)$. All moments of the gamma distribution are finite and its moment generating function uniquely determines the distribution. In the case of $f_{\alpha,1}$, the associated orthonormal sequence is the generalised Laguerre polynomial sequence $\{Q_k^{(\alpha)}\}_{k\geq 0}$ defined for $\alpha>-1$ by

$$Q_0^{(\alpha)}(t)=1 \quad \text{and} \quad Q_k^{(\alpha)}(t) = (-1)^k \left(\frac{\Gamma(\alpha+1+k)}{k!\Gamma(\alpha+1)}\right)^{-\frac{1}{2}} L_k^{(\alpha)}(t), \quad t\in[0,\infty), \tag{6.7}$$

where

$$L_0^{(\alpha)}(t)=1 \quad \text{and} \quad L_k^{(\alpha)}(t) = \sum_{i=0}^{k} \binom{k+\alpha}{k-i} \frac{(-t)^i}{i!}, \quad t\in[0,\infty). \tag{6.8}$$

A summary of their properties which are used in this chapter is given in Section 6.A of the appendix. It is well known that this polynomial sequence is complete (classical results can be found, for instance, in Sansone, 1991). To be able to exploit the additional parameter β, we consider the general gamma density $f_{\alpha,\beta}$ for $\alpha>-1$ and $\beta>0$. With a slight abuse of notation, we let $L^2_{f_{\alpha,\beta}}$ be

$$L^2_{f_{\alpha,\beta}} = \left\{ h:\mathbb{R}\to\mathbb{R} \,|\, \left(\int_{-\infty}^{\infty} h^2(x) f_{\alpha,\beta}(x)\,dx\right)^{1/2} < \infty \right\} \tag{6.9}$$

and the associated inner product will hence be denoted by $<h,l>_{f_{\alpha,\beta}}$ for h,l in $L^2_{f_{\alpha,\beta}}$. In this case, the associated family of orthogonal polynomials is $\{L_k^{(\alpha,\beta)}\}_{k\geq 0}$, defined for $\alpha>-1$ and $\beta>0$ as

$$L_k^{(\alpha,\beta)}(t) = L_k^{(\alpha)}(\beta t), \quad t>0, \tag{6.10}$$

and the orthonormal polynomials are $\{Q_k^{(\alpha,\beta)}\}_{k\geq 0}$, defined for $\alpha > -1$ and $\beta > 0$ as

$$Q_k^{(\alpha,\beta)}(t) = Q_k^{(\alpha)}(\beta t), \quad t > 0. \tag{6.11}$$

At the best of our knowledge, the latter seems to not have been thoroughly investigated in the literature. However, an application similar to ours can be found in Oakley (1990) where they are referred to as *extended* Laguerre polynomials. It is clear that when $\beta = 1$ they reduce to the generalised Laguerre polynomials. In the upcoming proposition we collect some of their properties (two of which have already been mentioned above) which will be fundamental in the following. They are almost a direct consequence of the properties of $\{Q_k^{(\alpha)}\}_{k\geq 0}$ and $\{L_k^{(\alpha)}\}_{k\geq 0}$.

Proposition 6.1.3. *Consider* $\{L_k^{(\alpha,\beta)}\}_{k\geq 0}$ *and* $\{Q_k^{(\alpha,\beta)}\}_{k\geq 0}$. *We have that*

1. $< L_n^{(\alpha,\beta)}, L_m^{(\alpha,\beta)} >_{f_{\alpha,\beta}} = \frac{\Gamma(\alpha+1+n)}{n!\Gamma(\alpha+1)}$ *if* $m = n$ *and zero otherwise, with* $m, n \in \mathbb{N}$;
2. $< Q_n^{(\alpha,\beta)}, Q_m^{(\alpha,\beta)} >_{f_{\alpha,\beta}} = \delta_{mn}, \quad m, n \in \mathbb{N}$;
3. *the set* $\{Q_k^{(\alpha,\beta)}\}_{k\geq 0}$ *is complete in* $L^2_{f_{\alpha,\beta}}$.

Proof. Consider the operator $B : L^2_{f_{\alpha,\beta}} \to L^2_{f_{\alpha,1}}$ defined as

$$B(h) = h_{\frac{1}{\beta}}, \text{ with } h_{\frac{1}{\beta}}(t) = h(t/\beta), \quad t > 0.$$

Furthermore, note that

$$B(h) = 0 \text{ if and only if } h = 0. \tag{6.12}$$

Then, as for any $t > 0$ we have $f_{\alpha,1}(t) = \frac{1}{\beta} f_{\alpha,\beta}(\frac{t}{\beta})$, by using the change of variable $x \mapsto \frac{x}{\beta}$ it can be shown that for all h and l in $L^2_{f_{\alpha,\beta}}$ we have

$$< h, l >_{f_{\alpha,\beta}} = < B(h), B(l) >_{f_{\alpha,1}}. \tag{6.13}$$

Then, 1. and 2. follow by applying (6.13) and the well known properties of the generalised Laguerre polynomials in (6.7) and (6.8) (see Section 6.A in the appendix). Statement 3. is proven by using (6.12), (6.13) and the completeness of $\{Q_k^{(\alpha)}\}_{k\geq 0}$ to obtain 1. in Definition 6.1.1. □

For the reasons detailed in the preceding paragraphs, in the following we consider the gamma pdf as reference density. In the next subsection, we elaborate further on this and we provide the desired approximation for a FPT pdf g and cdf G of a one dimensional diffusion process.

6.2 The FPT Pdf and Cdf Series Representations

Before proceeding, we specify the setting in which the upcoming results shall be given. Consider a one dimensional diffusion process $\{Y(t)\}_{t\geq 0}$ whose state space is any type of interval I of the real line, with endpoints $l \leq r$. We also allow the possibilities $l = -\infty$ and/or $r = +\infty$. In the following, we shall naturally deal only with regular diffusion processes. More precisely, let $\tau \geq 0$ and suppose $Y(\tau) = y$, $y \in I$. Define

$$T_z = \begin{cases} \inf_{t\geq\tau}\{Y(t) > z\}, & Y(\tau) = y < z \\ \inf_{t\geq\tau}\{Y(t) < z\}, & Y(\tau) = y > z \end{cases} \tag{6.14}$$

where z is in the interior of I. The rv T_z describes the first time the process attains the value z. We call the process *regular* if

$$\mathbb{P}(T_z < \infty | Y(0) = y) > 0$$

whenever $l < y, z < r$. A *regular process* is such that starting from any point in the interior of I, the process can reach any other point in the interior of I with positive probability. The FPT rv with respect to a constant boundary S we are considering is defined as an instance of (6.14) with $\tau = 0$, that is

$$T = \begin{cases} \inf_{t\geq 0}\{Y_t > S\}, & \text{if } y_0 < S, \\ \inf_{t\geq 0}\{Y_t < S\}, & \text{if } y_0 > S, \end{cases} \tag{6.15}$$

where $y_0 = Y(0)$ is the so called starting position. We are interested in applying the contents of Sections 6.1 and 6.1.1 to provide a series expansion for the FPT cdf and pdf, whose definitions we respectively recall as

$$G(S,t|y_0) = \mathbb{P}(T \leq t | Y_0 = y_0), \quad t \in [0,\infty),$$

and

$$g(S,t|y_0) = \frac{d\mathbb{P}(T \leq t | Y_0 = y_0)}{dt}, \quad t \in [0,\infty).$$

In the remainder of this chapter, unless explicitly specified otherwise, for ease of notation, we will refer to the latter two always with simply G and g and as functions of t only, since we will consider fixed S, y_0 and the eventual parameters of the underlying diffusion $\{Y(t)\}_{t\geq 0}$. Note that, by definition, $g(0) = 0$. We shall now state some further definitions, in order to provide a condition which will come in handy later. Let $t \geq 0$. For all $0 < \tau < t$ and

$x, y \in I$ we define the *transition cdf* and *transition pdf* of $\{Y(t)\}_{t\geq 0}$ to be, respectively

$$F(x,t|y,\tau) = \mathbb{P}(Y(t) < x | Y(\tau) = y) \tag{6.16}$$

and

$$f(x,t|y,\tau) = \frac{\partial}{\partial x} F(x,t|y,\tau)\,. \tag{6.17}$$

Condition 6.2.1. *The diffusion $\{(Y_t)\}_{t\geq 0}$ admits a stationary distribution. That is, there exist a pdf Ψ such that*

$$\lim_{t\to\infty} f(x,t|y_0,t_0) = \Psi(x) \tag{6.18}$$

independently of the time and the initial conditions.

The latter is used in the following proposition where we collect two useful implications.

Proposition 6.2.2. *Assume Condition 6.2.1 holds for $\{Y(t)\}_{t\geq 0}$. Then we have that*

1. *the FPT pdf g is unimodal,*
2. *the FPT pdf g has exponential long-time behavior with parameter the inverse mean FPT, that is*

$$g(t) \approx \frac{1}{\mathbb{E}(T)} e^{-\frac{t}{\mathbb{E}(T)}}, \text{ as } t \to \infty,$$

Proof. The proof of 1. can be found in Rosler (1980). For 2. see, for instance, Nobile et al. (1985). □

Recall that we consider the gamma pdf $f_{\alpha,\beta}$ in the following form

$$f_{\alpha,\beta}(t) = \beta(\beta t)^{\alpha} \frac{e^{-\beta t}}{\Gamma(\alpha+1)}, \quad t > 0, \tag{6.19}$$

with scale parameter $\alpha + 1 > 0$ and shape parameter $\beta > 0$. The following simple proposition provides the integral condition that g must satisfy in order to have the desired series representation of the type described in Section 6.1.

Proposition 6.2.3. *Let $a_k^{(\alpha,\beta)} = \mathbb{E}[Q_k^{(\alpha,\beta)}(T)]$ for $k \geq 0$ with $Q_k^{(\alpha,\beta)}$ as in (6.11). We have the following series representation*

$$\frac{g}{f_{\alpha,\beta}} = \sum_{k\geq 0} a_k^{(\alpha,\beta)} Q_k^{(\alpha,\beta)}, \tag{6.20}$$

if and only if

$$\int_0^{\infty} t^{-\alpha} e^{\beta t} g(t)^2 \, \mathrm{d}t < \infty.$$

Proof. From 3. in Proposition 6.1.3, we have that $\{Q_k^{(\alpha,\beta)}\}_{k\geq 0}$ form an orthonormal basis of $L^2_{f_{\alpha,\beta}}$. Thus, to obtain (6.21) it suffices to show that $\frac{g}{f_{\alpha,\beta}} \in L^2_{f_{\alpha,\beta}}$. Then the result follows from (6.9). □

Remark 6.2.4. A starting point of a detailed study of the pointwise convergence of the series representation in (6.21) could be, for instance, the note by Hille (1926). A similar approach has been taken in Theorem 2 of Di Nardo and D'Onofrio (2021). However, such an investigation is out of scope for this thesis. □

Additionally assuming Condition 6.2.1, the next result provides simple sufficient conditions on α, β and on the FPT pdf g so that g admits the series expansion (6.21).

Proposition 6.2.5. *Assume Condition 6.2.1. Suppose* $\beta < 2/\mathbb{E}[T]$ *and* $g(t) = o\left(t^{\delta}\right)$ *with* $\delta > \frac{\alpha}{2}$. *Then* $\frac{g}{f_{\alpha,\beta}} \in L^2_{f_{\alpha,\beta}}$ *and we have*

$$\frac{g}{f_{\alpha,\beta}} = \sum_{k\geq 0} a_k^{(\alpha,\beta)} Q_k^{(\alpha,\beta)}, \tag{6.21}$$

where $a_k^{(\alpha,\beta)} = \mathbb{E}[Q_k^{(\alpha,\beta)}(T)]$ *for* $k \geq 0$ *with* $Q_k^{(\alpha,\beta)}$ *as in* (6.11).

Proof. From Proposition 6.2.3, it suffices to show

$$\int_0^\infty t^{-\alpha} e^{\beta t} g(t)^2 \, \mathrm{d}t = \int_0^1 t^{-\alpha} e^{\beta t} g(t)^2 \, \mathrm{d}t + \int_1^\infty t^{-\alpha} e^{\beta t} g(t)^2 \, \mathrm{d}t = I_1 + I_2 < \infty.$$

Since Condition 6.2.1 holds, from 2. in Proposition 6.2.2 we have that the FPT pdf g has exponential long-time behaviour with parameter the inverse mean FPT, i.e.

$$g(t) \approx \frac{1}{\mathbb{E}(T)} e^{-\frac{t}{\mathbb{E}(T)}}, \text{ as } t \to \infty.$$

Hence, it is necessary and sufficient to have $\beta - \frac{2}{\mathbb{E}[T]} < 0$ for I_2 to be finite. Under the assumption $g(t) = o\left(t^{\delta}\right)$ with $\delta > \frac{\alpha}{2}$, we see that the integrand in I_1 is bounded and we get the desired result. □

Some algebra allows us to rewrite expression (6.21) as

$$\frac{g}{f_{\alpha,\beta}} = \sum_{k\geq 0} \mathcal{B}_k^{(\alpha,\beta)} L_k^{(\alpha,\beta)}, \tag{6.22}$$

with coefficients $\mathcal{B}_0^{(\alpha,\beta)} = 1$ and

$$\mathcal{B}_k^{(\alpha,\beta)} = \Gamma(\alpha+1)(-1)^k a_k^{(\alpha,\beta)} \left(\frac{\Gamma(\alpha+1)\Gamma(\alpha+1+k)}{k!} \right)^{-1/2} = \tag{6.23}$$

$$= \frac{\Gamma(\alpha+1)k!}{\Gamma(\alpha+1+k)} \mathbb{E}[L_k^{(\alpha,\beta)}(T)] = 1 + \sum_{j=1}^{k} \binom{k}{j} \frac{(-\beta)^j \mathbb{E}[T^j]}{(\alpha+j)_j}, \tag{6.24}$$

depending on the moments of T as expected. For $j \in \mathbb{N}$, with $(\alpha+j)_j$ we denote the falling (or descending) factorial, which for $\alpha \in \mathbb{R}$ is defined as $(\alpha+j)_j = \frac{\Gamma(\alpha+j+1)}{\Gamma(\alpha+1)}$. The coefficients $\{\mathcal{B}_k^{(\alpha,\beta)}\}_{k\geq 0}$ satisfy the following useful recurrence relation.

Proposition 6.2.6. *For all $k \geq 1$ we have*

$$\mathcal{B}_k^{(\alpha,\beta)} = \sum_{j=1}^{k} \binom{k}{j} (-1)^{j+1} \mathcal{B}_{k-j}^{(\alpha,\beta)} + \frac{(-\beta)^k \mathbb{E}[T^k]}{(\alpha+k)_k}. \tag{6.25}$$

Proof. By plugging (6.24) into $\mathcal{B}_{k-j}^{(\alpha,\beta)}$ for $j = 1, \ldots, k$ we get

$$\sum_{j=0}^{k} \binom{k}{j} (-1)^j \mathcal{B}_{k-j}^{(\alpha,\beta)} = \mathcal{B}_k^{(\alpha,\beta)} + \sum_{j=1}^{k} \binom{k}{j} (-1)^j \left[1 + \sum_{i=1}^{k-j} \binom{k-j}{i} \frac{(-\beta)^i \mathbb{E}[T^i]}{(\alpha+i)_i} \right]. \tag{6.26}$$

Moreover, by expanding the inner sum in the rhs of (6.26) and grouping with respect to the j-th moment $\mathbb{E}[T^j]$ we have

$$\sum_{j=1}^{k} \binom{k}{j} (-1)^j \left[\sum_{i=1}^{k-j} \binom{k-j}{i} \frac{(-\beta)^i \mathbb{E}[T^i]}{(\alpha+i)_i} \right] = \sum_{j=1}^{k-1} \frac{(-\beta)^j \mathbb{E}[T^j]}{(\alpha+j)_j} \left[\sum_{i=1}^{k-j} (-1)^i \binom{k}{i} \binom{k-i}{j} \right].$$

Since

$$\sum_{i=1}^{k-j} (-1)^i \binom{k}{i} \binom{k-i}{j} = \binom{k}{j} \sum_{i=1}^{k-j} \binom{k-j}{i} (-1)^i = -\binom{k}{j}$$

and $\sum_{j=0}^{k} \binom{k}{j} (-1)^j = 0$ which gives $\sum_{j=1}^{k} \binom{k}{j} (-1)^j = -1$, from (6.26) we get

$$\sum_{j=0}^{k} \binom{k}{j} (-1)^j \mathcal{B}_{k-j}^{(\alpha,\beta)} = \mathcal{B}_k^{(\alpha,\beta)} - 1 - \sum_{j=1}^{k-1} \binom{k}{j} \frac{(-\beta)^j \mathbb{E}[T^j]}{(\alpha+j)_j}. \tag{6.27}$$

Plugging (6.24) into $\mathcal{B}_k^{(\alpha,\beta)}$ after some algebraic manipulation, we get

$$\frac{(-\beta)^k \mathbb{E}[T^k]}{(\alpha+k)_k} = \sum_{j=0}^{k} \binom{k}{j} (-1)^j \mathcal{B}_{k-j}^{(\alpha,\beta)}, \tag{6.28}$$

from which (6.25) follows. □

Using (6.22), a series representation of the FPT cdf G is given in the following statement. Consider $\Phi(a,b,z) = {}_1F_1(a;b;z)$, the confluent hypergeometric function of the first kind, and define $\Phi^{(-\beta)}(a,b,\cdot)$ as $z \mapsto \Phi(a,b,-\beta z)$ for $z \in [0,\infty)$.

Proposition 6.2.7. *Suppose* $\frac{g}{f_{\alpha,\beta}} \in L^2_{f_{\alpha,\beta}}$. *Then the FPT cdf* G *has the following series representation*

$$G = h_{\alpha,\beta} \sum_{k\geq 0} \frac{\Gamma(\alpha+k+1)}{k!} \mathcal{B}_k^{(\alpha,\beta)} \Phi^{(\beta)}(\alpha+k+1, \alpha+2, \cdot), \tag{6.29}$$

where $h_{\alpha,\beta}$ *is defined as* $t \mapsto \frac{(\beta t)^{\alpha+1}}{\Gamma(\alpha+1)}$.

Proof. Note that $G(t) = \int_0^t \frac{g(x)}{f_{\alpha,\beta}(x)} f_{\alpha,\beta}(x)dx$, with $t \in (0,\infty)$. For $n \in \mathbb{N}$, let now $p_n(t) = \sum_{k=0}^n \mathcal{B}_k^{(\alpha,\beta)} L_k^{(\alpha,\beta)}(t)$ and

$$G_n(t) = \int_0^t f_{\alpha,\beta}(x) p_n(x) dx, \quad t \in (0,\infty).$$

Then, it is easily shown that

$$\lim_{n\to\infty} \|G - G_n\|^2_{f_{\alpha,\beta}} = 0.$$

Indeed, we have that

$$\begin{aligned}
\|G - G_n\|^2_{f_{\alpha,\beta}} &= \int_0^\infty \left[\int_0^t \frac{g(x)}{f_{\alpha,\beta}(x)} f_{\alpha,\beta}(x)dx - \int_0^t f_{\alpha,\beta}(x) p_n(x) dx \right]^2 f_{\alpha,\beta}(t)dt = \\
&= \int_0^\infty \left[\int_0^t \left(\frac{g(x)}{f_{\alpha,\beta}(x)} - p_n(x) \right) f_{\alpha,\beta}(x)dx \right]^2 f_{\alpha,\beta}(t)dt \leq \\
&\leq \left\| \frac{g}{f_{\alpha,\beta}} - p_n \right\|_{f_{\alpha,\beta}} \int_0^\infty f_{\alpha,\beta}(t)dt = \left\| \frac{g}{f_{\alpha,\beta}} - p_n \right\|_{f_{\alpha,\beta}},
\end{aligned}$$

which goes to zero as n goes to infinity by the assumption of $\frac{g}{f_{\alpha,\beta}} \in L^2_{f_{\alpha,\beta}}$. Hence, by making use of the formula (Wolfram Research, Inc., 2024)

$$\int_0^t \tau^\alpha e^{-\tau} L_k^{(\alpha)}(\tau)\, \mathrm{d}\tau = \frac{t^{\alpha+1}\Gamma(\alpha+k+1)}{k!} \frac{\Phi(\alpha+k+1;\alpha+2;-t)}{\Gamma(\alpha+2)},$$

to simplify G_n, we obtain that G has the desired $L^2_{f_{\alpha,\beta}}$ series representation. □

6.3 The FPT Approximation

As explained in Section 6.1, an approximation of the FPT pdf g can be then recovered from (6.21) by using a truncation of the series (6.21) up to an order $n \in \mathbb{N}$. Then, let the latter be fixed. We define our approximant $\hat{g}_n$ as

$$\hat{g}_n(t) = \frac{\beta(\beta t)^\alpha e^{-\beta t}}{\Gamma(\alpha+1)} \sum_{k=0}^{n} a_k^{(\alpha,\beta)} Q_k^{(\alpha,\beta)}(t), \quad t>0, \tag{6.30}$$

or, equivalently from (6.22), as

$$\hat{g}_n(t) = \frac{\beta(\beta t)^\alpha e^{-\beta t}}{\Gamma(\alpha+1)} \sum_{k=0}^{n} \mathcal{B}_k^{(\alpha,\beta)} L_k^{(\alpha,\beta)}(t), \quad t>0. \tag{6.31}$$

From the classical Parseval identity in 2. in Definition 6.1.1, the $L^2_{f_{\alpha,\beta}}$-error in replacing g with its approximation $\hat{g}_n$ given in (6.30) is

$$\left\| \frac{g-\hat{g}_n}{f_{\alpha,\beta}} \right\|_{\alpha,\beta} = \left[\sum_{k\geq n+1} \left(a_k^{(\alpha,\beta)}\right)^2 \right]^{1/2} \tag{6.32}$$

where $\| \|_{\alpha,\beta}$ denotes the norm in $L^2_{f_{\alpha,\beta}}$. Thus the error may be estimated by calculating the rate of decrease of $a_k^{(\alpha,\beta)}$ when $k \to \infty$. In this regard, under suitable assumptions on g, we have the following proposition, whose proof relies on classical results (see, e.g., Shohat 1935, or, for a more recent treatment, Gottlieb and Orszag 1977).

Theorem 6.3.1. *Assume Condition 6.2.1 and that the FPT pdf $g \in C^2[0,+\infty)$. If*

$$\beta < \frac{2}{\mathbb{E}[T]} \quad \text{and} \quad g(t) = o(t^\delta) \quad \text{for } t \to 0, \quad \text{with } \delta > \frac{\alpha}{2}+1,$$

then in (6.32) $a_k^{(\alpha,\beta)} = \mathbb{E}[Q_k^{(\alpha,\beta)}(T)] = O(k^{-1})$ as $k \to \infty$.

Proof. Observe that $a_k^{(\alpha,\beta)} = \mathbb{E}[Q_k^{(\alpha,\beta)}(T)] = \mathbb{E}[Q_k^{(\alpha)}(\beta T)]$ gives, after a suitable change of variables,

$$a_k^{(\alpha,\beta)} = \frac{1}{\beta} \int_0^\infty Q_k^{(\alpha)}(t)\, \tilde{g}_{\alpha,\beta}(t)\, t^\alpha e^{-t}\, \mathrm{d}t \quad \text{where} \quad \tilde{g}_{\alpha,\beta}(t) = \frac{g(t/\beta)}{t^\alpha e^{-t}}. \tag{6.33}$$

As the generalised Laguerre polynomials $\{L_k^{(\alpha)}\}$ are eigenfunctions of a Sturm-Liouville problem (see, for instance, Sansone, 1991) with associated eigenvalues $\lambda_k = k$

$$\frac{\mathrm{d}}{\mathrm{d}t}\left(t^{\alpha+1} e^{-t} y'\right) + k t^\alpha e^{-t} y = 0 \quad \text{with} \quad y = y(t), \quad k \geq 1$$

the same happens for the linearly transformed polynomials $\{Q_k^{(\alpha)}\}$ in (6.7). Therefore in (6.33), replace

$$Q_k^{(\alpha)}(t)t^\alpha e^{-t} \quad \text{with} \quad -\frac{1}{k}\frac{\mathrm{d}}{\mathrm{d}t}\left(t^{\alpha+1}e^{-t}y'\right).$$

Integrating by parts the integral in (6.33) and neglecting the constants, the rhs of (6.33) reads

$$a_k^{(\alpha,\beta)} \approx \frac{1}{k}\int_0^\infty t^{\alpha+1}e^{-t}\frac{\mathrm{d}}{\mathrm{d}t}[Q_k^{(\alpha)}(t)]\frac{\mathrm{d}}{\mathrm{d}t}[\tilde{g}_{\alpha,\beta}(t)]\,\mathrm{d}t. \tag{6.34}$$

Indeed we have

$$\lim_{t\to 0}\tilde{g}_{\alpha,\beta}(t)t^{\alpha+1}e^{-t}\frac{\mathrm{d}}{\mathrm{d}t}[Q_k^{(\alpha)}(t)]=0 \quad \text{and} \quad \lim_{t\to\infty}\tilde{g}_{\alpha,\beta}(t)t^{\alpha+1}e^{-t}\frac{\mathrm{d}}{\mathrm{d}t}[Q_k^{(\alpha)}(t)]=0. \tag{6.35}$$

The first limit in (6.35) results by the hypothesis $g(t)=o(t^\delta)$ for $t\to 0$. The second limit in (6.35) follows by taking into account that, since Condition 6.2.1 holds, from 2. in Proposition 6.2.2, we have that g is approximately exponential for $t\to\infty$, with parameter $\mathbb{E}[T]^{-1}$. Integrating by parts the integral in (6.34) and neglecting the constants, the rhs of (6.34) reads

$$a_k^{(\alpha,\beta)} \approx -\frac{1}{k}\int_0^\infty Q_k^{(\alpha)}(t)\frac{\mathrm{d}}{\mathrm{d}t}\left(t^{\alpha+1}e^{-t}\frac{\mathrm{d}}{\mathrm{d}t}[\tilde{g}_{\alpha,\beta}(t)]\right)\mathrm{d}t, \tag{6.36}$$

where similar considerations done for (6.35) apply for recovering

$$\lim_{t\to 0}t^{\alpha+1}e^{-t}Q_k^{(\alpha)}(t)\frac{\mathrm{d}}{\mathrm{d}t}[\tilde{g}_{\alpha,\beta}(t)]=0 \quad \text{and} \quad \lim_{t\to\infty}t^{\alpha+1}e^{-t}Q_k^{(\alpha)}(t)\frac{\mathrm{d}}{\mathrm{d}t}[\tilde{g}_{\alpha,\beta}(t)]=0.$$

Now, in (6.36) set

$$h(t)=\frac{1}{w(t)}\frac{\mathrm{d}}{\mathrm{d}t}\left(t^{\alpha+1}e^{-t}\frac{\mathrm{d}}{\mathrm{d}t}[\tilde{g}_{\alpha,\beta}(t)]\right) \quad \text{with} \quad w(t)=t^\alpha e^{-t}.$$

Applying the Cauchy-Schwarz inequality to the rhs of (6.36), we get

$$\left|\int_0^\infty \sqrt{w(t)}Q_k^{(\alpha)}(t)h(t)\sqrt{w(t)}\,\mathrm{d}t\right|^2 \le \left(\int_0^\infty w(t)[Q_k^{(\alpha)}(t)]^2\,\mathrm{d}t\right)\left(\int_0^\infty [h(t)]^2 w(t)\,\mathrm{d}t\right)$$

which is finite and not depending on the order k, if the same is true for both integrals on the lhs. Observe that the first integral corresponds to the orthonormality condition of the family $\{Q_k^{(\alpha)}\}_{k\ge 0}$ and so it is finite and not depending on k. The second integral does not depend on k and is finite if

the integrand is smooth and the limits for $t \to 0$ and $t \to \infty$ are finite. Note that

$$\frac{h(t)}{\sqrt{w(t)}} = \frac{(\alpha+1-t)\frac{\mathrm{d}}{\mathrm{d}t}[\tilde{g}_{\alpha,\beta}(t)] + t\frac{\mathrm{d}^2}{\mathrm{d}t^2}[\tilde{g}_{\alpha,\beta}(t)]}{t^{-\alpha/2}e^{t/2}}.$$

Thus for $t \to \infty$ and $\beta < 2/\mathbb{E}[T]$ we have

$$\lim_{t\to\infty} \frac{h(t)}{\sqrt{w(t)}} = \lim_{t\to\infty} \frac{e^{z[1-(\beta\mathbb{E}[t])^{-1}]}}{z^{\alpha/2}e^{z/2}} + \lim_{t\to\infty} \frac{e^{z[1-(\beta\mathbb{E}[t])^{-1}]}}{z^{\alpha/2-1}e^{z/2}} = 0$$

using again that $g(t) = O(e^{-t/(\beta\mathbb{E}[T])})$. Instead for $t \to 0$ the limit reduces to

$$\lim_{t\to 0} \frac{h(t)}{\sqrt{w(t)}} = \lim_{t\to 0} e^{t/2}t^{\delta-\alpha/2-1} + \lim_{t\to 0} e^{t/2}t^{\delta-\alpha} = 0$$

if $\delta > 1 + \alpha/2$. □

The request of the existence of the second derivative in Theorem 6.3.1 is a reasonable assumption for our application in 6.5.1, as shall become clearer therein. Combining the coefficients of $L_k^{\alpha}(\beta t)$ with the same power of t, expression (6.31) can be rearranged as follows

$$\hat{g}_n(t) = f_{\alpha,\beta}(t)\, p_n(t), \quad t \geq 0, \tag{6.37}$$

where

$$p_n(t) = \sum_{k=0}^{n} h_{n,k} \frac{(-\beta t)^k}{k!} \quad \text{with} \quad h_{n,k} = \sum_{j=k}^{n} \mathcal{B}_j^{(\alpha,\beta)} \binom{\alpha+j}{j-k}, \tag{6.38}$$

$\mathcal{B}_0^{(\alpha,\beta)} = 1$ and

$$\binom{\alpha+j}{j-k} = \begin{cases} 1, & j = k, \\ \frac{(\alpha+j)(\alpha+j-1)\cdots(\alpha+k+1)}{(j-k)!}, & j > k. \end{cases} \tag{6.39}$$

Recall that from 1. in Proposition 6.1.2 we have

$$\int_0^{\infty} \hat{g}_n(t)\mathrm{d}t = 1 \tag{6.40}$$

for all $n \geq 0$, and that from 2. in the same proposition the first n moments of $\hat{g}_n$ are the same of T. Unfortunately, as already said in the beginning of this chapter and in the end of Section 6.1, $\hat{g}_n$ is not guaranteed to be a pdf since negative values can occur. Indeed, the values assumed by the polynomial

p_n in (6.37) are not necessarily non-negative. Before dealing with this issue more in detail in Section 6.4.5, we provide a simple preliminary result which might be helpful. Indeed, p_n may hold non-negative values on the for high values of t and in a right-handed neighbourhood of the origin, depending on the sign of some coefficients in (6.38). These conditions are established in the following proposition.

Proposition 6.3.1. *Suppose $p_n(t) > 0$ for all $t > 0$. Then $(-1)^n h_{n,n} \geq 0$ and $h_{n,0} \geq 0$. Conversely, if $h_{n,0} > 0$ and $(-1)^n h_{n,n} > 0$, there exist $t_1 > 0$ and $t_2 > 0$ such that $p_n(t) > 0$ in $(0, t_1) \cup (t_2, +\infty)$, with t_1 and t_2 not necessarily distinct or finite.*

Proof. From (6.38) we have

$$p_n(0) = h_{n,0} \text{ and } p_n(t) \sim (-1)^n h_{n,n} t^n \frac{\beta^n}{n!}, \text{ for } t \to \infty.$$

Since $\beta > 0$ and $p_n \in C(0, \infty)$, the results follow from the sign permanence theorem. □

Finally, note that by integrating $\hat{g}_n$ in (6.37) over $(0, t)$, or by truncating (6.29), brief calculations reveal that an approximation $\hat{G}_n$ of the FPT cdf G can be defined as

$$\hat{G}_n(t) = \frac{1}{\Gamma(\alpha+1)} \sum_{k=0}^{n} \frac{(-1)^k}{k!} h_{n,k} \left[\Gamma(\alpha+k+1) - \Gamma(\alpha+k+1, \beta t)\right], \quad t \in [0, \infty) \tag{6.41}$$

where $\Gamma(a, t) = \int_t^\infty \tau^{a-1} e^{-\tau} \mathrm{d}\tau$ is the incomplete Gamma function.

In the following sections we will deal with the actual computation of $\hat{g}_n$ in (6.37).

6.4 Computational Issues

6.4.1 An Iterative Procedure

In the following we propose a procedure to recover the n-th approximation $\hat{g}_n$ which reduces the overall computational time to evaluate (6.37) and can be implemented iteratively. This procedure relies on a recurrence relation obtained using nested products. Indeed, thanks to the representation (6.37), we can exploit the fact that $p_n(t)$ can be efficiently evaluated for any $t > 0$ using the recurrence relation

$$d_{n,i}(t) = h_{n,i-1} - \frac{\beta t}{i} d_{n,i+1}(t) \quad \text{for } i \in \{n, n-1, \ldots, 1\}, \tag{6.42}$$

with the initial condition $d_{n,n+1}(t) = h_{n,n}$, since the last value gives $d_{n,1}(t) = p_n(t)$.

We stress out that in practice we cannot take n arbitrarily large, due to numerical errors which may incur in the computation of the coefficients $\{\mathcal{B}_j^{(\alpha,\beta)}\}_{j\geq 0}$ as in (6.24). Clearly, using infinite precision operations would avoid this issue. Software tools like `Mathematica` and `R` allow for arbitrarily large but finite precision. This swiftly becomes prohibitively slow. In some scenarios, resorting to this may be needed to obtain a satisfactory approximation, as will be seen in the numerical examples of Section 6.6.3. However, in most of the cases we can rely on the following iterative procedure to compute $\hat{g}_n$ in (6.37).

To this aim, let us first observe that (6.42) can be easily updated to $n+1$, returning $d_{n+1,1}(t) = p_{n+1}(t)$ from the initial condition $d_{n+1,n+2}(t) = h_{n+1,n+1}$. Indeed, from the second equation in (6.38), a little thought let us recover $\{h_{n+1,i}\}_{i=0}^{n+1}$ from $\{h_{n,i}\}_{i=0}^{n}$ using

$$h_{n+1,i} = \begin{cases} \mathcal{B}_{n+1}^{(\alpha,\beta)} & \text{for } i = n+1 \\ h_{n,i} + \mathcal{B}_{n+1}^{(\alpha,\beta)} \binom{\alpha+n+1}{n+1-i} & \text{for } i = 0,\dots,n. \end{cases} \tag{6.43}$$

The coefficient $\mathcal{B}_{n+1}^{(\alpha,\beta)}$ can be computed from $\{\mathcal{B}_j{}^{(\alpha,\beta)}\}_{j=1}^{n}$ as well, using the recursion formula in Proposition 6.25. Algorithm 2 provides a pseudo-code explanation of the iterative procedure.

6.4.2 On the Order n of the Approximation

It is clear that the choice of the truncating order n is of upmost importance. Since in the previous section we have provided an iterative procedure to compute $\hat{g}_n$ in (6.37), the next step is choosing a stopping criterion for Algorithm 2. Naturally, one should start by considering a convergence-based stopping criterion, which is commonly seen in the literature along with graphical checks (see, e.g., Provost and Ha, 2016). That is, the practitioner shall run the iterative procedure until the smallest n such that, for a user chosen tolerance $\epsilon > 0$, it holds that

$$\|\hat{g}_n - \hat{g}_{n-1}\|_{L^2} < \epsilon, \tag{6.44}$$

where L^2 is the usual space of square integrable functions. However, for our two applications in Sections 6.6.2 and 6.6.2, this criterion proved to be affected by numerical instability. The reason is that condition (6.44) may be satisfied for large values of n where the corresponding accumulation of

Algorithm 2: Iterative procedure for computing $\hat{g}_n$

1: **For** $n=1$**:**

1. Compute $\mathbb{E}[T]$
2. Set $\mathcal{B}_0^{(\alpha,\beta)}=1$ and $\mathcal{B}_1^{(\alpha,\beta)}=1-\beta\frac{\mathbb{E}[T]}{(\alpha+1)_1}$
3. Compute $h_{0,1}$ and $h_{1,1}$ using (6.38), from which p_1 and $\hat{g}_1$ are computed

2: **For** $n>1$**:**

1. Compute $\mathbb{E}[T^n]$
2. Compute $\mathcal{B}_n^{(\alpha,\beta)}$ by the recurrence formula

$$\mathcal{B}_n^{(\alpha,\beta)}=\sum_{j=1}^{n}\binom{n}{j}(-1)^{j+1}\mathcal{B}_{n-j}^{(\alpha,\beta)}+\frac{(-\beta)^n\mathbb{E}[T^n]}{(\alpha+n)_n}$$

3. Compute $\hat{g}_n$

Set $h_{n,n}=\mathcal{B}_n^{(\alpha,\beta)}$

Update $h_{n-1,i}$ to $h_{n,i}+\mathcal{B}_n^{(\alpha,\beta)}\binom{\alpha+n}{n-i}$

Run recursion (6.42) and set $\hat{g}_n=f_{\alpha,\beta}\,d_{n,1}$.

numerical error has already compromised the approximation. Therefore, we looked into another way to stop the iteration procedure in the previous section at a n large enough such that a satisfactory approximation is obtained, while, at the same time, numerical errors are not yet encountered. Our choice, which will be detailed below, relies on the subsequent normalization condition satisfied by the sequence $\{h_{n,i}\}$ in (6.38).

Proposition 6.4.1. *For all $n \geq 0$ we have*

$$h_{n,0} + \sum_{i=1}^{n} \frac{(-1)^i}{i!} h_{n,i} (\alpha + i)_i = 1. \tag{6.45}$$

Proof. From 1. in Proposition 6.1.2, we have

$$\sum_{i=0}^{n} \frac{(-1)^i}{i!} \beta^i h_{n,i} \int_0^\infty t^i f_{\alpha,\beta}(t) \mathrm{d}t = 1 \quad \text{for all } \ n \geq 0.$$

The result follows by observing that the integrals in the lhs of (6.101) are the moments of the gamma rv with pdf $f_{\alpha,\beta}$, that is

$$\int_0^\infty t^i f_{\alpha,\beta}(t) \mathrm{d}t = \frac{\Gamma(\alpha + 1 + i)}{\beta^i \Gamma(\alpha + 1)}, \quad i \in \mathbb{N}.$$

□

As a result of an extended number of numerical experiments, some of which will be reported later in Section 6.6, increasing n in (6.37) as long as (6.45) is satisfied with a fixed level of tolerance proved to guarantee a satisfactory approximation. With the more conservative aim of obtaining a positive approximant, we further propose to join the normalisation condition (6.45) with an additional condition following from Proposition 6.3.1. This condition guarantees an order n of approximation such that $p_n(t)$ is positive close to the origin and as $t \to +\infty$. Set

$$\hat{h}_n = h_{n,0} + \sum_{i=1}^{n} \frac{(-1)^i}{i!} h_{n,i} (\alpha + 1)_i. \tag{6.46}$$

For a fixed $\varepsilon > 0$, the iterative procedure for computing $\hat{g}_n$ in (6.37) is run as long as

$$(|\hat{h}_{n+1} - 1| < \varepsilon \ \text{ for a fixed } \ \varepsilon > 0) \text{ or } (h_{n,0} > 0 \text{ and } (-1)^n h_{n,n} > 0) \tag{6.47}$$

is fulfilled.

6.4.3 On the Choice of α and β

Using a reference density such as $f_{\alpha,\beta}$ introduces two additional parameters that can be adjusted to enhance the approximant $\hat{g}_n$. Denote with $X_{\alpha,\beta}$ the rv having pdf $f_{\alpha,\beta}$ in (6.19). Heuristically, one would like the reference density $f_{\alpha,\beta}$ to resemble as much as possible the unknown objective pdf g. Since we suppose that g is unknown while the moments (resp. cumulants) of T are known, a classical method (see, e.g., Provost and Ha, 2016) consists in matching the first two moments (resp. cumulants) of $X_{\alpha,\beta}$ with the first two moments (resp. cumulants) of T. Let us underline that a range of possibilities was investigated for α and β in this regard. However, what proved to be the more consistent in terms of the resulting approximation was the following "simpler" choice. That is we set

$$\alpha = \frac{c_1^2[T]}{c_2[T]} - 1 = \frac{\mathbb{E}^2[T]}{\mathbb{V}[T]} - 1 \quad \text{and} \quad \beta = \frac{c_1[T]}{c_2[T]} = \frac{\mathbb{E}[T]}{\mathbb{V}[T]}, \tag{6.48}$$

since with these choices we have

$$\mathbb{E}[X_{\alpha,\beta}] = \frac{\alpha+1}{\beta} = \mathbb{E}[T] \quad \text{and} \quad \mathbb{E}[X_{\alpha,\beta}^2] = \frac{(\alpha+1)(\alpha+2)}{\beta^2} = \mathbb{E}[T^2].$$

As mentioned above, this choice conceptually falls within the classical method of moments, and, moreover, from 3. in Proposition 6.1.2, also simplifies the computation of $\hat{g}_n$ since its first two coefficients will be zero, that is $\mathcal{B}_1^{(\alpha,\beta)} = \mathcal{B}_2^{(\alpha,\beta)} = 0$ in (6.31).

Remark 6.4.2. Note that, under the choice in (6.48), the conditions on g involving α and β in Proposition 6.2.5 become conditions involving the first two moments (or cumulants) of T, which we assume to be known. □

Different choices are also suggested in Belt and den Brinker (1997) where results concerning the determination of the two parameters α and β are presented. Unfortunately, adopting these choices requires a knowledge of the FPT pdf not depending on α and β, which is in general not true.

Now, note that the choices of α and β in (6.48) return

$$c_\nu[X_{\alpha,\beta}] = c_\nu[T] = \frac{\sqrt{c_2[T]}}{c_1[T]}$$

where c_ν denotes the coefficient of variation. In such a case the first equation in (6.48) reads $\alpha + 1 = (c_\nu[T])^{-2}$. Thus an higher coefficient of variation of T

reduces $\alpha+1$ and increases the chance of a vertical asymptote of the gamma pdf $f_{\alpha,\beta}$ in 0. As we know that $g(0)=0$, the occurrence of this vertical asymptote may cause of numerical issues which are further exacerbated if the FPT pdf is flat with a large mean value and a heavy right tail with a large variance. To address this issue, an useful tool is the employment of a suitable standardization technique. The concept is straightforward and consists in constructing the approximation $\tilde{g}_n$ of the pdf $\tilde{g}$ corresponding to

$$\tilde{T} = T/\sigma_T, \tag{6.49}$$

where σ_T is the standard deviation of T. The approximated FPT pdf and cdf can be recovered as

$$\hat{g}_n(t) = \frac{1}{\sigma_T}\tilde{g}_n\left(\frac{t}{\sigma_T}\right) \quad \text{and} \quad \hat{G}_n(t) = \tilde{G}_n\left(\frac{t}{\sigma_T}\right)$$

respectively. As $c_2[\tilde{T}] = \mathrm{Var}[\tilde{T}] = 1$, from (6.48) the parameters of the gamma pdf are

$$\tilde{\alpha} := (\mathbb{E}[\tilde{T}])^2 - 1 \text{ and } \tilde{\beta} := \mathbb{E}[\tilde{T}]$$

so that $\tilde{\alpha}+1 = \tilde{\beta}^2 = (c_\nu[X_{\tilde{\alpha},\tilde{\beta}}])^{-2} = (c_\nu[\tilde{T}])^{-2}$ and $\mathrm{Var}[X_{\tilde{\alpha},\tilde{\beta}}] = 1$. The advantage of this strategy is that it should provide an "initial" approximation (the gamma pdf) with a shape resembling the desired shape of the FPT pdf. Moreover, the moments $\mathbb{E}[\tilde{T}^n]$ grow slower than the moments $\mathbb{E}[T^n]$, leading to an observed improvement of numerical stability.

Since the pdf g has support $(0,\infty)$, further information on the shape of the density can be recovered using some dispersion indices that work as the coefficient of variation but provide further global statistical information (Kostal et al., 2011). In particular we consider the relative entropy based dispersion coefficient

$$c_h := \frac{\sigma_h}{\mathbb{E}(T)} = \frac{1}{\mathbb{E}(T)} \exp\left\{\int_0^\infty g(t)\ln g(t)\,\mathrm{d}t - 1\right\}. \tag{6.50}$$

The value of σ_h quantifies how evenly is the pdf over $(0,\infty)$. Moreover, in logarithmic scale, c_h is inversely proportional to the Kullback-Leibler dinstance of the pdf g from the exponential density with mean $\mathbb{E}[T]$. For densities resembling the exponential distribution, the coefficients c_ν and c_h are approximately equal to 1.

6.4.4 On the Monotonicity of $\hat{G}_n$

Fix $\Delta t > 0$. The computation of the FPT cdf approximant $\hat{G}_n$ as in (6.41) can benefit from the iterative calculation of the increments

$$\Delta\hat{G}_n(t) = \frac{1}{\Gamma(\alpha+1)} \sum_{k=0}^{n} \frac{(-1)^k}{k!} h_{n,k} \left[\Gamma(\alpha+k+1,\beta t) - \Gamma(\alpha+k+1,\beta(t+\Delta t))\right] \tag{6.51}$$

where

$$\Delta\hat{G}_n(t) = \hat{G}_n(t+\Delta t) - \hat{G}_n(t), \quad t > 0.$$

Since there could exists $t > 0$ such that $\hat{g}_n(t) < 0$, the corresponding increments $\Delta\hat{G}_n(t)$ in (6.51) might be negative and, consequently, the approximated cdf may turn out to be decreasing in these intervals. Furthermore, if in a right-hand neighborhood of the origin the first increments are already negative, this might cause negative values of $\hat{G}_n$ itself in the same neighborhood. We propose a simple correction to this last scenario. Set $\tau_0 = 0$ or $\tau_0 = \min\{t > 0 | \Delta\hat{G}_n(t) < 0\}$, depending if $\Delta\hat{G}_n(\Delta t)$ is negative or not, and $\tau_1 = \min\{t > \tau_0 | \hat{G}_n(t) > \hat{G}_n(\tau_0)\}$. Then, iteratively find the intervals $[\tau_i, \tau_{i+1}]$ such that

$$\tau_i = \min\{t > \tau_{i-1} | \Delta\hat{G}_n(t) < 0\} \quad \text{and} \quad \tau_{i+1} = \min\{t > \tau_i | \hat{G}_n(t) > \hat{G}_n(\tau_i)\}$$

in order to replace $\hat{G}_n(t)$ with an increasing line for $t \in [\tau_i, \tau_{i+1}]$. Let us refer with $\hat{G}_n^{corr}$ to the result of this procedure when applied to $\hat{G}_n$. We stress out that for the cases here considered in Section 6.6.3, if present, the intervals $t \in [\tau_i, \tau_{i+1}]$ had a small amplitude and that is why a simple line has been considered. The usefulness of this procedure is twofold. Firstly, one obtains an approximated cdf $\hat{G}_n^{corr}$ which is positive and increasing. Secondly, carrying out a numerical differentiation of $\hat{G}_n^{corr}$ would provide an approximation to the FPT pdf g which is positive. However, since by construction $\hat{G}_n^{corr}(t)$ is linear for $t \in [\tau_i, \tau_{i+1}]$, the pdf will be constant in $[\tau_i, \tau_{i+1}]$. This construction of a positive approximant for the FPT pdf g is computationally simple, but it is clear that it might fail to fit some properties which a FPT pdf should enjoy. The following subsection provides a numerical procedure which is equally simple but aims to provide a positive approximant to g that takes more into account the properties of a FPT pdf.

6.4.5 On the Positivity of $\hat{g}_n$

Although the stopping criteria in (6.47) take into account Proposition 6.3.1, there is no guarantee that $\hat{g}_n(t)$ is non-negative for t in $(0, \infty)$ depending on

the values assumed by $p_n(t)$ for $t \in (t_1, t_2)$ in Proposition 6.3.1. We propose the following strategy to address this issue.

Suppose then there exists $t_{1,\text{neg}}, t_{2,\text{neg}}$ in $(0,\infty)$ such that $(t_{1,\text{neg}}, t_{2,\text{neg}}) \subseteq (t_1, t_2)$ and $\hat{g}_n(t) < 0, t \in (t_{1,\text{neg}}, t_{2,\text{neg}})$. In what follows we develop a simple numerical procedure whose aim is to replace $\hat{g}_n(t)$ with a suitable positive function $p(t)$, for t in a generic interval $(t_1', t_2') \supseteq (t_{1,\text{neg}}, t_{2,\text{neg}})$. As we continue, we shall further assume Condition 6.2.1. Then, from 1. in Proposition 6.2.2 we have that g is unimodal and, as will be clear later, the following construction can be made conceptually easier. Indeed, suppose then $m^* = \max_{t\in(0,\infty)} \hat{g}_n(t)$ be the unique approximated mode of the FPT rv T. Consequently, $\hat{g}_n$ could be negative in an interval located to the right or/and to the left of m^*. In both cases, a natural approach would consist in building a fourth-degree polynomial p interpolating smoothly $(t_1', \hat{g}_n(t_1'))$ and $(t_2', \hat{g}_n(t_2'))$, fulfilling the additional constraints imposed by the conservation of probability mass

$$\int_{t_1'}^{t_2'} p(t)\,\mathrm{d}t = \int_{t_1'}^{t_2'} \hat{g}_n(t)\,\mathrm{d}t, \tag{6.52}$$

as well as positivity and monotonicity. Since such a polynomial is unique, the last two remaining conditions could only possibly be satisfied by a computationally heavy choice of t_1' and t_2' or by selecting an higher degree polynomial. In both cases, the procedure could become cumbersome to carry out. Therefore, in the following we propose a different lighter approach both to determine numerically (t_1', t_2') and to correct $\hat{g}_n$, taking into account the conditions required on the FPT pdf g in Theorem 6.3.1.

In accord with the discussion above, two possible scenarios might occur:

a) $\hat{g}_n(t) < 0$ for $t \in (t_{1,\text{neg}}, t_{2,\text{neg}})$ with $t_{2,\text{neg}} < m^*$

b) $\hat{g}_n(t) < 0$ for $t \in (t_{1,\text{neg}}, t_{2,\text{neg}})$ with $t_{1,\text{neg}} > m^*$.

Therefore, two procedures were developed which take into account the behavior of the FPT pdf either in a right-handed neighborhood of the origin - Case *a)* - and on the tail - Case *b)* - with the aim of minimizing the number of parameters involved while streamlining the writing and the procedure. Note that no conservation of the probability mass (6.52) is required in this strategy. Empirically, this is justified by the fact that in all the numerical examples in Section 6.6.3 the negative areas are so small that the mass involved can be considered negligible.

Case a)

Since $g(0)=0$, we set $t_1' := 0$. To reduce the number of parameters, we assume $p(t)=at^\delta$ with $a>0$ and $\delta > \frac{\alpha}{2}+1$ according to Theorem 6.3.1. Therefore, the correction of $\hat{g}_n(t)$ is defined as

$$\hat{g}_n^{corr}(t) = \begin{cases} at^\delta & 0 \le t \le t_2', \\ \hat{g}_n(t) & t > t_2'. \end{cases} \tag{6.53}$$

In order to achieve a certain level of smoothness, a and δ in (6.53) are chosen such that

$$p(t_2') = \hat{g}_n(t_2') \quad \text{and} \quad \frac{\mathrm{d}}{\mathrm{d}t}p(t)\Big|_{t=t_2'} = \frac{\mathrm{d}}{\mathrm{d}t}\hat{g}_n(t)\Big|_{t=t_2'}.$$

Finally, we set

$$t_2' := \min\left\{ t \in (0, m^*) \,\middle|\, \int_0^t \hat{g}_n(t)\,\mathrm{d}t > 0 \right\}$$

to avoid an excessive increment of the probability mass when $\hat{g}_n$ is replaced by $\hat{g}_n^{corr}$. Fig. 6.4.1-a) shows an example of negative $\hat{g}_n(t)$ for t close to the origin together with its correction $\hat{g}_n^{corr}(t)$, as obtained by the procedure described above.

Case b)

We assume $p(t)=ae^{bt}$, with $a>0$ and $b<0$, according to 2. in Proposition 6.2.2, since we have assumed Contion 6.2.1. Therefore, the correction of $\hat{g}_n(t)$ is defined as

$$\hat{g}_n^{corr}(t) = \begin{cases} \hat{g}_n(t) & t < t_1' \\ ae^{bt} & t_1' \le t \le t_2' \\ \hat{g}_n(t) & t > t_2'. \end{cases} \tag{6.54}$$

In this case, a and b are chosen such that

$$p(t_1') = \hat{g}_n(t_1') \quad \text{and} \quad p(t_2') = \hat{g}_n(t_2').$$

In order to fit a decreasing exponential function in the interval (t_1', t_2'), the endpoints t_1' and t_2' are chosen such that

$$t_2' = \min\left\{ t > t_2 \,\middle|\, \frac{\mathrm{d}}{\mathrm{d}t}\hat{g}_n(t) < 0 \right\} \quad \text{and} \quad t_1' = \min\left\{ t < t_1 \,|\, \hat{g}_n(t) > \hat{g}_n(t_2') \right\}.$$

Fig. 6.4.1-b) shows an example of negative $\hat{g}_n(t)$ for values of t larger than the mode together with its correction $\hat{g}_n^{corr}(t)$, as obtained by the procedure just described.

We end this subsection by stating that a similar procedure could clearly be followed even when g is not unimodal, albeit with less "automation": graphical checks of $\hat{g}_n$ could be used to individuate the possible intervals of negativity and a suitable positive function p could then be built on a case to case basis.

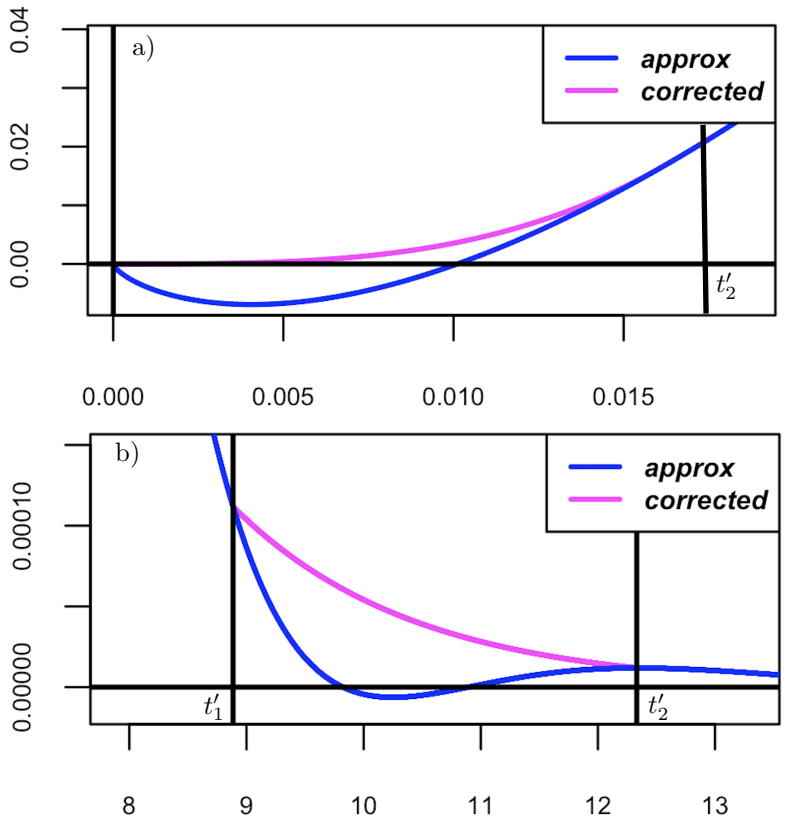

Figure 6.4.1: In a) plots of the approximation $\hat{g}_n$ and of its correction $\hat{g}_n^{corr}$ (6.53) over the interval $(0,t_2')$ are given for $n=10$ and parameters $y_0=0$, $\mu=3$, $S=10$, $c=-10$, $\sigma=1.2,\tau=0.2$ (see case C_{cir} in Section 6.6.3). In b), plots of the approximation $\hat{g}_n$ and of its correction $\hat{g}_n^{corr}$ (6.54) over the interval (t_1',t_2') are given for $n=9$ and parameters $y_0=0.2$, $\mu=0.9$, $S=1$, $c=0$, $\sigma=1.2,\tau=2/3$ (see case A_{cir} in Section 6.6.3).

6.5 The CIR Process, the GBM and their First Passage Time Problems

In the following two subsections we briefly recall the definition and some facts about the FPT problem for the two processes on which we have applied the above explained approximation method.

6.5.1 The CIR Process

As already mentioned in the introduction, we denote by CIR process the unique strong solution of the following stochastic differential equation (see, e.g., Feller, 1951)

$$dY(t) = (-\tau Y(t) + \mu)\,\mathrm{d}t + \sigma\sqrt{Y(t) - c}\,\mathrm{d}W(t), \quad t \geq 0, \tag{6.55}$$

where $\{W(t)\}_{t\geq 0}$ is a standard Brownian motion, $c \leq 0$, $\tau > 0$, $\mu \in \mathbb{R}$, $\sigma > 0$ and $Y_0 = y_0$. The state space of the process is the interval $(c, +\infty)$. Depending on the underlying parameters, the endpoints c and $+\infty$ can or cannot be reached in a finite time. According to the Feller classification of boundaries (Karlin and Taylor, 1981), c is an entrance boundary if it cannot be reached by $\{Y(t)\}_{t\geq 0}$ in finite time, and there is no probability flow to the outside of the interval $(c, +\infty)$. In particular,

$$c \text{ is an entrance boundary if } s := 2(\mu - c\tau)/\sigma^2 \geq 1.$$

This will be assumed to hold as we continue.

In the following, we denote with T_{cir} the FPT rv for the CIR process and with g_{cir} (resp. G_{cir}) the corresponding pdf (resp. cdf). The Laplace transform $LT[g_{cir}](z)$ of g_{cir} is such that $LT[g_{cir}](z) = 1$ if $y_0 \equiv S$ and $LT[g_{cir}](z) < +\infty$ for any different y_0 (see, for example, Masoliver and Perelló, 2012). A closed form expression of $LT[g_{cir}](z)$ can be found in D'Onofrio et al. (2018) and reads as

$$LT[g_{cir}](z) = \frac{\Phi\left(\frac{z}{\tau}, s, \frac{2\tau(y_0 - c)}{\sigma^2}\right)}{\Phi\left(\frac{z}{\tau}, s, \frac{2\tau(S - c)}{\sigma^2}\right)}, \quad z > 0 \tag{6.56}$$

where $\Phi(a, b, z) = {}_1F_1(a; b; z)$ is the confluent hypergeometric function of the first kind (or Kummer's function). Except for the case $S = 0$, the Laplace transform in (6.56) cannot be inverted explicitly, as explained for example in Martin et al. (2011). As known, information on the moments of T_{cir}

could be obtained by direct derivation of $LT[g_{cir}]$, a way which however is not feasible in general for the CIR process. Fortunately enough, expression (6.56) has been used in a clever way in Di Nardo and D'Onofrio (2021) to obtain the cumulants of T_{cir}, which are given in the following. Recall that, if T has moment generating function $\mathbb{E}[e^{zT}] < \infty$ for all z in an open interval about 0, then its cumulants $\{c_k(T)\}_{k\geq 1}$ are such that

$$\sum_{k\geq 1} c_k(T)\frac{z^k}{k!} = \log \mathbb{E}[e^{zT}]$$

for all z in some (possibly smaller) open interval about 0. Using the logarithmic polynomials $\{\mathcal{L}_k\}$ as in (2.12), the FPT cumulants of the CIR process can be expressed as (Di Nardo and D'Onofrio, 2021)

$$c_k(T_{cir}) = (-\tau)^{-k}[c_k^*(y_0) - c_k^*(S)], \quad k \geq 1, \tag{6.57}$$

where

$$c_k^*(w) = \mathcal{L}_k\left[h_1\left(\frac{2\tau(w-c)}{\sigma^2}\right), h_2\left(\frac{2\tau(w-c)}{\sigma^2}\right), \ldots, h_k\left(\frac{2\tau(w-c)}{\sigma^2}\right)\right], \tag{6.58}$$

with $h_j(y) = j!\sum_{n\geq j} \begin{bmatrix} n \\ j \end{bmatrix} \frac{y^n}{n!\langle s\rangle_n}$, for $j \in \{1,2,\ldots,k\}$, $\begin{bmatrix} n \\ j \end{bmatrix}$ the unsigned Stirling numbers of first type and $\langle\cdot\rangle_n$ the n-th rising factorial.

Then, FPT moments of the CIR process can be obtained from FPT cumulants (Di Nardo and D'Onofrio, 2021) using the complete Bell polynomials $\{B_k\}$ as in (2.5) and $\{c_k^*\}$ given in (6.58), that is

$$\begin{aligned}&\mathbb{E}\left[T_{cir}^k\right] \\ &= \frac{(-1)^k}{\tau^k}\sum_{i=0}^{k}\binom{k}{i} B_{k-i}\left[c_1^*(y_0), \ldots, c_{k-i}^*(y_0)\right] B_i[-c_1^*(S), \ldots, -c_i^*(S)]\end{aligned} \tag{6.59}$$

for $k \in \mathbb{N}^+$. The following well-known (see, e.g., Di Nardo and Senato, 2006) recursion formula, which can be found in Section 2.A of the Appendix of Chapter 2, is particularly convenient from a computational point of view and can be used to recover FPT moments from the knowledge of cumulants. For $k \in \mathbb{N}^+$ we have

$$\mathbb{E}\left[T^k\right] = c_k(T) + \sum_{i=1}^{k-1}\binom{k-1}{i-1} c_i(T)\mathbb{E}\left[T^{k-i}\right]. \tag{6.60}$$

In the absence of a threshold, the CIR process admits a stationary distribution which is a shifted gamma distribution with the following shape, scale and location parameters

$$Y_\infty \sim \text{Gamma}\left(s, \frac{1}{2}\frac{\sigma^2}{\tau}, c\right).$$

Therefore, Condition 6.2.1 holds and, provided that $\beta < 2/\mathbb{E}[T]$ and $g(t) = o\left(t^\delta\right)$ with $\delta > \frac{\alpha}{2}$, we can apply Proposition 6.2.5 to obtain that g_{cir} has the following series representation

$$g_{cir} = f_{\alpha,\beta}\left(1 + \sum_{k\geq 1} \mathcal{B}_k^{(\alpha,\beta)} L_k^{(\alpha,\beta)}\right) \text{ with } \mathcal{B}_k^{(\alpha,\beta)} = 1 + \sum_{j=1}^{k} \binom{k}{j} \frac{(-\beta)^j \mathbb{E}[T_{cir}^j]}{(\alpha+j)_j}. \tag{6.61}$$

Remark 6.5.1. Recall that to be able to enjoy the information on the approximation error given in Proposition 6.3.1, we should additionally have $g_{cir} \in C^2([0,\infty))$. Following (Pauwels, 1987, Section 3), who has studied the smoothness of FPT pdfs, we could have alternatively asked that there exists $\varepsilon > 0$ such that $\sigma\sqrt{x} \geq \varepsilon$, for all x in the state space of the process. This condition implies, in this case, the existence and boundedness of at least the first two derivatives of an FPT pdf g. In the present context of the CIR process, to investigate Pauwels' condition, one could follow Feller's classification of the boundaries (Feller, 1952). Using the transition densities of the CIR process (Masoliver and Perelló, 2012), one has to show that the flux through the value ε is zero or that the capacity of the interval $[0,\varepsilon)$ vanishes (Bertini and Passalacqua, 2008). Proceeding analitically proved to be a complex task. However, for a fixed and small $\varepsilon > 0$ we have observed, at least numerically, that the capacity of the interval $[0,\varepsilon)$ (see formula 19 in Bertini and Passalacqua, 2008) goes to zero as μ increases. This means that for a "large enough" choice of μ, the mentioned assumption in Theorem 6.3.1 should be satisfied. □

6.5.2 The Geometric Brownian Motion

The GBM is a regular diffusion process on $(0,+\infty)$. It can be obtained as a transformation of a Brownian motion and for this reason it is also called exponential Brownian motion (see, for instance Yor, 1992). Indeed, if $X(t) = (\mu - \frac{\sigma^2}{2})t + \sigma W(t)$ for $t \geq 0$ is a drifted Brownian motion, then the stochastic process

$$Y(t) = y_0\, e^{(\mu-\frac{\sigma^2}{2})t+\sigma W(t)} = y_0\, e^{X(t)} \qquad \text{with } \mu > \frac{\sigma^2}{2} \tag{6.62}$$

is said GBM with starting point $y_0 \in \mathbb{R}$ and with infinitesimal mean and variance $(\mu - \frac{\sigma^2}{2})y$ and $\sigma^2 y^2$, respectively (see, e.g., Karlin and Taylor, 1981).

The transition pdf of the GBM is known to be a lognormal pdf with parameters $\mu - \frac{\sigma^2}{2}$ and $\sigma\sqrt{t}$. It can be obtained by solving the Fokker-Planck equation, which in the case of the GBM reduces to the canonical form of the heat equation. As we continue, we denote with T_{gbm} the FPT rv for the GBM and with g_{gbm} the corresponding pdf.

The Laplace transform $LT[g_{gbm}](z)$ of g_{gbm} is known and can be found, for instance, in Borodin and Salminen (2002). We have

$$LT[g_{gbm}](z) = \left(\frac{y_0}{S}\right)^{k(z)} \quad \text{with} \quad k(z) = \frac{1}{\sigma^2}\left[\sqrt{\left(\mu - \frac{\sigma^2}{2}\right)^2 + 2\sigma^2 z} + \left(\frac{\sigma^2}{2} - \mu\right)\right].$$

A little thought reveals that it can be rewritten as

$$LT[g_{gbm}](z) = \exp\left\{k(z)\ln\left(\frac{y_0}{S}\right)\right\} = \exp\left\{\frac{a}{b}\left(1 - \sqrt{1 + \frac{2b^2}{a}z}\right)\right\} \tag{6.63}$$

with

$$a = \frac{(\ln S - \ln y_0)^2}{\sigma^2} > 0 \quad \text{and} \quad b = \frac{\ln S - \ln y_0}{\mu - \frac{\sigma^2}{2}} > 0. \tag{6.64}$$

From the well known relationship $M_{T_{gbm}}(z) = LT[g_{gbm}](-z)$ between the moment generating function $M_{T_{gbm}}(z) = \mathbb{E}[\exp(zT_{gbm})]$ and the laplace transform, (6.63) implies that T_{gbm} has an inverse gaussian distribution $IG(a,b)$ of parameters a (the shape) and b (the mean), that is

$$g_{gbm}(t) = \sqrt{\frac{a}{2\pi t^3}}\exp\left(-\frac{a(t-b)^2}{2b^2 t}\right), \quad t > 0. \tag{6.65}$$

The statistical properties of $IG(a,b)$ have been examined in Tweedie (1957). Therein, the moments of T_{gbm} are shown to be

$$\mathbb{E}[T^n] = b^n \sum_{k=0}^{n-1} \frac{(n-1+k)!}{k!(n-1-k)!}\frac{b^k}{(2a)^k} = \frac{\exp(a/b)}{b^{\frac{1}{2}-n}}\mathcal{K}_{n-\frac{1}{2}}\left(\frac{a}{b}\right)\sqrt{\frac{2a}{\pi}} \tag{6.66}$$

where $\mathcal{K}_\nu(z)$ is the modified Bessel function of second type (Gradshteyn and Ryzhik, 1994)

$$\mathcal{K}_{\pm\nu}(z) = \frac{1}{2}\left(\frac{z}{2}\right)^\nu \int_0^\infty x^{-\nu-1}\exp\left(-x - \frac{z^2}{4x}\right)\mathrm{d}x \tag{6.67}$$

under the conditions $|\arg z| < \pi/2$ and $\mathrm{Re}\, z^2 > 0$. Since $z\mathcal{K}_{\nu-1}(z) - z\mathcal{K}_{\nu+1}(z) = -2\nu\mathcal{K}_\nu(z)$ (see, for instance, Gradshteyn and Ryzhik, 1994), the following recursion formula for the FPT moments can be readily derived:

$$\mathbb{E}[T_{gbm}^{n+1}] = \frac{(2n-1)b^2}{a}\mathbb{E}[T_{gbm}^n] + b^2\mathbb{E}[T_{gbm}^{n-1}], \quad n \geq 1 \tag{6.68}$$

with $\mathbb{E}[T_{gbm}^0] = 1$ and $\mathbb{E}[T_{gbm}] = b$.

A new and alternative expression of the FPT moments for the GBM can be obtained by using the partition polynomial as in (2.11), that is

$$G_n(y; x_1, \ldots, x_n) = \sum_{j=1} y^j B_{n,j}(x_1, \ldots, x_{n-j+1}) \tag{6.69}$$

where $\{B_{n,j}\}$ are the partial exponential Bell polynomials as in (2.1). The latter, combined with the relatively simple expression of the cumulants of T_{gbm} that shall be recalled in the proof of the next proposition, provides a (polynomial) closed-form expression of the moments of T_{gbm} which lets one avoid the evaluation of $\mathcal{K}_\nu$ in (6.66).

Proposition 6.5.2.

$$\mathbb{E}[T_{gbm}^n] = \left(-\frac{2b^2}{a}\right)^n G_n\left(-\frac{a}{b}; \frac{1}{2}, \left(\frac{1}{2}\right)_2, \ldots, \left(\frac{1}{2}\right)_{n-j+1}\right) \tag{6.70}$$

Proof. From the power series expansion of $\ln[M_{T_{gbm}}(\lambda) - 1]$ with $M_{T_{gbm}}(\lambda)$ the moment generating function of T_{gbm}, cumulants $\{c_n[T_{gbm}]\}$ of T_{gbm} result to be (Tweedie, 1957)

$$c_n[T_{gbm}] = (2n-3)!!\frac{b^{2n-1}}{a^{n-1}} \quad \text{for } n \geq 1 \tag{6.71}$$

with $(2n-3)!! = (2n-3)\cdots 5\cdot 3\cdot 1$. Since $(2n-3)!! = (-1)^{n-1}2^n\left(\frac{1}{2}\right)_n$ with $\left(\frac{1}{2}\right)_n = \Pi_{j=0}^{n-1}(\frac{1}{2} - j)$ the lowering factorial, from (6.71) we also have

$$c_n[T_{gbm}] = \left(-\frac{a}{b}\right)\left(-\frac{2b^2}{a}\right)^n\left(\frac{1}{2}\right)_n \quad \text{for } n \geq 1. \tag{6.72}$$

As known, moments $\mathbb{E}[T_{gbm}^n]$ can be recovered from cumulants through the partial exponential Bell polynomials (see Section 2.A in the Appendix of Chapter 2)

$$\mathbb{E}[T^n] = \sum_{j=1}^n B_{n,j}(c_1[T_{gbm}], \ldots, c_{n-j+1}[T_{gbm}]), \qquad \text{for } n \geq 1, \tag{6.73}$$

with $\{B_{n,j}\}$ given in (2.1). The result follows replacing (6.72) in (6.73) and using the well known property $B_{n,j}(pqx_1, pq^2x_2, \ldots) = p^j q^n B_{n,j}(x_1, x_2, \ldots)$ (Comtet, 1970). □

From (6.71), the following recursion holds for the FPT cumulants

$$c_n[T_{gbm}] = \frac{(2n-3)b^2}{a} c_{n-1}[T_{gbm}]$$

starting with $c_1[T_{gbm}] = \mathbb{E}[T_{gbm}] = b$.

The GBM does not have a stationary distribution, thus we cannot apply directly Proposition 6.2.5. However, thanks to the fact that in this case g_{gbm} is known, we can state a simpler version of Proposition 6.2.3 and an analogue of Proposition 6.2.5.

Proposition 6.5.3. *If*

$$I = \int_0^\infty t^{-(\alpha+3)} \exp\left(-At - \frac{a}{t}\right) \mathrm{d}t < \infty, \tag{6.74}$$

with $A = \frac{a}{b^2} - \beta$, *then*

$$g_{gbm} = f_{\alpha,\beta}\left(1 + \sum_{k\geq 1} \mathcal{B}_k^{(\alpha,\beta)} L_k^{(\alpha,\beta)}\right) \text{ with } \mathcal{B}_k^{(\alpha,\beta)} = 1 + \sum_{j=1}^{k} \binom{k}{j} \frac{(-\beta)^j \mathbb{E}[T_{gbm}^j]}{(\alpha+j)_j}. \tag{6.75}$$

Proof. The result is immediate after using Proposition 6.2.3, noting that in this case

$$\int_0^\infty t^{-\alpha} e^{\beta t} g_{gbm}(t)^2 \, \mathrm{d}t = \frac{a\Gamma(\alpha+1)e^{2a/b}}{2\pi\beta^{\alpha+1}} I.$$

□

Corollary 6.5.4. *Condition* (6.74) *is fulfilled if and only if*

$$\beta \leq \frac{1}{\sigma^2}\left(\mu - \frac{\sigma^2}{2}\right)^2. \tag{6.76}$$

Proof. Note that (6.76) is equivalent to have $A = \frac{1}{\sigma^2}\left(\mu - \frac{\sigma^2}{2}\right)^2 - \beta \geq 0$ in (6.74). If $A > 0$ from 3.471 no. 9 in Gradshteyn and Ryzhik (1994), we get

$$\int_0^\infty t^{-(\alpha+3)} \exp\left(-At - \frac{a}{t}\right) \mathrm{d}t = 2\left(\frac{A}{a}\right)^{\frac{2+\alpha}{2}} \mathcal{K}_{\alpha+2}(2\sqrt{aA}) < \infty$$

where $\mathcal{K}_\nu$ is the modified Bessel function of second type. If $A=0$, by a suitable change of variable and using 3.478 no. 1 in Gradshteyn and Ryzhik (1994), we have

$$\int_0^\infty t^{-(\alpha+3)}\exp\left(-At-\frac{a}{t}\right)\mathrm{d}t=\int_0^\infty x^{(\alpha+2)-1}e^{-ax}\mathrm{d}x=a^{-(\alpha+2)}\Gamma(\alpha+2)<\infty.$$

If $A<0$, the integrand function in I grows with t and the condition (6.74) is not fulfilled. □

The following theorem provides a more precise bound on $\left\|\frac{g_{gbm}-\hat{g}_n}{f_{\alpha,\beta}}\right\|_{\alpha,\beta}$ exploiting the knowledge of g_{gbm}.

Theorem 6.5.1. *If* $\beta<\frac{1}{\sigma^2}\left(\mu-\frac{\sigma^2}{2}\right)^2$ *and* $k\in\mathbb{N}$, *then there exists a constant* $C_k>0$ *such that*

$$\left\|\frac{g_{gbm}-\hat{g}_n}{f_{\alpha,\beta}}\right\|_{\alpha,\beta}\leq C_k\left(\frac{1}{\sqrt{n}}\right)^k \quad \text{for all } n>k \tag{6.77}$$

with $C_k=O(k^k)$.

Proof. Set $\tilde{a}=\frac{a}{2}$ and $\tilde{A}=\frac{a}{2b^2}-\beta$ in order to write

$$\frac{g_{gbm}(t)}{f_{\alpha,\beta}(t)}=D\exp\left(-\tilde{A}t-\frac{\tilde{a}}{t}\right)t^{-\frac{3}{2}-\alpha} \quad \text{with} \quad D=\sqrt{\frac{\tilde{a}}{\pi}}\frac{e^{a/b}\Gamma(\alpha+1)}{\beta^{\alpha+1}}.$$

According to Theorem 6.2.5 in Funaro (1992), if for a fixed $k\in\mathbb{N}$

$$t^{m/2}\frac{\mathrm{d}^m}{\mathrm{d}t^m}\left[\frac{g_{gbm}(t)}{f_{\alpha,\beta}(t)}\right]\in\mathcal{L}^2_{f_{\alpha,\beta}(t)}(0,\infty) \quad \text{for} \quad 0\leq m\leq k \tag{6.78}$$

then there exists a constant $C>0$ such that

$$\left\|\frac{g_{gbm}-\hat{g}_n}{f_{\alpha,\beta}}\right\|_{\alpha,\beta}\leq C\left(\frac{1}{\sqrt{n}}\right)^k\left\|t^{k/2}\frac{\mathrm{d}^k}{\mathrm{d}t^k}\left[\frac{g_{gbm}(t)}{f_{\alpha,\beta}(t)}\right]\right\|_{\alpha,\beta} \quad \text{for all } n>k. \tag{6.79}$$

By recursion, we have

$$\frac{\mathrm{d}^m}{\mathrm{d}t^m}\left[\frac{g_{gbm}(t)}{f_{\alpha,\beta}(t)}\right]=\frac{(-1)^m}{2^m}D\exp\left(-\tilde{A}t-\frac{\tilde{a}}{t}\right)t^{-2m-\frac{3}{2}-\alpha}q_{2m}(t) \quad \text{for } m\geq 0 \tag{6.80}$$

where $q_{2m}(t)$ is a polynomial of degree $2m$ such that

$$q_{2m}(t)=\sum_{j=0}^{2k}c_{2m,j}t^j=q_{2m-2}(t)[2\tilde{A}t^2+(4m-1+2\alpha)t-2\tilde{a}]-2t^2\frac{\mathrm{d}}{\mathrm{d}t}q_{2m-2}(t) \tag{6.81}$$

with $q_0(t)=1$ and $c_{2m,2m}=(2\tilde{A})^m\neq 0, c_{2m,0}=(-2\tilde{a})^m$.

Now set

$$\tilde{q}_{4m}(t)=[q_{2m}(t)]^2=\sum_{j=0}^{4m}\tilde{c}_{4m,j}t^j \quad \text{and} \quad I_m=\int_0^\infty t^m\left(\frac{\mathrm{d}^m}{\mathrm{d}t^m}\left[\frac{g_{gbm}(t)}{f_{\alpha,\beta}(t)}\right]\right)^2 f_{\alpha,\beta}(t)\mathrm{d}t.$$

Fix an integer $0\leq k<n$. Since $2\tilde{A}+\beta=\frac{c_1[T]}{c_2[T]}-\beta>0$ we have $I_m<\infty$ for all $m\geq 0$ and in particular condition (6.78) holds for $0\leq m\leq k$. Indeed using the modified Bessel function of second type (6.67) and the integral 3.471 no. 9 in Gradshteyn and Ryzhik (1994), we have

$$\begin{aligned} I_m&=\frac{D^2}{2^{2m}}\sum_{j=0}^{4m}\tilde{c}_{4m,j}\int_0^\infty \exp\left(-(2\tilde{A}+\beta)t-\frac{a}{t}\right)t^{(-3m-2-\alpha+j)-1}\mathrm{d}t \qquad (6.82)\\ &=\frac{D^2}{2^{2m-1}}\sum_{j=0}^{4m}\tilde{c}_{4m,j}\mathcal{K}_{3m-j+(\alpha+2)}\left(\sqrt{4a(2\tilde{A}+\beta)}\right)\left[\sqrt{\frac{a}{2\tilde{A}+\beta}}\right]^{-3m-2-\alpha+j}<\infty. \end{aligned}$$

Eq. (6.77) follows from (6.79) setting $m=k$ and

$$C_k=CI_k\propto\frac{1}{2^{2k}}\sum_{j=0}^{4k}\tilde{c}_{4k,j}\mathcal{K}_{3k-j+(\alpha+2)}\left(\sqrt{4a(2\tilde{A}+\beta)}\right)\left[\sqrt{\frac{a}{2\tilde{A}+\beta}}\right]^{-3k-2-\alpha+j}. \tag{6.83}$$

In (6.83), note that $3k-j+\alpha+2>0$ for $j=0,\ldots,3k+1$ as $\alpha+1>0$. For $3k+2\leq j\leq 4k$, the order of the modified Bessel function involved in C_k might be positive, depending on the magnitude of α. Let us first suppose $\alpha>k-2$ such that $3k-j+\alpha+2>0$ for all $j=0,\ldots,4k$. As for $\nu\to\infty$ one has

$$\mathcal{K}_\nu(z)\sim\sqrt{\frac{\pi}{2\nu}}\left(\frac{2\nu}{ez}\right)^\nu \tag{6.84}$$

then

$$C_k\propto\frac{1}{2^{2k}}\sum_{j=0}^{4k}\frac{\tilde{c}_{4k,j}}{\sqrt{3k-j+\alpha+2}}\left(\frac{3k-j+\alpha+2}{ea}\right)^{3k-j+\alpha+2}. \tag{6.85}$$

When k grows, the dominant term in (6.85) is for $j=0$, and the result follows.

If $\alpha < k-2$, then C_k might be splitted in $C_k \propto C_{k,1} + C_{k,2}$ with

$$\begin{aligned} C_{k,1} &= \frac{1}{2^{2k}} \sum_{j=0}^{k^*} \tilde{c}_{4k,j} \mathcal{K}_{3k-j+(\alpha+2)}\left(\sqrt{4a(2\tilde{A}+\beta)}\right)\left[\sqrt{\frac{a}{2\tilde{A}+\beta}}\right]^{-3k-2-\alpha+j} \\ C_{k,2} &= \frac{1}{2^{2k}} \sum_{j=k^*+1}^{4k} \tilde{c}_{4k,j} \mathcal{K}_{3k-j+(\alpha+2)}\left(\sqrt{4a(2\tilde{A}+\beta)}\right)\left[\sqrt{\frac{a}{2\tilde{A}+\beta}}\right]^{-3k-2-\alpha+j} \end{aligned}$$

where k^* is such that $3k-k^*+\alpha+2>0$ and $3k-k^*+\alpha+1<0$. For $\alpha \in (-1,0)$, we have $3k-j+(\alpha+2)<0$ for $3k+2 \leq j \leq 4k$ and $C_{k,2}$ includes the maximum number of terms, that is

$$\begin{aligned} C_{k,2} &= \frac{1}{2^{2k}} \sum_{j=3k+1}^{4k} \tilde{c}_{4k,j} \mathcal{K}_{j-3k-(\alpha+2)}\left(\sqrt{4a(2\tilde{A}+\beta)}\right)\left[\sqrt{\frac{a}{2\tilde{A}+\beta}}\right]^{-3k-2-\alpha+j} \\ &\sim \frac{1}{2^{2k}} \sum_{j=3k+1}^{4k} \frac{\tilde{c}_{4k,j}}{\sqrt{j-3k-\alpha-2}} \left(\frac{j-3k-\alpha-2}{e(2\tilde{A}+\beta)}\right)^{j-3k-\alpha-2}. \end{aligned} \tag{6.86}$$

The dominant term in (6.86) is for $j=4k$ and the asymptotic behaviour of C_k is still of order k^k. □

Observe that for $k=2$, from (6.77) we recover a similar result to the one obtained in Theorem 6.3.1.

Remark 6.5.5. *Note that if* $\beta = \frac{1}{\sigma^2}\left(\mu - \frac{\sigma^2}{2}\right)^2$ *the integral* I_m *in* (6.82) *converges if and only if* $\alpha > m-2$. *Indeed in such a case* $2\tilde{A}+\beta=0$ *and* I_m *in* (6.82) *reduces to*

$$I_m = \sum_{j=0}^{4m} \tilde{c}_{4k,j} \int_0^\infty \exp(-ay)\, y^{(3m+2+\alpha-j)-1}\, \mathrm{d}t. \tag{6.87}$$

The integral on the rhs of (6.87) *is convergent if and only if* $3m+2+\alpha-j>0$ *for all* $j=0,\ldots,4m$, *that is if and only if* $\alpha > m-2$. *Therefore* (6.77) *still holds with* $k<\alpha+2$. *In such a case we have*

$$\begin{aligned} C_k &= CI_k \propto \frac{1}{2^{2k}} \sum_{j=0}^{4k} \tilde{c}_{4k,j} a^{3k+\alpha+2-j} \Gamma(3k-j+\alpha+2) \\ &\propto \frac{1}{2^{2k}} \sum_{j=0}^{4k} \tilde{c}_{4k,j} a^{3k+\alpha+2-j} \frac{(3k-j)^{3k-j+\alpha+3/2}}{e^{3k-j}} \end{aligned} \tag{6.88}$$

as $\Gamma(z+b) \sim \sqrt{2\pi} e^{-z} z^{z+b-1/2}$. *As the leading term in* (6.88) *is for* $j=0$, *we still have* $C_k = O(k^k)$.

6.6 Numerical Results

In this section, the functions $\hat{G}_n$ in (6.41) and $\hat{g}_n$ in (6.37) are used to approximate the FPT cdf G_{cir} (resp. G_{gbm}) and the FPT pdf g_{cir} (resp. g_{gbm}) of a CIR process as in (6.55) (resp. GBM process as in (6.62)) with the corrections suggested in subsections 6.4.4 and 6.4.5, given the series representation of g_{cir} in (6.61) (resp. of g_{gbm} in (6.75)). In each of the two cases, we will stress out different strengths and weaknesses of the proposed method. For instance, while in the GBM case we will try to highlight how differently shaped reference and objective density may impact the approximation, focusing on the pdf, in the CIR case we will show how the standardization explained in the second part of Section 6.4.3 may be an useful tool to deal with the same issue and we will also deal with the cdf.

Before examining specific numerical examples, in the following subsection we provide some details on how the goodness of approximation can be assessed.

6.6.1 Comparisons With Alternative Approximation Methods

GBM Since g_{gbm} is known as explained in Section 6.5.2, it was possible to test the accuracy and efficiency of the approximation by comparing $\hat{g}_n$ in (6.37) with the true FPT pdf g_{gbm} in (6.65) using the following two approaches:

(a) a graphical approach by simply comparing the plot of $\hat{g}_n$ and g_{gbm},

(b) a quantitative approach by computing $|g_{gbm}(t) - \hat{g}_n(t)|$ for $t > 0$.

CIR As a closed form of the FPT pdf g_{cir} and cdf G_{cir} for a CIR process is in general not available, the validity of proposed approximations needs to be evaluated by comparing it with other estimates obtained using different techniques.

As mentioned in the introduction, for one-dimensional diffusion processes, the FPT pdf through a time-dependent boundary can be recovered as the solution of a Volterra integral equation of the second kind (Buonocore et al., 1987). With appropriate numerical methods for approximating the integral, a discrete numerical evaluation of this solution can be performed to obtain an approximation of g_{cir}. However, during the implementation of these methods, we encountered significant challenges. When the coefficient of variation is large, this procedure was prone to overflow failures, potentially

leading to highly inaccurate results. This issue is likely exacerbated by the inevitable propagation of numerical errors, as in this procedure the new approximated values of the FPT pdf rely on those computed in previous steps. Nevertheless, even for FPT pdfs with a small coefficient of variation, we faced several implementation difficulties. Indeed, in any case, the presence of the Bessel function in the transition pdf of the process complicates the numerical evaluation of its derivatives. Moreover, similar issues arose when using the `R` package `fptdApprox` (Román-Román et al., 2023). Despite these challenges, the numerical results, excluding these pathological cases, are comparable to those obtained through classical Monte Carlo methods. Consequently, we have opted to use Monte Carlo methods as alternative tools to assess the accuracy of our approximations.

In this context, with Monte Carlo method we refer to simulating sample paths of the CIR process and look for their FPTs over the given threshold. Within this procedure, we first implemented the Milstein algorithm (Platen and Kloeden, 1992) to sample the paths of the CIR process, which generates a trajectory by a suitable discretisation of the stochastic differential equation (6.55). As it is well-known, this procedure may easily become time-consuming. For instance, to get a FPT sample of size $N>0$, it is clear that at least N different trajectories of the CIR process must be generated. Indeed, not all the generated trajectories may reach the threshold in a "reasonable" time. This also implies that the FPT pdf can be underestimated if a finite time interval has been set for the simulation, as usually happens. Moreover, the fixed time step determines how accurately the dynamics can be described and the computational time increases as this time step gets smaller. Therefore, it is necessary to choose a very small step size and simulate many trajectories of the CIR process to obtain significant results. This can become computationally intense, especially if the coefficient of variation of T is large, as a consequence of the likelihood of a large time span length over which the trajectories must be simulated.

Similar problems may arise when sampling trajectories of the CIR process using its transition pdf, a non-central chi-square distribution (see, e.g., Feller, 1951). Set $\Delta t>0$. In such a case, starting from $Y_0=y_0$, an instance of $Y_{\Delta t}$ is generated from the conditional distribution of $Y_{\Delta t}|Y_0=y_0$, an instance of $Y_{2\Delta t}$ from the conditional distribution of $Y_{2\Delta t}|Y_{\Delta t}$ and so on. The results obtained by the two Monte Carlo methods are comparable. However we have used the Milstein algorithm, since the computational time is lower in all the cases examined.

Once a sample of independent and identically distributed (i.i.d.) FPTs, denoted by $\mathcal{T}=\{T_1,\ldots,T_N\}$, $N\in\mathbb{N}^+$, has been collected, the empirical cdf (ecdf) can be used to obtain a sufficiently reliable estimate of the cdf shape. Among the available nonparametric methods for estimating the pdf, the histogram is the most widely used. However, its limitations, such as the significant dependence on bandwidth choice, are well-documented in the literature. Kernel density estimators (KDEs) are often cited as simple alternatives to histograms. Nevertheless, if the unknown pdf is supported on the positive half-line and lacks smoothness at the origin, the kernel method may be inefficient (Hall, 1980). For instance, the KDE might obscure the mode of the unknown pdf by assigning positive mass to negative values (see Fig. 6.6.9). To achieve the smoothness characteristic of KDEs while obtaining an adequate estimated density with support $(0,\infty)$, an estimator based on an orthogonal series can be highly competitive (Efromovich, 2010). This latter approach turns out to be an intriguing application of the approximation $\hat{g}_n$ itself and thus will be discussed in more detail in Section 6.7.2. Consequently, to evaluate the effectiveness of these approximations, histograms and classical KDEs have been employed despite their known shortcomings.

6.6.2 Numerical Examples for the GBM

Recall that for the GBM the FPT pdf is known and we have

$$g_{gbm}(t)=\sqrt{\frac{a}{2\pi t^3}}\exp\left(-\frac{a(t-b)^2}{2b^2t}\right),\quad t>0. \tag{6.89}$$

where

$$a=\frac{(\ln S-\ln y_0)^2}{\sigma^2}>0\quad\text{and}\quad b=\frac{\ln S-\ln y_0}{\mu-\frac{\sigma^2}{2}}>0.$$

We consider three different cases of the GBM defined as in (6.62). Throughout the whole section we fix the boundary $S=10$ and the starting position $y_0=1$.

case A_{gbm}: $\mu=4$ and $\sigma=1.4$,

case B_{gbm}: $\mu=2.2$ and $\sigma=1.4$,

case C_{gbm}: $\mu=1.4$ and $\sigma=1.4$.

Note that as shown in Fig. 6.6.1, as the parameters change, the FPT pdf g_{gbm} and the reference pdf $f_{\alpha,\beta}$ can be remarkably different. From now onward, whenever mentioned, the stopping criteria in (6.47) are employed

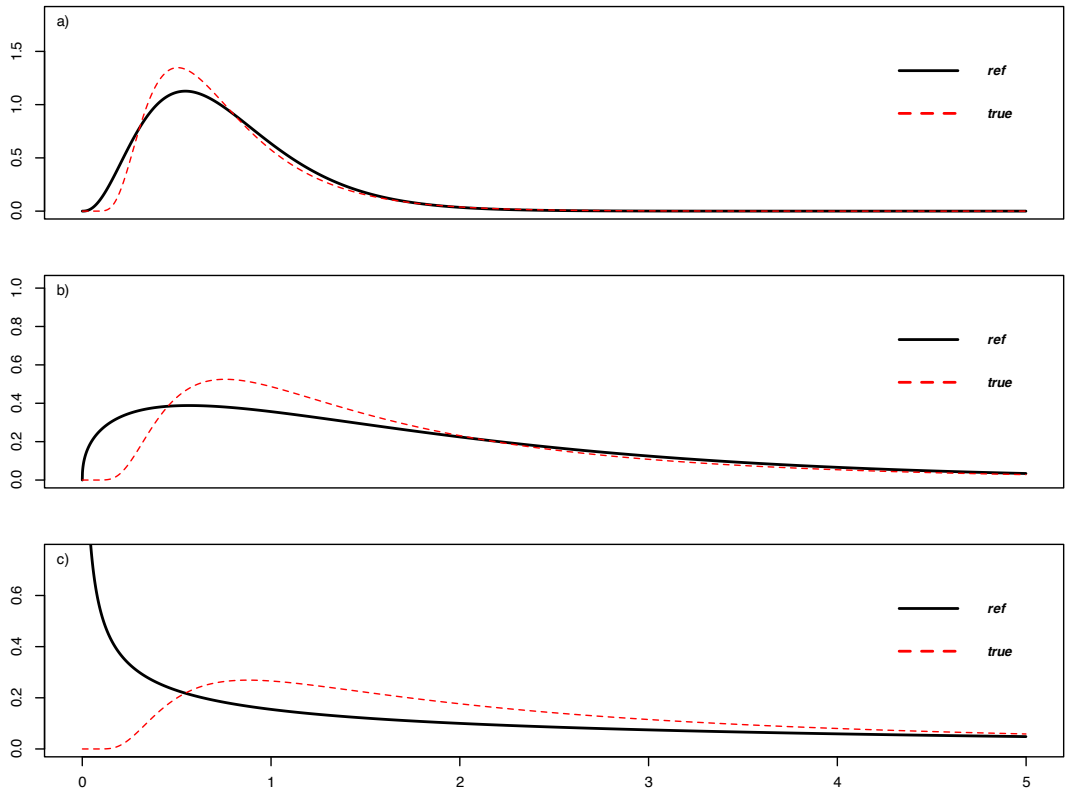

Figure 6.6.1: Plots of the FPT pdf g_{gbm} (red dashed line) and the corresponding reference pdf $f_{\alpha,\beta}$ (black solid line), in a) for case A_{gbm} with $\alpha = 2.54$ and $\beta = 4.65$, in b) for case B_{gbm} with $\alpha = 0.43$ and $\beta = 0.76$ and in c) for case C_{gbm} with $\alpha = -0.05$ and $\beta = 0.09$.

with $\varepsilon = 10^{-14}$. Furthermore, for all the three considered cases the parameters α and β of $f_{\alpha,\beta}$ in (6.37) have been chosen according to Section 6.4.3. In Figs. 6.6.2, 6.6.3 and 6.6.4 we have plotted the polynomial approximation $\hat{g}_n$ (blue solid line) and the true density g_{gbm} (red dashed line) for the three mentioned cases using four different orders of approximation each time.

By comparing the three figures we can make the following remarks.

(a) In case A_{gbm}, where the reference pdf $f_{\alpha,\beta}$ and the FPT pdf g_{gbm} have a similar behaviour, a low degree n already guarantees a good approximation,

(b) as μ diminishes and the computed α decreases, indicating an increased c_v for T_{gbm} as explained in Section 6.4.3, the reference pdf loses its typical bell shape and deviates further away from the FPT pdf g_{gbm},

(c) in all the considered cases the goodness of the approximation increases

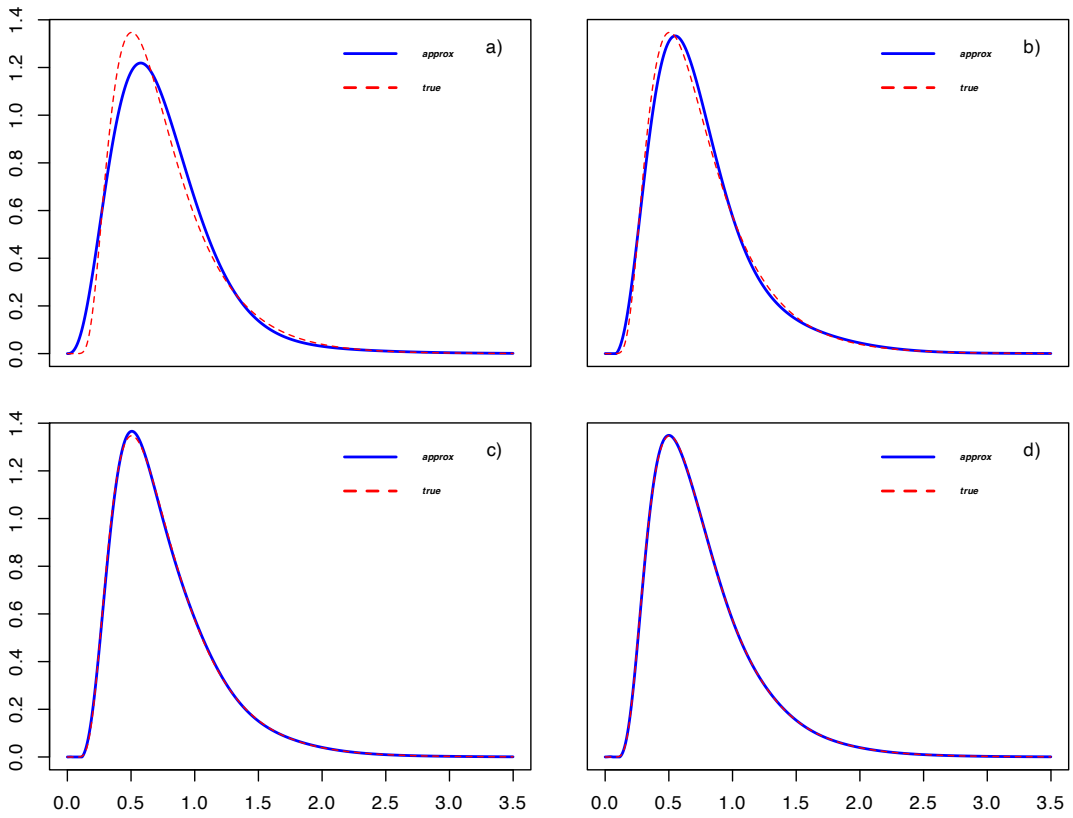

Figure 6.6.2: Plots of the polynomial approximation $\hat{g}_n$ (blue solid line) and of the true density g_{gbm} (red dashed line) in case A_{gbm} with $S=10$, $y_0=1$, $\mu=4$ and $\sigma=1.4$ for $n=3$ in *a)* $n=5$ in *b)* $n=16$ in *c)* and $n=30$ in *d)*, where the last n is the minimum integer s.t. conditions (6.47) do not hold.

with n as long as conditions (6.47) are satisfied,

(d) compared with cases A_{gbm} and B_{gbm}, case C_{gbm} proves to be the hardest to tackle; in fact even for $n=36$ the approximant $\hat{g}_n$ does not match the peak of g_{gbm}. A cause of this could be the fact that, for $\alpha<0$, the mode of the reference pdf $f_{\alpha,\beta}$ is not well defined and thus the initial approximation is very different from g_{gbm}. Moreover, for the choice of parameters $\sigma^2=1.4$ and $\mu=1.4$ the stochastic component of the dynamics prevails over the deterministic one, resulting in a flatter FPT pdf g_{gbm}. To get a better approximation we should consider even higher values of n, but this would be prevented by the increasing numerical errors, unless one employs an increased numerical precision (an example of this will be given in the CIR scenario).

In Fig. 6.6.5 we have plotted the absolute error between the true FPT pdf

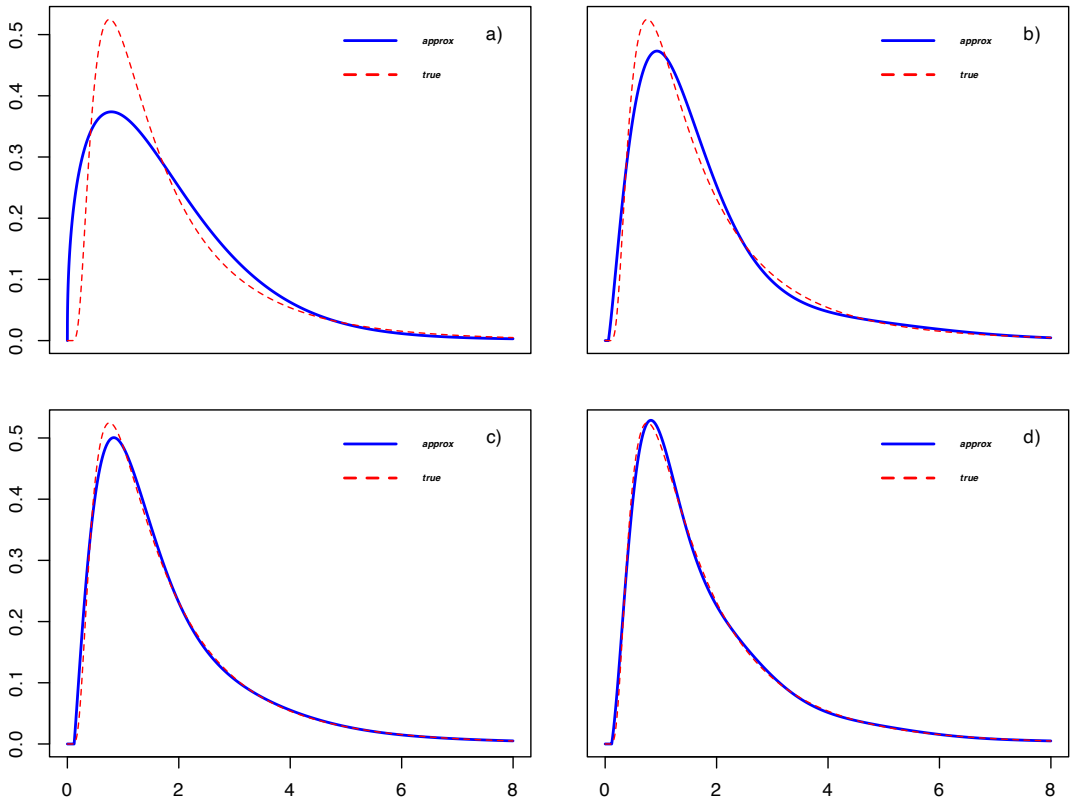

Figure 6.6.3: Plots of the polynomial approximation $\hat{g}_n$ (blue solid line) and of the true density g_{gbm} (red dashed line) in case B_{gbm} with $S=10$, $y_0=1$, $\mu=2.2$ and $\sigma=1.4$ for $n=3$ in *a)* $n=8$ in *b)* $n=16$ in *c)* and $n=29$ in *d)*, where the last n is the minimum integer s.t. conditions (6.47) do not hold.

g_{gbm} and the approximated $\hat{g}_n$, that is the function $|g_{gbm}(t)-\hat{g}_n(t)|$ for $t>0$, for the three considered cases and for the smallest n such that conditions (6.47) do not hold anymore.

6.6.3 Numerical Examples for the CIR

Before proceeding further, note that, unlike for the GBM, in the case of the CIR process an additional issue arises in computing (6.25). As they are not available in a closed form, the moments of T_{cir} are calculated from its cumulants (6.57) using the recursion (6.60). Because of the series involved in (6.57), a truncation order m should be chosen before computing the moments through (6.60). Here, a standard approach has been used, which involves computing partial sums of the series as long as their difference exceeds an input tolerance. To analyse the efficiency and the usefulness of the proposed

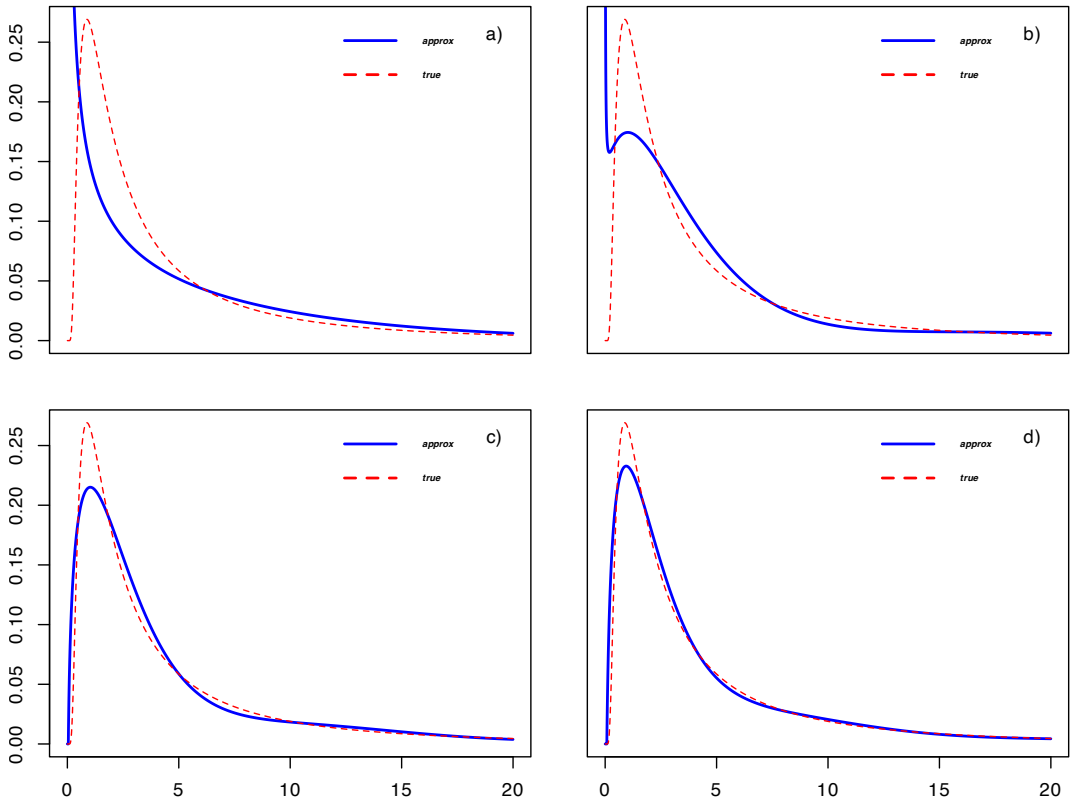

Figure 6.6.4: Plots of the polynomial approximation $\hat{g}_n$ (blue solid line) and of the true density g (red dashed line) in case C with $S=10$, $y_0=1$, $\mu=1.4$ and $\sigma=1.4$ for $n=3$ in *a)* $n=15$ in *b)* $n=25$ in *c)* and $n=36$ in *d)*, where the last n is the minimum integer s.t. conditions (6.47) do not hold.

method in the CIR process case we consider three different sets of parameters for (6.55).

case A_{cir}: $y_0=0.2$, $\mu=0.9$, $S=1$, $c=0$, $\sigma=1.2$ and $\tau=2/3$,

case B_{cir}: $y_0=0.01$, $\mu=0.005$, $S=0.02$, $c=0$, $\sigma=0.1$ and $\tau=0.25$,

case C_{cir}: $y_0=0$, $\mu=3$, $S=10$, $c=-10$, $\sigma=1.2$ and $\tau=0.2$.

We stress out again that in these three cases a closed form of the FPT pdfs g_{cir} and cdfs G_{cir} is not available. Nevertheless, by observing plots of empirical FPT cdfs in Fig. 6.6.6, we can infer that the different considered parameters result in FPT pdfs g_{cir} and cdfs G_{cir} with various forms and statistical properties.

The empirical cdfs have been constructed after using the Milstein method to simulate a sample of 10^4 FPTs for each case, as said in Section 6.6.1.

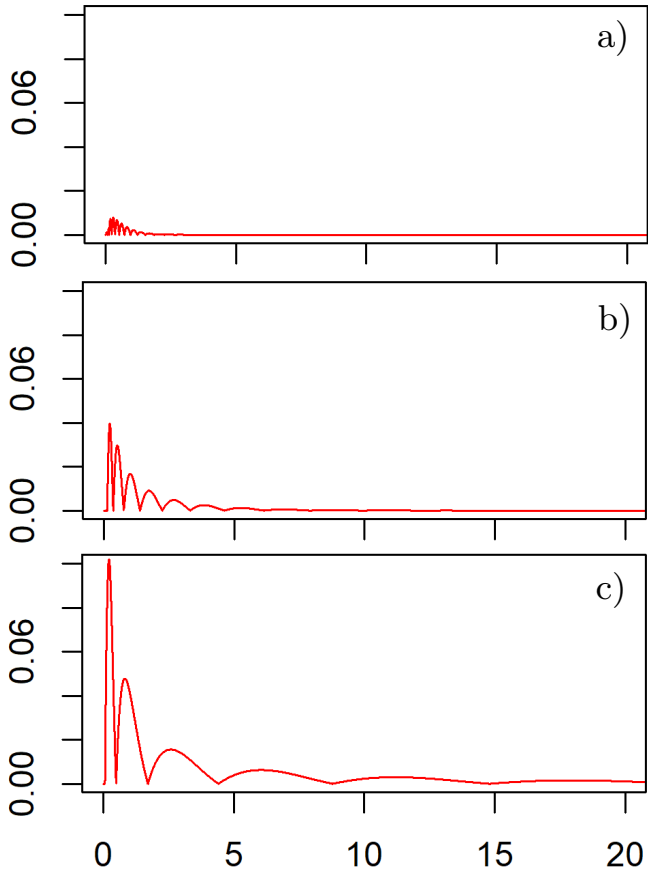

Figure 6.6.5: Plots of the absolute error $|g_{gbm} - \hat{g}_n|$ between the true FPT pdf g_{gbm} and the approximation $\hat{g}_n$ for the smallest n s.t. conditions (6.47) do not hold in case A_{gbm} with $S = 10$, $y = 1$, $\mu = 4$ and $\sigma = 1.4$ in *a)*, case B_{gbm} with $S = 10$, $y_0 = 1$, $\mu = 2.2$ and $\sigma = 1.4$ in *b)* and case C_{gbm} with $S = 10$, $y = 1$, $\mu = 1.4$ and $\sigma = 1.4$ in *c)*.

As we continue, these three samples are denoted by $\mathcal{T}_{A_{cir}}$, $\mathcal{T}_{B_{cir}}$ and $\mathcal{T}_{C_{cir}}$ respectively.

For each case, as further information on the shape of the unknown FPT pdf g_{cir}, we have computed the FPT dispersion coefficients as given in Section 6.4.3. The coefficient of variation is computed using the theoretical FPT mean and variance, since they are known from Section 6.5.1. The estimated relative entropy based dispersion coefficient $\hat{c}_h$ is computed with the Vasicek estimator (Kostal and Pokora, 2012) using the samples $\mathcal{T}_{A_{cir}}$, $\mathcal{T}_{B_{cir}}$ and $\mathcal{T}_{C_{cir}}$ respectively. The results are given in Table 6.6.1.

As in the previous subsection, note that, whenever mentioned, the stopping criteria in (6.47) are employed with $\varepsilon = 10^{-14}$ and recall that parameters α and β in (6.37) have been chosen according to the first part of Section

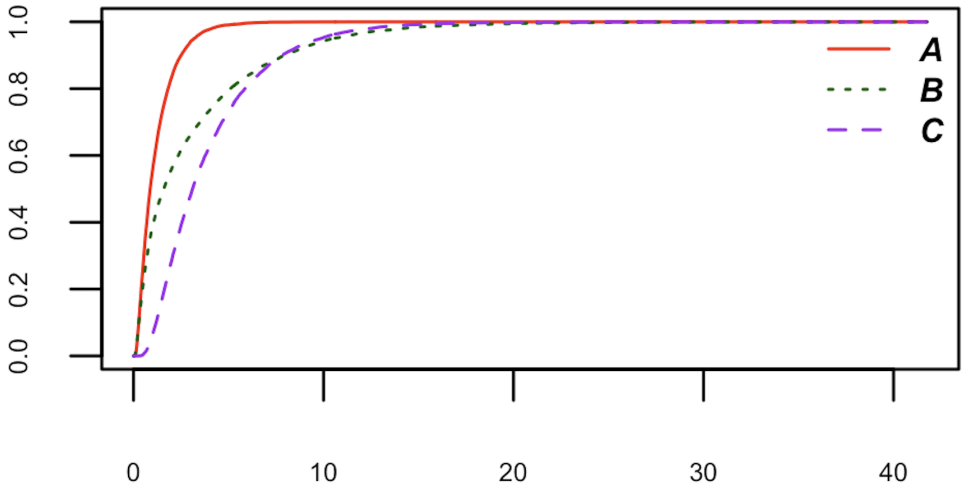

Figure 6.6.6: Plots of the empirical (not standardized) FPT cdfs for cases A_{cir} (in solid red), B_{cir} (in dash green), and C_{cir} (in dashed purple).

Table 6.6.1: FPT dispersion indexes c_v and $\hat{c}_h$ togheter with mean, standard deviation, skewness and kurtosis for the cases A, B and C.

	c_v	$\hat{c}_h$	$\mathbb{E}[T]$	$\sqrt{\mathrm{Var}[T]}$	γ_1	κ_1
A	0.855	0.909	1.16	0.984	1.968	5.9862
B	1.231	0.916	2.991	13.56	2.39	8.118
C	0.765	0.855	3.937	9.084	1.905	5.572

6.4.3. Unlike in the previous section, in this case we have employed the standardization procedure detailed in the second part of Section 6.4.3. Indeed, Figs. 6.6.7, 6.6.8, 6.6.9 and 6.6.10 refer to the standardized FPT rv $\tilde{T}$ in (6.49). In Fig. 6.6.7 we have plotted the empirical cdfs G_e, corrisponding to the samples $\mathcal{T}_{A_{cir}}$, $\mathcal{T}_{B_{cir}}$ and $\mathcal{T}_{C_{cir}}$, together with the approximated cdfs $\tilde{G}_n$, as obtained using (6.51) and the corrections described in Section 6.4.4, for $\Delta t = 10^{-5}$, normalized by the corresponding standard deviations (see Table 6.6.1). Moreover each figure displays the maximum absolute error defined as $\varepsilon_a = \max_{t\geq 0} |\tilde{G}_n(t) - G_e(t)|$. Figs. 6.6.8, 6.6.9 and 6.6.10, correspond to cases A_{cir}, B_{cir} and C_{cir} respectively. To emphasise the differences in density estimations, as discussed in Section 6.6.1, we have plotted in these figures a classical KDE and the standardised approximated pdf $\tilde{g}_n$, corrected according to Section 6.4.5. Since the three considered instances correspond to FPT pdfs with different shapes, these comparisons should help in providing a further comprehensive picture of the strengths and weaknesses of the proposed method in addition to the already treated GBM case. They are discussed in

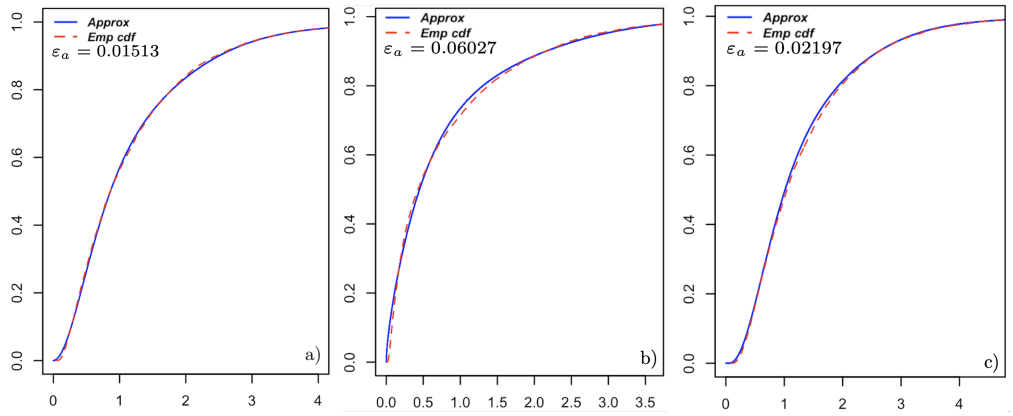

Figure 6.6.7: Plots of the approximated $\tilde{G}_n$ (in solid blue) and of the empirical cdf (in dashed red) together with the corresponding maximum absolute error ε_a. The plots refer: to case A_{cir} in a) with $n=10$, $\alpha=0.367$ and $\beta=1.17$; to case B_{cir} in b) with $n=10$, $\alpha=-0.34$ and $\beta=0.812$; to case C_{cir} in c) with $n=9$, $\alpha=0.7$ and $\beta=1.306$. Note that $\tilde{G}_n(t)$ is obtained using the stopping criteria (6.47) and corrected according to Section 6.4.4 while the empirical cdf is obtained from the standardized samples $\mathcal{T}_{A_{cir}}, \mathcal{T}_{B_{cir}}$ and $\mathcal{T}_{C_{cir}}$ respectively.

the following.

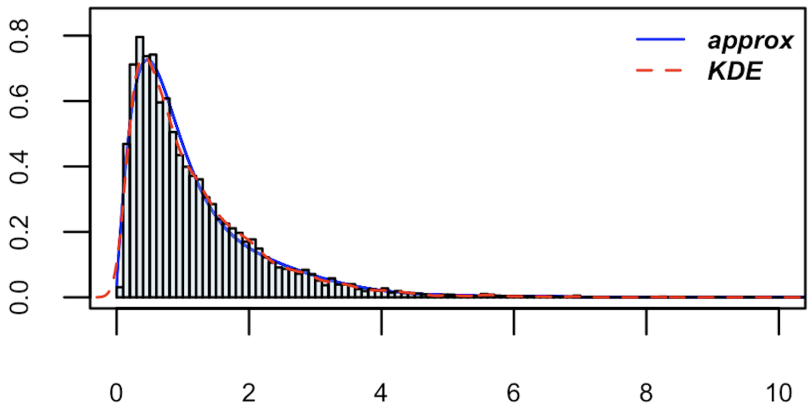

Figure 6.6.8: plot of $\tilde{g}_n$ (in solid blue) in case A_{cir} with $n=10$, $\alpha=0.367$ and $\beta=1.17$, obtained with the stopping criteria (6.47) and corrected to ensure positivity as in (6.54), together with a KDE (in dashed red) and a histogram both computed with the standardized sample $\mathcal{T}_{A_{cir}}$.

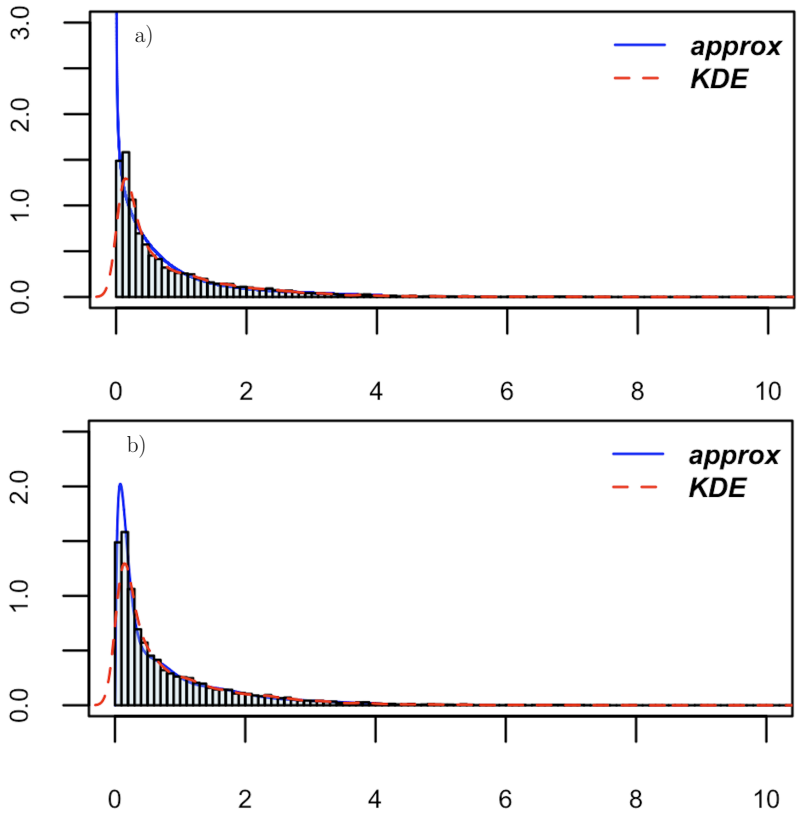

Figure 6.6.9: In a), plot of $\tilde{g}_n$ (in solid blue) in case B_{cir} with $n=10$, $\alpha=-0.34$ and $\beta=0.812$, obtained with the stopping criteria (6.47) together with a KDE (in dashed red) and a histogram both computed with the standardized sample $\mathcal{T}_{B_{cir}}$. In b), a plot of $\tilde{g}_n$ (in solid blue) in case B with $n=55$, $\alpha=-0.34$ and $\beta=0.812$, obtained without any stopping criterion, increasing the numerical precision.

Case A_{cir}

Among the three instances taken into consideration, case A_{cir} has the lightest tail (see Fig. 6.6.6). Intuitively, this should result in an accurate approximation even with a small value of n. Indeed the maximum absolute error between the empirical and approximated cdf is low as shown in Fig. 6.6.7-a). It is assumed at $t=0.134$. For the pdf, the suggested approach combined with the stopping criteria (6.47) yield an approximation $\tilde{g}_n$ with $n=10$, $\alpha=0.367$ and $\beta=1.17$. In this case, $\tilde{g}_n$ is negative on a small interval after the mode, as shown in Fig. 6.4.1. Therefore, the correction outlined in Section 6.4.5 has been implemented. This corrected approximation is plotted in Fig. 6.6.8. An estimated pdf obtained with a classical KDE and a histogram, both computed on the standardized sample $\mathcal{T}_{A_{cir}}$, are shown on the same

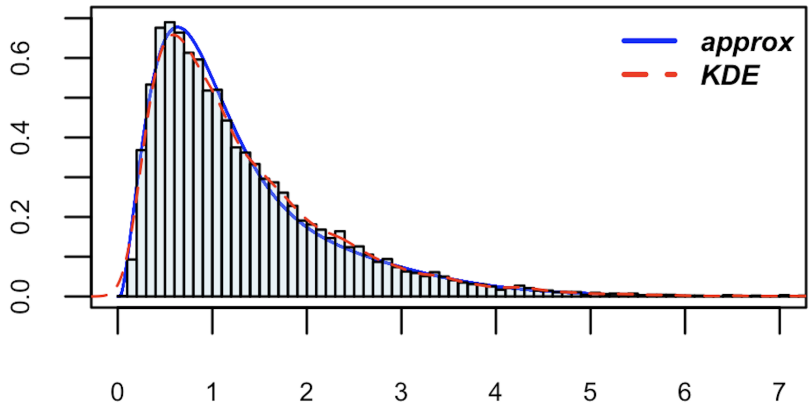

Figure 6.6.10: plot of $\tilde{g}_n$ (in solid blue) for case C_{cir} with $n=9$, $\alpha=0.7$ and $\beta=1.306$, obtained with the stopping criteria (6.47) and corrected to ensure positivity as in (6.53), together with a KDE (in dashed red) and a histogram both computed with the sample $\mathcal{T}_{C_{cir}}$, after the latter has been standardized.

figure.

Case B_{cir}

In this case, different considerations are required. From Table 6.6.1, the FPT rv T_{cir} has a coefficient of variation larger than 1. This makes the approximation more challenging because the distribution seems to have a significant tail (see Fig. 6.6.6). The stopping criteria (6.47) yield an approximated FPT pdf $\tilde{g}_n$ with $n=10$, $\alpha=-0.34$ and $\beta=0.812$, which can be seen in Fig. 6.6.9 a). When compared with a KDE and a histogram, it seems that $n=10$ is not enough to recover $\tilde{g}_{cir}$. A relatively low n is caused by the first conservative stopping criterion in (6.47), proposed to avoid numerical instability, which, in this trickier case, seems to arise somehow quickly. To underline that the behavior of $\tilde{g}_n$ is significantly influenced by the numerical precision, Fig. 6.6.9-b) shows the remarkably good approximation obtained when computing $\tilde{g}_n$ with a very high numerical precision, allowing to push the iterative procedure up to $n=55$. The numerical precision has been raised using the `R`-package `bignum` (Hall, 2023). Still, it is worth to mention that the approximation $\tilde{g}_n$ in Fig. 6.6.9-a) yields a satisfactory result for the tail of $\tilde{g}_{cir}$. Note that the maximum absolute error between empirical and approximated cdf, assumed at $t=0.042$, is higher compared to case A (see Fig. 6.6.7-b)).

Case C_{cir}

In this case, the FPT pdf has a coefficient of variation less than 1 along with a tail whose heaviness lies between cases A_{cir} and B_{cir}, as can be seen in Fig. 6.6.6 and Table 6.6.1. As in case A_{cir}, this value of the coefficient of variation should intuitively ensure a good approximation. The suggested approach yields an approximated FPT pdf $\tilde{g}_n$ with $n=9$, $\alpha=0.7$ and $\beta=1.306$. In this case, $\tilde{g}_n$ is negative on a small interval before the mode and close to the origin (see Fig. 6.4.1). Therefore, also in this case, the correction of $\tilde{g}_n$ described in Section 6.4.5 has been used. The result is shown in Fig. 6.6.10. Furthemore, as expected, the maximum absolute error between the empirical and approximated cdfs (see Fig. 6.6.7-c)) is low. It is assumed at $t=1.365$.

6.7 Applications

6.7.1 An Acceptance-Rejection Type Algorithm

In this section we delineate an application of the polynomial FPT pdf approximation described previously. Let $n\in\mathbb{N}$. It consists in a modified acceptance rejection algorithm which exploits the form of the approximant $\hat{g}_n$ in (6.37).

The acceptance-rejection method is a well known technique for sampling from a distribution that is unknown or difficult to simulate through the classical inverse transformation. Under such circumstances, samples are collected from an auxiliary density by making use of a probability of acceptance (Chib and Greenberg, 1995; Casella et al., 2004).

More in details, let Z be an absolutely continuous rv with pdf π. Suppose there exists a constant $M>0$ and a pdf q such that

$$q(x)>0 \quad \text{and} \quad \frac{\pi(x)}{q(x)}\leq M, \quad \forall x\in \operatorname{supp}(\pi). \tag{6.90}$$

The acceptance-rejection method exploits the condition in (6.90) to sample from the support of Z. Clearly, the rv X_q having density q should be easy to sample from. Suppose that, as in the case of the CIR process, the FPT pdf g is unknown and the moments of T (or the cumulants of T) are known. We then assume that $\frac{g}{f}$ has the $L^2_{f_{\alpha,\beta}}$ series representation

$$\frac{g}{f_{\alpha,\beta}}=\sum_{k\geq 0} a_k^{(\alpha,\beta)} Q_k^{(\alpha,\beta)}. \tag{6.91}$$

Recall that from the latter we obtain the approximant $\hat{g}_n$ as

$$\frac{\hat{g}_n(t)}{f_{\alpha,\beta}(t)} = p_n(t), \quad t \in (0,\infty), \tag{6.92}$$

where p_n is given in (6.38). We further assume $\hat{g}_n$ is non-negative for all $t>0$. Note that, in practice, the latter assumption is not restrictive since we can always provide a positive approximant using the methodology proposed in Section 6.4.5. Expression (6.92) clearly resembles (6.90). Unfortunately, since p_n is a polynomial for any $n>0$, the right hand side of (6.92) is unbounded on $(0,\infty)$ and thus condition 6.90 is not satisfiable.

By suitably correcting for the unboundedness of p_n, we provide a modification of the standard acceptance-rejection method with the aim of sampling from the FPT rv T using (6.92). Let T_n be the rv with pdf $\hat{g}_n$ over $(0,\infty)$. As we continue, Condition 6.2.1 is assumed to hold, so that we can make use of statements 1. and 2. of Proposition 6.2.2, as shall become clearer in a moment. The main steps of the method we propose can be summarized as follows:

i) find a constant $C>0$ such that $\mathbb{P}(T>C)\leq\varepsilon$, for a fixed, small $\varepsilon>0$;

ii) for $t\leq C$ apply the classical acceptance-rejection method to a truncated approximant $\hat{g}_{T_n|C}$ using the ratio

$$\frac{\hat{g}_{T_n|C}(t)}{\tilde{f}_{\alpha,\beta|C}(t)} \leq M \quad \text{where} \quad M = \frac{\mathbb{P}(T_n\leq C)}{\mathbb{P}(X_{\alpha,\beta}\leq C)} \max_{t\in[0,C]} p_n(t) \tag{6.93}$$

and

$$\tilde{f}_{\alpha,\beta|C}(t) = \frac{f_{\alpha,\beta}(t)}{\mathbb{P}(X_{\alpha,\beta}\leq C)}\mathbf{1}_{(0,C]}(t), \quad \hat{g}_{T_n|C}(t) = \frac{\hat{g}_n(t)}{\mathbb{P}(T_n\leq C)}\mathbf{1}_{(0,C]}(t); \tag{6.94}$$

iii) for $t>C$ sample from a truncated exponential rv $\bar{T}$ with pdf

$$g_{\bar{T}}(t) = \frac{1}{\mathbb{E}[T]}\exp\left(-\frac{t-C}{\mathbb{E}[T]}\right)\mathbf{1}_{(C,+\infty)}(t). \tag{6.95}$$

Note that the last step takes into account 2. of Proposition 6.2.2, that is, the FPT pdf's exponential asymptotic behaviour for one-dimensional diffusion processes with steady-state distribution (see, for instance, Masoliver and Perelló, 2014; Nobile et al., 1985).

Remark 6.7.1. *Note that the last step could be implemented whenever we know that there exists a well known pdf h, whose corresponding rv X_h should be easy to sample from, such that*

$$g(t) \approx h(t), \quad t \to \infty.$$

Exploring this could be subject of future work (see Section 7.2.3 in Chapter 7).

Algorithm 3 outlines the proposed method for constructing an "approximated" sample S of size N from T.

Algorithm 3: Modified acceptance-rejection method

Set the parameters α, $\beta > 0$, $n \in \mathbb{N}^+$ and $\varepsilon > 0$.
Initialise
  Find a constant $C > 0$ such that $\mathbb{P}(T > C) \leq \varepsilon$.
  Set M as in (6.93).
  $j \leftarrow 1$.
  $S \leftarrow \{\}$.
While $j < N$
  With probability ε
    Generate $\bar{T}$ from the truncated exponential pdf in (6.95).
    $S \leftarrow S \cup \bar{T}$.
    $j \leftarrow j + 1$.
  With probability $1 - \varepsilon$
    Generate G from the truncated gamma pdf $\tilde{f}_{\alpha,\beta|C}$ in (6.94).
    Generate $U \sim U(0,1)$.
    While $U > \frac{\hat{g}_{T_n|C}(G)}{M \tilde{f}_{\alpha,\beta|C}(G)}$
      Generate G from the truncated gamma pdf $\tilde{f}_{\alpha,\beta|C}$ in (6.94).
      Generate $U \sim U(0,1)$.
    $S \leftarrow S \cup G$.
    $j \leftarrow j + 1$.
Return S.

In addition to the user-specified input parameters, the constant C in *i)* must be chosen for the algorithm initialization. The so-called *Vysochanskij-Petunin* inequality for one-sided tail bounds (Mercadier and Strobel, 2021) can be employed to achieve that.

Theorem 6.7.1 (Vysochanskij-Petunin inequality)**.** *If $r \geq 0$ and X is a rv*

with unimodal density, finite mean μ and finite variance σ^2, then

$$\mathbb{P}(X-\mu\geq r)\leq\begin{cases}\frac{4}{9}\frac{\sigma^2}{\sigma^2+r^2} & \text{if } 3r^2\geq 5\sigma^2,\\ \frac{4}{3}\frac{\sigma^2}{\sigma^2+r^2}-\frac{1}{3} & \text{otherwise.}\end{cases} \tag{6.96}$$

Since, under our assumptions, the moments of T are known and the FPT rv g is unimodal from 1. in Proposition 6.2.2, the Vysochanskij-Petunin inequality can be applied. Indeed, to recover a C such that $\mathbb{P}(T>C)\leq\varepsilon$, one can proceed as follows. Fix $\varepsilon>0$. Setting $4\sigma^2/[9(\sigma^2+r^2)]=\varepsilon$ in the first inequality (6.96), we recover $r=r(\varepsilon)$ as a function of ε and get the condition $\varepsilon\leq 1/6$ from $3r(\varepsilon)^2\geq 5\sigma^2$. Then set $C=\mu+r(\varepsilon)$ where

$$r(\varepsilon)=\begin{cases}\sqrt{\frac{4\sigma^2}{9\varepsilon}-\sigma^2} & \text{if } \varepsilon\leq 1/6,\\ \sqrt{\frac{4\sigma^2}{1+3\varepsilon}-\sigma^2} & \text{if } 1/6<\varepsilon\leq 1.\end{cases} \tag{6.97}$$

It is clear that the quality of the outcome of Algorithm 3 relies on the approximation $\hat{g}_n$ and the selection of an exponential distribution for $t>C$. A theoretical justification for it is provided in the following. At first, in the next result we compute the cdf of the rv Y whose observations are generated by Algorithm 3, using the classical proof of the validity of the acceptance-rejection method.

Lemma 6.7.2. *If Y denotes the rv sampled at the end of each cycle of Algorithm 3, then*

$$\mathbb{P}(Y\leq t)=\varepsilon\left[1-\exp\left(-\frac{t-C}{\mathbb{E}(T)}\right)\right]+(1-\varepsilon)\mathbb{P}(T_n\leq t\,|\,T_n\leq C),\qquad t>0$$

where C, n and ε are given in Algorithm 3, and T_n is the rv with pdf $\hat{g}_n$.

Proof. According to Algorithm 3, we have

$$\mathbb{P}(Y\leq t)=\mathbb{P}(X=1)\mathbb{P}(\tilde{T}\leq t)+\mathbb{P}(X=0)\mathbb{P}(G\leq t\,|\,G\text{accepted}) \tag{6.98}$$

where X is a Bernoulli rv of parameter $\varepsilon\in(0,1)$ independent from the rv $\tilde{T}$, with truncated exponential pdf $g_{\bar{\tau}}$ in (6.95), and the rv G with truncated gamma pdf $\tilde{f}_{\alpha,\beta|C}$ in (6.94). Thus from (6.95) and (6.98), we have

$$\mathbb{P}(Y\leq t)=\varepsilon\left[1-\exp\left(-\frac{t-C}{\mathbb{E}(T)}\right)\right]+(1-\varepsilon)\mathbb{P}(G\leq t\,|\,G\text{accepted}). \tag{6.99}$$

For the latter term in (6.99) observe that

$$\mathbb{P}(G\leq t\,|\,G\text{ accepted})=M\,\mathbb{P}(G\leq t, G\text{ accepted}) \tag{6.100}$$

since $\mathbb{P}(G \text{ accepted}) = 1/M$ with M given in (6.93). Moreover

$$\begin{aligned}\mathbb{P}(G \le t, G \text{ accepted}) &= \int_0^\infty \mathbb{P}(G \le t, G \text{ accepted} \,|\, G = x)\, \tilde{f}_{\alpha,\beta|C}(x)\, \mathrm{d}x \\ &= \int_0^\infty \mathbb{P}(G \le t \,|\, G = x)\, \mathbb{P}(G \text{ accepted} \,|\, G = x)\, \tilde{f}_{\alpha,\beta|C}(x)\, \mathrm{d}x \end{aligned} \tag{6.101}$$

since $(G \le t)$ and $(G \text{ accepted})$ are conditionally independent events. For any $x \in (0,\infty)$, observe that $\mathbb{P}(G \le t \,|\, G = x) = \mathbf{1}_{x \le t}$ and

$$\begin{aligned}\mathbb{P}(G \text{ accepted} \,|\, G = x) &= \mathbb{P}\left(U \le \frac{\hat{g}_{T_n|C}(G)}{M\, \tilde{f}_{\alpha,\beta|C}(G)} \,\middle|\, G = x \right) \\ &= \mathbb{P}\left(U \le \frac{\hat{g}_{T_n|C}(x)}{M\, \tilde{f}_{\alpha,\beta|C}(x)} \right) = \frac{\hat{g}_{T_n|C}(x)}{M\, \tilde{f}_{\alpha,\beta|C}(x)} \end{aligned} \tag{6.102}$$

since U is a rv with uniform distribution over (0,1). Plugging (6.102) in (6.101) and the resulting integral in (6.100), we get

$$\mathbb{P}(G \le t \,|\, G \text{ accepted}) = M \int_0^\infty \mathbf{1}_{x \le t} \frac{\hat{g}_{T_n|C}(x)}{M\, \tilde{f}_{\alpha,\beta|C}(x)}\, \tilde{f}_{\alpha,\beta|C}(x)\, \mathrm{d}x = \int_0^t \hat{g}_{T_n|C}(x)\, \mathrm{d}x,$$

and the result follows from (6.94). □

The following is a technical lemma necessary for the forthcoming proposition. For any $t \le C$, it provides an upper bound of the error in approximating the truncated FPT cdf $\mathbb{P}(T \le t \,|\, T \le C)$ with the corresponding approximated $\mathbb{P}(T_n \le t \,|\, T_n \le C)$ obtained using the Laguerre-Gamma expansion $\hat{g}_n$ in (6.37).

Lemma 6.7.3. *Under the same hypotheses as Lemma 6.7.2, for any $t \le C$ we have*

$$\begin{aligned}&\Big| \mathbb{P}(T_n \le t \,|\, T_n \le C) - \mathbb{P}(T \le t \,|\, T \le C) \Big| \\ &\qquad \le \frac{1}{\mathbb{P}(T_n \le C)} \left[\sum_{k \ge n+1} (a_k^{(\alpha,\beta)})^2 \right]^{\frac{1}{2}} \left(1 + \frac{1}{\mathbb{P}(T \le C)} \right) \end{aligned} \tag{6.103}$$

where $a_k^{(\alpha,\beta)} = \mathbb{E}[Q_k^{(\alpha,\beta)}(T)]$ for $k \ge 0$, with $Q_k^{(\alpha,\beta)}$ as in (6.11).

Proof. For $t \leq C$, let $g_{T|C}(t) = \mathbf{1}_{(0,C]}(t) g(t)/\mathbb{P}(T \leq C)$ be the truncated FPT pdf and $\hat{g}_{T_n|C}(t)$ as in (6.94). We have

$$\begin{aligned}\left|\hat{g}_{T_n|C}(t) - g_{T|C}(t)\right| &= \left|\frac{\hat{g}_n(t)}{\mathbb{P}(T_n \leq C)} - \frac{g(t)}{\mathbb{P}(T \leq C)}\right| \\ &\leq \frac{1}{\mathbb{P}(T_n \leq C)}\left|\hat{g}_n(t) - g(t)\right| + g(t)\left|\frac{1}{\mathbb{P}(T_n \leq C)} - \frac{1}{\mathbb{P}(T \leq C)}\right|.\end{aligned}$$

Using the previous inequality, the difference between the truncated FPT cdf and its approximation by means of the Laguerre-Gamma expansion may be bounded as follows:

$$\begin{aligned}&\left|\mathbb{P}(T_n \leq t \,|\, T_n \leq C) - \mathbb{P}(T \leq t \,|\, T \leq C)\right| = \left|\int_0^t [\hat{g}_{T_n|C}(s) - g_{T|C}(s)]\,\mathrm{d}s\right| \\ &\leq \int_0^\infty \left|\hat{g}_{T_n|C}(t) - g_{T|C}(t)\right]|\mathrm{d}t \leq \frac{1}{\mathbb{P}(T_n \leq C)}\int_0^\infty \frac{|\hat{g}_n(s) - g(s)|}{\varphi_{\alpha,\beta}(s)}\varphi_{\alpha,\beta}(s)\,\mathrm{d}s \\ &+ \left|\frac{1}{\mathbb{P}(T_n \leq C)} - \frac{1}{\mathbb{P}(T \leq C)}\right| \underbrace{\int_0^\infty g(s)\,\mathrm{d}s}_{=1} \\ &= \frac{1}{\mathbb{P}(T_n \leq C)}\left\|\frac{\hat{g}_n - g}{f_{\alpha,\beta}}\right\|_{L^1(\nu)} + \left|\frac{1}{\mathbb{P}(T_n \leq C)} - \frac{1}{\mathbb{P}(T \leq C)}\right| \qquad (6.104)\end{aligned}$$

where L^1_ν is the Hilbert space of the integrable functions with respect to the measure ν having density $f_{\alpha,\beta}$. Observe that

$$\left\|\frac{\hat{g}_n - g}{f_{\alpha,\beta}}\right\|_{L^1(\nu)} \leq \left\|\frac{\hat{g}_n - g}{f_{\alpha,\beta}}\right\|_{L^2_\nu} = \left[\sum_{k \geq n+1} (a_k^{(\alpha,\beta)})^2\right]^{1/2} \qquad (6.105)$$

where the last equality follows from (6.32). Finally, since

$$\left|\frac{1}{\mathbb{P}(T_n \leq C)} - \frac{1}{\mathbb{P}(T \leq C)}\right| \leq \frac{1}{\mathbb{P}(T_n \leq C)}\left(1 + \frac{\mathbb{P}(T_n \leq C)}{\mathbb{P}(T \leq C)}\right) \qquad (6.106)$$

with $\mathbb{P}(T_n \leq C) \leq 1$, plugging (6.106) and (6.105) in (6.104) the result follows. □

Given the latter lemma, we are able to state the following proposition which provides a theoretical justification of the proposed acceptance-rejection Algorithm 3.

Proposition 6.7.4. *Under the same hypotheses as Lemma 6.7.2, for every $\delta > 0$, there exist a finite constant C_δ and an integer $n_\delta \in \mathbb{N}$ such that, for $n > n_\delta$ and $C > C_\delta$ in Algorithm 3, we have*

$$\left|\mathbb{P}(Y \leq t) - \mathbb{P}(T \leq t)\right| < \delta, \qquad t > 0.$$

Proof. Fix δ, $n>0$ and $\varepsilon>0$ and let $C>0$ be the corresponding quantity calculated in Algorithm 3. Let any $t>0$. By plugging $\mathbb{P}(T\leq t)=\mathbb{P}(T\leq C)\mathbb{P}(T\leq t|T\leq C)+\mathbb{P}(T>C)\mathbb{P}(T\leq t|T>C)$ in $|\mathbb{P}(Y\leq t)-\mathbb{P}(T\leq t)|$ and using Lemma 6.7.2, the following bound can be recovered:

$$|\mathbb{P}(Y\leq t)-\mathbb{P}(T\leq t)|\leq P_1+P_2$$

where

$$\begin{aligned} P_1 &:= |\varepsilon\mathbb{P}(\bar{T}\leq t\,|\,\bar{T}>C)-\mathbb{P}(T>C)\mathbb{P}(T\leq t\,|\,T>C)| \\ P_2 &:= |(1-\varepsilon)\mathbb{P}(T_n\leq t\,|\,T_n\leq C)-\mathbb{P}(T\leq C)\mathbb{P}(T\leq t\,|\,T\leq C)|, \end{aligned}$$

$\bar{T}$ is the truncated exponential rv with pdf in (6.95) and T_n is the rv with pdf $\hat{g}_n$. Now, let us bound P_1. Adding and subtracting the quantity $\varepsilon\mathbb{P}(T\leq t\,|\,T>C)$ in P_1, we get

$$P_1\leq P_{1,1}+P_{1,2} \quad \text{with} \quad \begin{cases} P_{1,1} := \varepsilon|\mathbb{P}(\bar{T}\leq t\,|\,\bar{T}>C)-\mathbb{P}(T\leq t\,|\,T>C)|, \\ P_{1,2} := \mathbb{P}(T\leq t\,|\,T>C)|\varepsilon-\mathbb{P}(T>C)|. \end{cases}$$

Since $\mathbb{P}(T>C)\leq\varepsilon$ by the definition of C and thanks to the exponential behaviour of the tails of the FPT pdf which holds by 2. of Proposition 6.2.2, by eventually decreasing ε it is always possible to find $C_{\delta,1}$ (big enough) such that $P_{1,1}<\frac{\delta}{4}$ and $P_{1,2}<\frac{\delta}{4}$. Let us apply the same strategy to P_2. Adding and subtracting the quantity $(1-\varepsilon)\mathbb{P}(T\leq t\,|\,T\leq C)$ in P_1, we get

$$P_2\leq P_{2,1}+P_{2,2} \quad \text{with} \quad \begin{cases} P_{2,1} := (1-\varepsilon)|\mathbb{P}(T_n\leq t\,|\,T_n\leq C)-\mathbb{P}(T\leq t\,|\,T\leq C)|, \\ P_{2,2} := \mathbb{P}(T\leq t\,|\,T\leq C)|\varepsilon-\mathbb{P}(T>C)|. \end{cases}$$

From Lemma 6.7.3, the quantity $P_{2,1}$ can be made sufficiently small by taking an $n_\delta\in\mathbb{N}$ large enough since the remainder (6.32) goes to zero as n increases, as we have assumed (6.109). Similarly as what has been done for P_1 above, thanks to the exponential behaviour of the FPT pdf tail, by eventually decreasing ε we can always find $C_{\delta,2}$ and n_δ such that $P_{2,1}<\frac{\delta}{4}$ and $P_{2,2}<\frac{\delta}{4}$. Setting $C_\delta=\max(C_{\delta,1},C_{\delta,2})$ concludes the proof. □

Remark 6.7.5. *Note that although Proposition 6.7.4 guarantees that we can always choose a constant C and an order n such that the cdf of Y sampled in each cycle of Algorithm 3 is sufficiently close to the cdf of T, as C increases by decreasing ε, as seen in (6.97), $M=\max_{t\in[0,C]}p_n(t)$ will increase or at most remain constant. However, the probability of acceptance in Algorithm 3 is $\frac{1}{M}$. Therefore, there is a trade-off between increasing C, to achieve better accuracy, and the running time of the proposed algorithm.*

The ability of Algorithm 3 to generate a satisfactory "approximated" sample from the FPT rv is shown in Figs. 6.7.1, 6.7.2, and 6.7.3. Since we have initially assumed Condition 6.2.1, the method has been hereby applied only to the CIR process. Additionally, in this case, the lack of a closed form of g_{cir} and of exact simulation methods are strong motivations for investigating this procedure. We consider cases A_{cir}, B_{cir} and C_{cir} as in Section 6.6.3 with the standardisation procedure outlined in Section 6.4.3. Hence in (6.94), instead of $\hat{g}_n$, we use the approximation $\tilde{g}_n$ of the pdf $\tilde{g}_{cir}$ corresponding to $\tilde{T} = T/\sigma_T$, where σ_T is the standard deviation of T. Moreover, whenever mentioned, n in $\tilde{g}_n$ has been computed using the stopping criteria in (6.47) with $\varepsilon = 10^{-14}$ and the parameters α and β in (6.37) have been chosen according to the first part of Section 6.4.3. Again, case B_{cir} requires a special attention for two main reasons. On one hand, the corresponding $\alpha = -0.34$ and $\beta = 0.812$ leads to a peaked reference distribution, which increases the probability of rejection. On the other hand, a high numerical precision is required to obtain a satisfactory $\tilde{g}_n$, as already mentioned. However, in our numerical experiments the main source of a lengthy computational time was the higher order n used for obtaining $\tilde{g}_n$, rather than the acceptance-rejection method itself. Despite this, the overall computational time remained competitive when confronted with the remaining cases.

6.7.2 Orthogonal Series Estimators

Suppose a sample of i.i.d. FPTs $\mathcal{T} = \{T_1, \ldots, T_N\}$ generated by a FPT rv T having unknown density g is available, arising either from simulations or from experiments. We additionally assume that not even the moments (or cumulants) of T are known. Nevertheless, we can still assume that $\frac{g}{f}$ has a $L^2_{f_{\alpha,\beta}}$ series representation and as usual from the latter we can formally obtain the approximant $\hat{g}_n$, even though it cannot be directly computed. However, it can be still exploited in this context as shown in the following. Indeed, given the availability of a sample it is intuitive to replace the theoretical moments appearing in the cofficients $\{\mathcal{B}_k^{(\alpha,\beta)}\}_{k\geq 0}$ of $\hat{g}_n$ as in (6.24) with the corresponding sample moments computed on $\mathcal{T}$. Proceeding in this way, a straightforward calculation shows that replacing FPT moments $\{\mathbb{E}[T^j]\}_{j=1}^n$ with sample moments is equivalent to replacing $\mathbb{E}[L_k^{(\alpha,\beta)}(T)]$ in (6.24) with

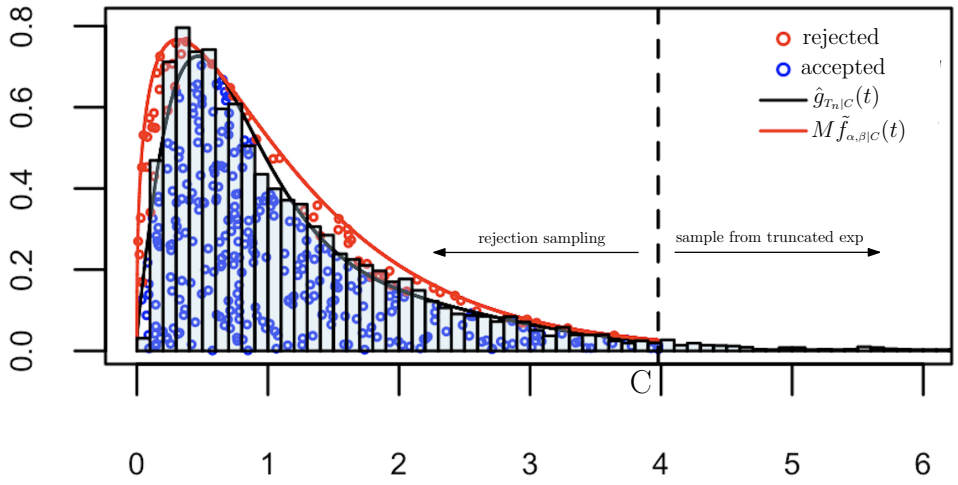

Figure 6.7.1: Plot of rejected draws (in red) and accepted draws (in blue) for the acceptance-rejection part of Algorithm 3 with $N=3000$, $\epsilon=0.05$ and $n=10$ applied to case A_{cir}, together with the corresponding $M\tilde{f}_{\alpha,\beta|C}$ and $\tilde{g}_{T_n|C}=\frac{\tilde{g}_n}{\mathbb{P}(T_n\leq C)}\mathbf{1}_{(0,C]}$ as in (6.94), with $C=3.97, \alpha=0.367$ and $\beta=1.17$. Only a portion (10%) of rejected and accepted draws has been plotted, for easier viewing. On the same plot an histogram computed on the sample of size $N=3000$ arising from Algorithm 3 with the aforementioned parameters is shown.

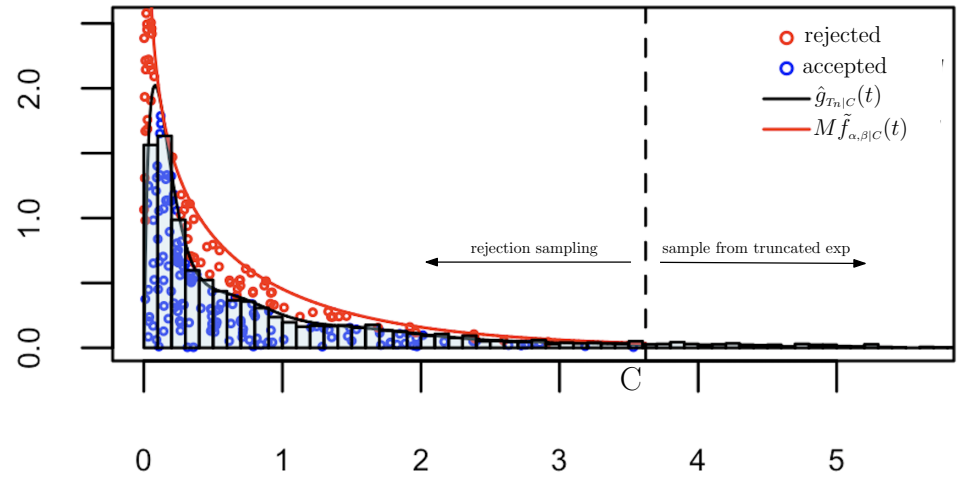

Figure 6.7.2: Plot of rejected draws (in red) and accepted draws (in blue) for the acceptance-rejection part of Algorithm 3 with $N=3000$, $\epsilon=0.05$ and $n=55$ applied to case B_{cir}, together with the corresponding $M\tilde{f}_{\alpha,\beta|C}$ and $\tilde{g}_{T_n|C}=\frac{\tilde{g}_n}{\mathbb{P}(T_n\leq C)}\mathbf{1}_{(0,C]}$ as in (6.94), with $C=3.621$, $\alpha=-0.34$ and $\beta=0.812$. Only a portion (10%) of rejected and accepted draws has been plotted, for easier viewing. On the same plot an histogram computed on the sample of size $N=3000$ arising from Algorithm 3 with the aforementioned parameters is shown.

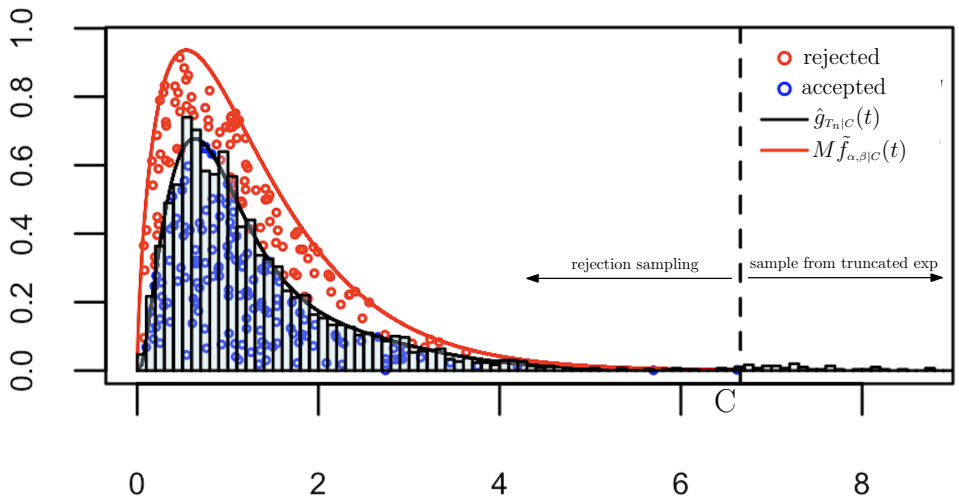

Figure 6.7.3: Plot of rejected draws (in red) and accepted draws (in blue) for the acceptance-rejection part of Algorithm 3 with $N=3000$, $\epsilon=0.015$ and $n=9$ applied to case C_{cir}, together with the corresponding $M\tilde{f}_{\alpha,\beta|C}$ and $\tilde{g}_{T_n|C}=\frac{\tilde{g}_n}{\mathbb{P}(T_n\leq C)}\mathbf{1}_{(0,C]}$ as in (6.94), with $C=6.65$, $\alpha=0.7$ and $\beta=1.306$. Only a portion (10%) of rejected and accepted draws has been plotted, for easier viewing. On the same plot an histogram computed on the sample of size $N=3000$ arising from Algorithm 3 with the aforementioned parameters is shown.

its sample mean estimator

$$\bar{l}_k=\frac{1}{N}\sum_{i=1}^{N}L_k^{(\alpha,\beta)}(T_i).$$

Then, the FPT pdf g which has generated the sample $\mathcal{T}$ can estimated for $t\in(0,\infty)$ and $n\in\mathbb{N}$ by

$$\bar{g}_n(t)=f_{\alpha,\beta}(t)\left(1+\sum_{k=1}^{n}\bar{l}_k b_k^{(\alpha,\beta)}L_k^{(\alpha,\beta)}(t)\right)\quad\text{with}\quad b_k^{(\alpha,\beta)}=\frac{k!\Gamma(\alpha+1)}{\Gamma(\alpha+1+k)},\tag{6.107}$$

which is known as an *orthogonal series estimator* of g (see, e.g., Efromovich, 2010). This observation reveals an additional advantage of using the Laguerre-Gamma approximation. If the FPT moments/cumulants are not known but a random sample is available, the Laguerre-Gamma approach offers the opportunity to recover an approximation of the FPT pdf similarly as an orthogonal series estimator. In such a case, the estimates carried out by sample moments or by k-statistics (Di Nardo and Guarino, 2022) replace the occurrences of FPT moments or cumulants respectively.

Figs 6.7.4 and 6.7.5 respectively show that the results obtained by the

orthogonal series method on a collected FPT sample are pretty much equivalent to the Laguerre-Gamma approximations in the case of known moments presented in Section 6.6.3 and 6.6.2 respectively, when the stopping criteria addressed in Section 6.4.2 are used for both the procedures.

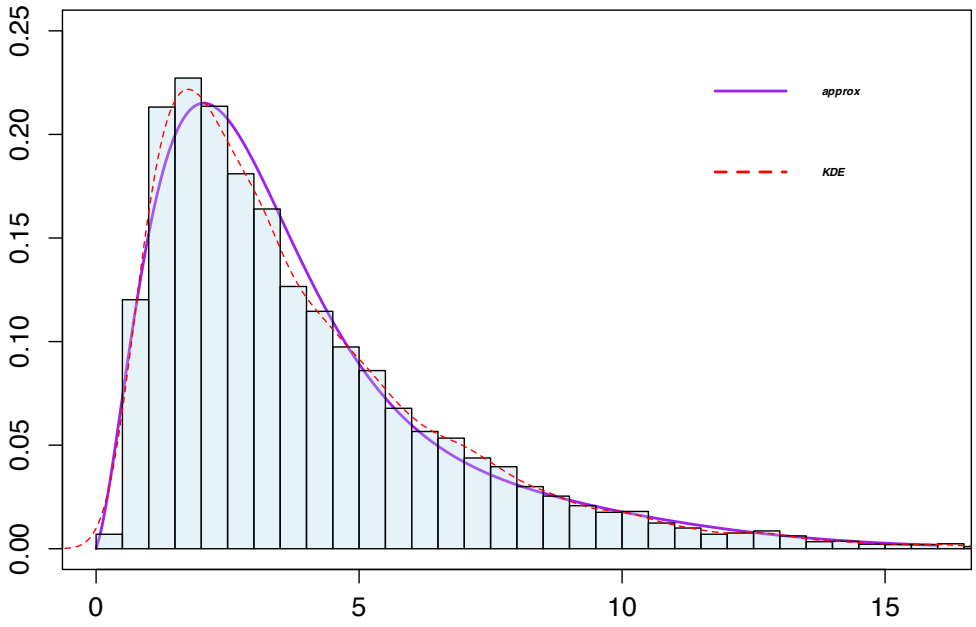

Figure 6.7.4: Plot of $\bar{g}_n$ (in solid purple) from the sample $\mathcal{T}_{C_{cir}}$ for case C_{cir} with $n=7$, $\alpha=0.8$ and $\beta=0.49$, obtained with the stopping criteria (6.47), together with a KDE (in dashed red) and a histogram both computed with the sample $\mathcal{T}_{C_{cir}}$

In this context, it is natural to ask for a method to choose n in $\bar{g}_n$ which depends on the sample $\mathcal{T}$ at hand. We briefly recall the general idea in the following. Under additional and somehow complicated hypotheses on the true pdf g, estimations of the convergence order of $\bar{g}_n$ to g are assessable through the mean integrated squared error (Hall, 1980). By using the orthogonality of the generalised Laguerre polynomial sequence, for $n \in \mathbb{N}$ the

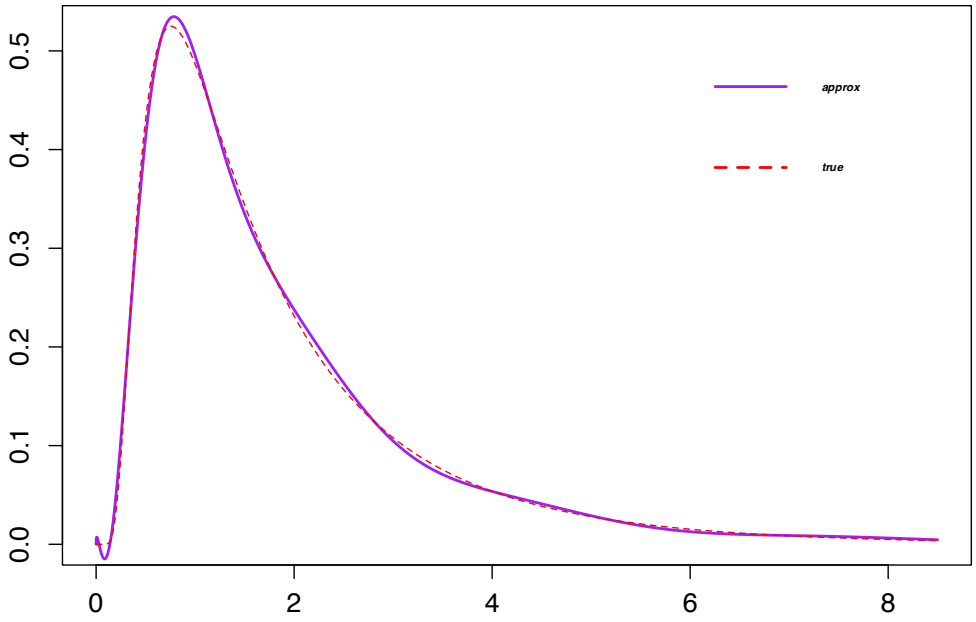

Figure 6.7.5: Plot of $\bar{g}_n$ (in solid purple) from a sample $\mathcal{T}_{B_{gbm}}$ of size $N = 10^4$ generated with the Milstein method for case B_{gbm} with $n = 34$, $\alpha = 0.39$ and $\beta = 0.73$, obtained with the stopping criteria (6.47), together with a plot of g_{gbm} (in dashed red).

mean integrated squared error is (Diggle and Hall, 1986)

$$\begin{aligned} J(n) &= \mathbb{E}\left[\int_0^\infty \left|\frac{\bar{g}_n(t) - g(t)}{f_{\alpha,\beta}(t)}\right|^2 f_{\alpha,\beta}(t)\,\mathrm{d}t\right] \\ &= \frac{1}{N}\sum_{k=1}^{n} b_k^{(\alpha,\beta)} \mathrm{Var}[L_k^{(\alpha,\beta)}(T)] + \sum_{k>n} b_k^{(\alpha,\beta)} \mathbb{E}[L_k^{(\alpha,\beta)}(T)]^2 \end{aligned} \tag{6.108}$$

with $b_k^{(\alpha,\beta)}$ as in (6.107). Thus increasing the sample size N leads to a reduction of the error $J(n)$ as expected from the estimation of moments with sample moments.

The vast majority of the strategies in the literature to choose the degree of the polynomial approximation n in (6.107) starts from the idea of minimising (6.108), see for example Diggle and Hall (1986). A discussion on which strategy is the most effective in our two applications, or, additionally, a general treatment of this procedure, goes beyond the scope of this thesis.

6.7.3 Approximated Maximum Likelihood Estimation

As in the previous context, suppose a sample of i.i.d. FPTs $\mathcal{T}=\{T_1,\dots,T_N\}$ generated by a FPT rv T having unknown density g is available, arising either from simulations or from experiments. The moments (or cumulants) of T are supposed to be known. As we continue, we will denote g by g_θ, to make explicit the fact that it depends on the parameters $\theta\in\Theta$ of the underlying process, considering the starting position and barrier to be fixed and known. For instance, $\theta=(\mu,\sigma^2)\in(-\infty,+\infty)\times(0,+\infty)=\Theta$ for the GBM. The goal is to estimate them with a suitably modified maximum likelihood estimation. In this case, we assume that for any $\theta\in\Theta$ the ratio $\frac{g_\theta}{f_{\alpha,\beta}}$ has a $L^2_{f_{\alpha,\beta}}$ series representation

$$\frac{g_\theta}{f_{\alpha,\beta}}=\sum_{k\geq 0}a_{k,\theta}^{(\alpha,\beta)}Q_k^{(\alpha,\beta)},\quad \theta\in\Theta. \tag{6.109}$$

From the latter we obtain the approximant $\hat{g}_{n,\theta}$ as

$$\frac{\hat{g}_{n,\theta}(t)}{f_{\alpha,\beta}(t)}=p_{n,\theta}(t),\quad t\in(0,\infty),\quad \theta\in\Theta, \tag{6.110}$$

where p_n is given in (6.38). We further assume that for any $\theta\in\Theta$ and $n\in\mathbb{N}$ the approximant $\hat{g}_{n,\theta}$ is non-negative for all $t>0$. Note that in the series representation (resp. the approximant) we have that the dependence on θ is found in the moments $\{\mathbb{E}[T^j]_\theta\}_{j\geq 0}$ which appear in the coefficients $a_{k,\theta}^{(\alpha,\beta)}$ of the series (resp. the sum). Then, it is intuitive to proceed by using the approximated density $\hat{g}_{n,\theta}$ in the likelihood function instead of the true unknown density g_θ. Note that n denotes the truncation order and N denotes the sample size. The approximate MLE (AMLE) of θ then is

$$\hat{\theta}_{n,N}=\operatorname*{argmax}_{\theta\in\Theta} l_{n,N}(\theta), \tag{6.111}$$

with $l_{n,N}(\theta)=\ln L_{n,N}(\theta)$, and

$$L_{n,N}(\theta)=\prod_{i=1}^{N}\hat{g}_{n,\theta}(T_i).$$

The maximisation problem in (6.111) must be numerically solved. Before proceeding, we make note to the reader that extensive additional studies must be set forth to theoretically ground this AMLE. This issue may be subject of future work and will be discussed in the appropriate section at the end of this thesis (see Section 7.2.3 in Chapter 7). Nevertheless, we provide

some numerical experiments in the case of the CIR. Following standard practices when using this process in the context of neuronal modelling (see, e.g., Ditlevsen and Lansky, 2006), we have only tackled the estimation of the parameters μ and σ in (6.55) and we consider the remaining ones fixed and known. The true parameters have been selected according to cases A_{cir}, B_{cir} and C_{cir} in Section 6.6.3. For each of the three combination of parameters, 1000 instances of the AMLE in (6.111) have been computed for three simulated samples $\mathcal{T}^N_{A_{cir}}$, $\mathcal{T}^N_{B_{cir}}$ and $\mathcal{T}^N_{C_{cir}}$ with sample size $N \in \{500, 1000, 5000\}$. In each case, the maximisation problem in (6.111) has been solved by using a global optimisation algorithm implemented in the R package `nloptr` (Johnson, 2024). Additionally, recall that the order of approximation n, α and β must be chosen when computing the approximant $\hat{g}_n$. These choices require a few more words. For the former, a maximum n_{max} is set at the beginning of the numerical optimisation involved in computing the AMLE for a given sample: whenever $\hat{g}_n$ needs to be evaluated, the iterative procedure for computing it is stopped if our stopping criteria (6.47) are met or if n_{max} is reached. For lightening the computational costs, n_{max} has been set to 10 in all cases considered. For the latter parameters, two possibilities have been considered. The first one involved fixing α and β at the beginning of the numerical optimisation involved in computing the AMLE for a given sample, using the choice in Section 6.4.3 employing sample moments computed from the considered sample instead of theoretical moments. The second one instead goes as follows. For a given sample, for each (μ', σ') in which one wants to evaluate $\hat{g}_n$ during the numerical optimisation, a couple (α', β') is computed using the choice in Section 6.4.3 with the theoretical moments corresponding to (μ', σ'). The latter choice proved to be more well behaved and it is the one chosen for the results presented here. However, we stress out again that all of these aspetcs must be further investigated. The bias and MSE resulting from these experiments are reported in Table 6.7.1. An inspection of the latter reveals satisfactory values which confirm that the proposed AMLE is a promising technique worthy of further studies.

6.8 Conclusions

In this chapter we have studied a method for approximating the FPT pdf and cdf of a one dimensional regular diffusion process that relies on a series expansion involving the generalised Laguerre polynomials and the gamma

	Bias		MSE	
	μ	σ	μ	σ
$\mathcal{T}^{500}_{A_{cir}}$	8.86×10^{-3}	-7.85×10^{-2}	4.88×10^{-3}	2.05×10^{-2}
$\mathcal{T}^{1000}_{A_{cir}}$	4.83×10^{-4}	-6.15×10^{-2}	1.30×10^{-3}	1.11×10^{-2}
$\mathcal{T}^{5000}_{A_{cir}}$	-2.73×10^{-3}	-6.68×10^{-2}	7.97×10^{-4}	1.05×10^{-2}
$\mathcal{T}^{500}_{B_{cir}}$	5.44×10^{-5}	1.36×10^{-1}	6.63×10^{-7}	6.77×10^{-1}
$\mathcal{T}^{1000}_{B_{cir}}$	3.07×10^{-5}	6.56×10^{-2}	3.23×10^{-7}	1.80×10^{-1}
$\mathcal{T}^{5000}_{B_{cir}}$	8.09×10^{-5}	3.93×10^{-2}	1.79×10^{-7}	1.93×10^{-3}
$\mathcal{T}^{500}_{C_{cir}}$	-2.71×10^{-2}	4.91×10^{-2}	4.97×10^{-2}	3.21×10^{-3}
$\mathcal{T}^{1000}_{C_{cir}}$	-3.23×10^{-2}	4.15×10^{-2}	1.04×10^{-2}	1.77×10^{-3}
$\mathcal{T}^{5000}_{C_{cir}}$	-3.74×10^{-2}	4.11×10^{-2}	5.77×10^{-3}	1.73×10^{-3}

Table 6.7.1: For cases A_{cir}, B_{cir} and C_{cir} in Section 6.6.3, the bias and MSE of the AMLE in (6.111) for (μ,σ), estimated from 1000 random samples of size $N\in\{100,500,1000\}$ generated, as explained in Section 6.6.1, with the Milstein method.

pdf. We began by recalling the more general method of moment based series representations for pdfs, which hinges upon expanding them with respect to a family of complete orthonormal polynomials and a reference density. Then, with this theory at hand, we have explained how it can be applied to approximating the FPT pdf and cdf, and, under appropriate assumptions, we have given a result on the approximation error. In the same spirit, considerations on the choice of the reference pdf parameters have been made as well and some conditions for the non-negativity of the approximant in a right hand side of the origin and in the tail have been given. Furthermore, methods of standardisation according to considerations on dispersion measures are proposed. Naturally, we have then discussed the computational aspects. An iterative procedure coupled with stopping criteria aimed to ensure numerical stability have been established. We have also considered a strategy to overcome the possible negativity of the resulting function, a feature that is clearly undesirable in the approximation of a pdf.

The accuracy of the approximation has been tested on the GBM and on the CIR process. The latter case where the true FPT density is not known provides a prototypical example for the usefulness of this tool. Three numerical examples, spanning various shapes of FPT pdfs and cdfs, have been investigated for both of them.

As a side result of the practicality of the approximation three applications have been considered. One is an acceptance-rejection-like method, that makes a clever use of the form of the approximation function. It allows the generation of FPT data when its distribution is unknown but the moments (or cumulants) are known.

When only a sample of FPT data is available, the second application shows that the structure of the Laguerre-Gamma series expansion can be used to construct an estimator of the FPT pdf. This technique is known in the literature as the method of orthogonal series estimators.

Finally, with the third application, we have set forth an example of how the Laguerre-Gamma series expansion could be used to perform parameter estimation, in the case where a sample of FPT data is available, and, unlike in the previous scenario, the moments are known.

Future work will be discussed in the ending section of this thesis.

Appendix

6.A Generalised Laguerre Polynomials

In the following we recall the definition and some properties of the generalised Laguerre polynomials $L_n^{(\alpha)}$. We refer the reader to two classical monographs Sansone (1991) and Szegő (1975) for more details. They are usually defined for $\alpha > -1$ and non-negative integers n by Rodrigues' formula

$$L_n^{(\alpha)}(x) = \frac{1}{n!} x^{-\alpha} e^x \frac{d^n}{dx^n}\left(e^{-x} x^{n+\alpha}\right), \quad x \in [0,\infty).$$

From the latter, the closed form expression for the generalised Laguerre polynomials can be recovered

$$L_n^{(\alpha)}(x) = \sum_{k=0}^{n} \binom{n+\alpha}{n-k} \frac{(-x)^k}{k!}, \quad x \in [0,\infty).$$

For instance, for $n \in \mathbb{N}$ and $x \in [0,\infty)$ they satisfy the following recurrence relations

1.
$$(n+1)L_{n+1}^{(\alpha)}(x) = (2n+\alpha+1-x)L_n^{(\alpha)}(x) - (n+\alpha)L_{n-1}^{(\alpha)}(x). \tag{6.112}$$

2.
$$L_{n+1}^{(\alpha)}(x) = \left(\frac{2n+\alpha+1-x}{n+1}\right) L_n^{(\alpha)}(x) - \left(\frac{n+\alpha}{n+1}\right) L_{n-1}^{(\alpha)}(x). \tag{6.113}$$

3.
$$\frac{d}{dx} L_n^{(\alpha)}(x) = -L_{n-1}^{(\alpha+1)}(x). \tag{6.114}$$

The polynomials $L_n^{(\alpha)}$ are orthogonal with respect to the weight function $x^\alpha e^{-x}$ on the interval $[0,\infty)$, that is

$$\int_0^\infty x^\alpha e^{-x} L_n^{(\alpha)}(x) L_m^{(\alpha)}(x)\, dx = \frac{\Gamma(n+\alpha+1)}{n!} \delta_{nm}. \tag{6.115}$$

Finally, we also recall that the generalised Laguerre polynomials are eigenfunctions of the following Sturm-Liouville problem

$$\frac{\mathrm{d}}{\mathrm{d}t}\left(t^{\alpha+1}e^{-t}y'\right)+kt^{\alpha}e^{-t}y=0, \quad \text{with} \quad y=y(t), \quad k\in\mathbb{N}^{+}. \tag{6.116}$$

Chapter 7

Conclusions and Future Works

7.1 Conclusions

In this thesis we have presented three main lines of research related by the underlying unifying thread of hydrological data analysis and modelling.

The first part was dedicated to discrete models for interarrival times and other related rainfall temporal variables, with the broader aim of providing a description of the rainfall occurrence process. We reviewed the HLZ distribution and expressed its properties in a more compact way, providing some new insights. A new result on its convolution was provided. To expand the range of options for fitting daily rainfall interarrival times data, we further explored two additional distributions which have never been considered for such variables and are closely associated to the Lerch family: the Poisson-stopped HLZ distribution, and a one inflated HLZ distribution. Subsequently, we introduced other temporal variables linked to daily rainfall interarrival times, specifically wet and dry spells as well as wet and dry chains. At the best of our knowledge, the latter have never been deeply investigated in the literature. We then demonstrated how their distributions can be derived from the distribution of interarrival times and vice versa. Hinging upon these relationships, firstly a procedure called direct method (DM) was presented, where the distribution of wet spells and dry spells (as well as of the corresponding chains) is derived as a consequence of the assumption of i.i.d. interarrival times. In other words, the rainfall occurrence process is described by a *renewal process*. Note that in this case, the wet spells will have a geometric distribution. Secondly, the latter assumption was relaxed in an indirect method (IM) where wet spells and dry spells are modelled separately, that is, under the assumption of an *alternating renewal*

process, hence including the possibility of a non-constant rain probability inside a rainfall cluster. Then we presented the empirical results obtained on daily rainfall data arising from measurements at 6 European stations which span a variety of rainfall regimes. Firstly, the three aforementioned discrete distributions were fitted and compared on rainfall depths and interarrival times obtained from the available dataset. On the one hand, the obtained results confirm the effectiveness and flexibility of the HLZ distribution, on the other, they showed how the Poisson stopped HLZ distribution may prove to be superior in some cases, where it is able to match the performance of the one inflated HLZ distribution. Secondly, we applied the DM and IM to the same data. Our findings indicate that the geometric distribution does not consistently replicate the ws frequencies accurately, even in cases where it frequencies are well modelled by the HLZ distribution. Enhanced performance is observed when using the IM, suggesting the existence of an underlying internal structure within multiday rainfall events. It is commonly assumed that the internal structure observed in sub-daily rainfall records (e.g., 10-minute intervals) vanishes when data is aggregated at daily levels (see, e.g., Ridolfi et al., 2011). However, the analysis presented here contest this assumption for certain studied locations, indicating that the internal dynamics of rainfall events may persist even at the daily time scale. In other words, these results highlights the need for more flexible modelling approaches, such as the IM, to accurately capture these complexities. An additional enhancement in the fittings is achieved when the datasets are divided into two periods, indicating that incorporating local seasonality could improve inferences. Interestingly, in most locations, the DM applied to seasonal data remains a viable and straightforward approach. The results of this study may help in scenario simulations of drought and flood events, considering that probabilistic functions, such as those applied in this work, are at the base of stochastic climate modelling.

Motivated by the problem of describing the dependence between some of the variables encountered in the previous chapters and inspired by the ideas contained in the essay on dependence between random vector Geenens (2020), in the second part of the thesis we proposed copula-like models for finitely supported bivariate discrete random vectors. Initially, we reviewed some of Geenens (2020)'s main ideas and explained how they form the conceptual building block of our results. Using the fundamental concept of I-projection (in the sense of Csiszár, 1975) on a Fréchet class of pmfs with fixed arbitrary (positive) margins, we then were able prove a Sklar-like decomposition for

a bivariate pmf p into its two margins and a unique copula pmf u, playing the role of a representative of dependence. Additionally supposing that p has a rectangular support (that is, p is strictly positive), this decomposition was then exploited to build a statistical model for p, hinging upon the possibility of separately specifying the margins and the associated copula pmf u. Regarding both the nonparametric and parametric estimation of the latter, inferential and goodness of fit tools were developed and studied both theoretically and empirically. Indeed, theoretical results were complemented by finite-sample experiments and a data example. As a result that is both key for the study of the asymptotics of these inferential procedures and of independent interest, we proved that I-projections (in the sense of Csiszár, 1975) on a Fréchet class of pmfs with fixed arbitrary (positive) margins are differentiable in a certain sense. Afterwards, we tackled the generalisation of this result in the following way. When substituting the Kullback–Leibler divergence, which underlies I-projections, with the more general class of the so-called ϕ-divergences (see, e.g., Ali and Silvey, 1966; Csiszár, 1967; Liese and Vajda, 1987; Csiszár and Shields, 2004, and the references therein), one obtains ϕ-projections. We have then established conditions under which ϕ-projections are continuously differentiable, in the broader context of finite measures on finite spaces. When the target set for the ϕ-projection is convex, we demonstrated that the necessary assumptions can be derived from straightforward and verifiable conditions. These findings were utilized to determine the asymptotics of ϕ-projection estimators (i.e., minimum ϕ-divergence estimators) for projections onto parametric sets of probability vectors, sets of probability vectors with specified fixed moments, and finally on Fréchet classes of bivariate probability arrays, as done previously, but with a more general point of view, particularly amenable to multivariate extensions. In these examples our theoretical findings were empirically tested with various simulations.

Since the study of the time that some hydrological variables take to reach certain thresholds holds an important role in applications, the third part of this thesis dealt with a method for approximating the first passage time probability density function and cumulative distribution function of some one dimensional diffusions. This approximation is obtained by truncating a series expansion involving the generalised Laguerre polynomials and the gamma probability density, and it relies on the knowledge of the moments, or equivalently, of the cumulants of the FPT random variable. After reviewing the more general procedure underlying this tool, sufficient conditions for

the existence of the series expansion in the FPT case were given and, under the latter, a result on the approximation error has been provided. We then naturally considered the computational aspects. In this regard, we provided an iterative procedure coupled with stopping criteria aimed to ensure numerical stability and we also considered a strategy to overcome the possible negativity of the resulting function. Furthermore, methods for choosing the parameters of the gamma density, along with a simple standardisation trick, were proposed, both developed according to considerations on dispersion measures. Various numerical experiments on the GBM and on the CIR process were performed to confirm the efficiency and practicality of the method. The case of the CIR process, where the true FPT density is not known, provided a prototypical example for the usefulness of this tool. As a byproduct of the approximation, we were able to present three applications. One is an acceptance-rejection-like method, that makes a clever use of the form of the approximation function. It allows the generation of FPT data when its distribution is unknown but the moments (or cumulants) are known. Again, we have shown its effectiveness in the CIR case. The second application was used to show that the structure of the Laguerre-Gamma series expansion can be used to construct an estimator of the FPT probability density function when only a sample of FPT data is available and not even the moments of the underlying random variable are known. Finally, with the third application, we have set forth an approximated maximum likelihood estimation which enlightens how Laguerre-Gamma series expansion can be used to perform parameter estimation, in the case where a sample of FPT data is available, and, unlike in the previous scenario, the moments are known.

7.2 Future Works

Similar to the previous section, following the conceptual subdivision of this thesis, the discussion on future works shall be divided in three main parts.

7.2.1 Direct and Indirect Method

The models proposed in this thesis are local. Due to the large distances between the stations considered in our dataset, we were able to assume spatio-temporal independence and to proceed with our methodology on each single station separately. However, when the network of stations at hand

is dense, spatial dependence of weather events might need to be taken into account. This problem could be addressed by combining our procedure with more complex stochastic space-time models. For example, in looking for spatial patterns, hidden Markov models (Robertson et al., 2004, 2006; Hughes and Guttorp, 1994) suppose the existence of hidden states as a link between synoptic atmospheric measurements and the local precipitation process. In our particular case, tools such as the latter could be employed to model spatial dependence between parameters.

Additionally, a more in depth statistical analysis of the dependence between the rainfall amounts and the corresponding wet spell length should be carried out. For instance, it would be of interest to investigate the possibility of modelling the rainfall depths conditioned to the length of the wet spell with the Lerch family. This is currently subject of ongoing research by us.

7.2.2 Copula Like Models for Discrete Random Vectors

Multivariate Extension We strongly believe that the proposed methodology can be fully extended to the multivariate case since the key results related to I-projections (and, more in general, ϕ-projections) are not limited to the bivariate case. A first effort in this context is merely notational, as to delineate a smoothly reading framework for copula arrays of arbitrary dimension. Afterwards, a fundamental step is to provide a multivariate version of the Sklar-like decomposition found in Proposition 4.4.1. To this aim, it seems that only minor modifications of its proof must be set forth. Furthermore, a multivariate generalisation of the odds ratios on which we based our concept of dependence structure should be sought. Fortunately, multidimensional odds ratios seem to appear in the literature and, naturally, in addition to them, conditional odds ratios (see, e.g., Rudas, 2018, Section 6.2) should also be taken into consideration, as can be for example inferred from a careful read of the recent work Perrone et al. (2024). Unsurprisingly, as hinted in the latter reference, a multivariate version of the IPFP (see Barthélemy and Suesse, 2018, Section 2.1, for an example) would maintain the multidimensional and conditional odds ratios constant when used to compute a copula array. For what concerns the statistical aspects, parametric models of the form in (4.47) would need to be developed. Considering multivariate families of classical copulas should do the job. Finally, note that the results stated in Section 5.3.3 of Chapter 5 could be used to derive

the limiting distributions of I-Projections on multivariate Fréchet classes. Indeed, it should suffice to consider an arbitrary dimensional version of the vectorisation operator to be able to employ those results into the array scenario, as exemplified by the last application of Chapter 5, modulo the increased notational effort. Hence, they could be used to investigate the asymptotics of nonparametric and parametric estimators of copula arrays.

Non Rectangular Supports The assumption of rectangular support could be replaced by an alternative postulate for the support of p (based for instance on domain knowledge). We briefly describe the practical challenges that would need to be addressed to adapt the proposed statistical modelling methodology. Following a (possibly multivariate version of) Proposition 4.4.1, one would need to verify the existence of a pmf with uniform margins that shares the same support as p. In this direction, a multivariate result analogous to Proposition 4.2.3 would be needed. To the best or our knowledge, it seems that a detailed reading of Csima (1970) should provide the necessary theory. Then, a suitable smoothed estimator of p, similar to (4.39), would need to be proposed so that the corresponding empirical copula pmf of the form in (4.40) could be practically computed via the IPFP. Furthermore, parametric models of the form in (4.47) with matching support for p would need to be developed. Finally, the results stated in Section 5.3.3 of Chapter 5, already mentioned in the item above, could also be used to investigate the asymptotics of nonparametric and parametric estimators of copula arrays with possibly non rectangular support, albeit with some minor required modifications.

The extensions to the multivariate and non rectangular support case are subject of a paper currently under writing.

Smoothing The proposed inference procedures rely on the initial smoothing in (4.39), which aims to ensure numerical convergence of the IPFP with "high probability". Our choice of smoothing is arbitrary, and alternative smoothing strategies should be empirically investigated.

Additional Simulations From a statistical practice perspective, the asymptotic results in Section 4.5 could be used to derive asymptotic confidence

intervals for various nonparametric and parametric estimates. With additional implementation effort, fitting and goodness-of-fit testing could be extended to multiparameter parametric families $\mathcal{J}$ in (4.47). This is reserved for a future project. Our focus on families $\mathcal{J}$ built via (4.48) was purely for practical convenience due to the available functionalities in R. Other classes of parametric families $\mathcal{J}$ were discussed in Section 7 of Geenens (2020).

Countably Infinite Supports Finally, the theoretical extension to the countably infinite support case could be tackled. In (Geenens, 2020, Section 8) we can find the following suggestion, based on a truncation and limiting argument. Let p be the pmf of a bivariate random vector (X,Y) supported on $\mathbb{N}\times\mathbb{N}$ and denote by p_N a suitable truncation of p (see the construction in Geenens, 2020, Section 8.1), which is a pmf with rectangular support $[N]\times[N]$, where $N\in\mathbb{N}^+$. Let $u_{1,N}$ and $u_{2,N}$ be uniform univariate pmfs on $[N]$. According to Proposition 4.2.1, since p_N has a rectangular support, $u_N = \mathcal{I}_{u_{1,N},u_{1,N}}(p_N)$ exists for all $N\in\mathbb{N}^+$ and, according to the initial Definition 4.1.2, it is the pmf of a bivariate random vector (U_N,V_N) having uniform margins on $\{N\} = \{\frac{1}{N+1},\dots,\frac{N}{N+1}\}$. It is easy to see that U_N and V_N both converge in distribution to a $U(0,1)$ as N goes to infinity. Ideally then, we should expect that (U_N,V_N) converges in law to a bona fide continuous copula C on $[0,1]^2$. Geenens (2020) affirms that such convergence is guaranteed by Theorem 3.1.8 in Durante and Sempi (2015). However, it does not seem that the main condition appearing in that theorem is straightforwardly satisfied by the sequence of discrete copulas associated to the sequence of copula pmfs $\{\mathcal{I}_{u_{1,N},u_{1,N}}(p_N)\}_{N\geq 1}$. Therefore, a more careful investigation of this convergence is needed. Supposing for a moment that it holds, we should then provide an explanation of how C would play the role of the representative of the dependence of (X,Y). Again, one could argue that it follows from the fact that C itself can be seen as a representation of the infinite odds ratios ω_{ij} for (i,j) in $\mathbb{N}^+\setminus\{1\}\times\mathbb{N}^+\setminus\{1\}$ of p. Finally, the biggest issue to face in this scenario seems to be, even conceptually, the obtainment of a Sklar-like decomposition between p and its margins: a uniform distribution on the integers does not exist.

An Initial Application To a Simple Daily Rainfall Generator As mentioned in the introduction, it may happen that rainfall records are too brief to conduct reliable and meaningful analyses. Hence, stochastic models are commonly used to simulate synthetic series and to generate longer alternative rainfall

scenarios that are statistically consistent with observed data. In this context, such tools are known as *rainfall generators*. One known and straightforward method of building a daily scale rainfall generator is by describing the alternation between rainy and non rainy days with a reliable pmf p_{ws} of the wet spells ws and p_{ds} of the dry spells ds, usually under their independence. Moreover, a distribution for the rainfall depths on the rainy days h, usually independent from ws, is also used, to generate the positive rainfall when needed. According to the setup seen in Chapter 2 and Chapter 3, we consider h as discrete and therefore we let p_h be the associated pmf. Then, very simply, instances of ws and ds are alternatively (independently) generated according to p_{ws} and p_{ds} (Wilks, 1999b). Given a wet spell of length $m \in \mathbb{N}^+$ and a dry spell of length $m' \in \mathbb{N}^+$, m rainfall depths are generated according to p_h for the wet days and m' zeroes are stored for the dry days.

However, if signs of associations between the aforementioned variables are present in the available data, one could use the methodology proposed in Chapter 4 to construct a more reliable rainfall generator. Let us delineate, with some degree of informality, how that could be achieved. We describe first the involved statistical model and the related fitting, then we explain the very simple algorithm.

Since ws, ds and h are theoretically supported on the integers, some sort of truncation would need be employed. However, note that, in practice, everything should work, as the support is forcibly limited to a finite number when implementing the procedure. For example, a multiple of the maximum value found in the available data. With this in mind, pardoning some lack of rigor, we can consider the bivariate pmf $p_{ws,ds}$ (resp. $p_{ws,h}$) associated to the bivariate discrete random vector (ws,ds)(resp. (ws,h), as constructed at the end of Section 3.1.1). Naturally, both these pmfs have a rectangular support, since, a priori, no values of the underlying random values have a reason to be excluded. Then, a three steps statistical model of the type in Section 4.5 could be postulated for $p_{ws,ds}$ (resp. $p_{ws,h}$) and fitted to available (ws,ds)(resp. (ws,h)) data arising from a series of daily rainfall records $\mathbf{h}$. Furthermore, the fitting could be seasonal. For instance, let $seas$ in $\{S1,S2\}$, where the latter are defined as in Chapter 3. Then, more in details, as described in Section 4.5, one could proceed as follows.

1. Estimate the univariate margins p_{ws}, p_{ds} and p_h parametrically; let $p_{ws}^{[\alpha_{seas}^{[n]}]}$, $p_{ds}^{[\beta_{seas}^{[n]}]}$ and $p_h^{[\gamma_{seas}^{[n]}]}$ be the resulting estimates. For instance, following the first two chapters, we could assume that ws, ds and h belong to the HLZ family of discrete distributions.

2. Estimate the copula pmfs $u_{ws,ds}$ and $u_{ws,h}$ parametrically; let $u_{ws,ds}^{[\theta^{[n]}]}$ and $u_{ws,h}^{[\eta^{[n]}]}$ be the resulting estimate. In this case, to chose an appropriate parametric discrete copula pmf, one could proceed with simple graphical checks or compare the values of the log-likelihood at the estimates, as hinted at the end of Section 4.7, as ad-hoc model selection procedures still need to be developed.

3. Form a final parametric estimate of $p_{ws,ds}$ and $p_{ws,h}$ as in (4.31) via an I-projection as

$$p_{ws,ds}^{[\alpha_{seas}^{[n]},\beta_{seas}^{[n]},\theta_{seas}^{[n]}]} = \mathcal{I}_{p_{ws}^{[\alpha_{seas}^{[n]}]},p_{ds}^{[\beta_{seas}^{[n]}]}}(u^{[\theta_{seas}^{[n]}]})$$

and

$$p_{ws,h}^{[\alpha_{seas}^{[n]},\gamma_{seas}^{[n]},\eta_{seas}^{[n]}]} = \mathcal{I}_{p_{ws}^{[\alpha_{seas}^{[n]}]},p_{h}^{[\gamma_{seas}^{[n]}]}}(u^{[\eta_{seas}^{[n]}]}).$$

Then, in the following, we propose a simple generalisation of the above briefly described algorithm for the rainfall generator, whose aim is to incorporate dependence between the variables involved (we omit the season for notational simplicity). For instance, fix a preceding dry spell $ds_{i-1} = l$, with $i, l \in \mathbb{N}^*$. Then

1. generate a ws_i according to $ws|ds = l$, whose pmf is easily obtained from $p_{ws,ds}^{[\alpha^{[n]},\gamma^{[n]},\theta^{[n]}]}$ and p_{ds}. Suppose $ws_i = m$, with $m \in \mathbb{N}^+$;

   (a) generate and store m instances of $h|ws = m$, whose pmf is easily obtained from $p_{ws,h}^{[\alpha^{[n]},\gamma^{[n]},\eta^{[n]}]}$ and p_h;

2. generate a ds_i according to $ds|ws = m$, whose pmf is easily obtained from $p_{ws,ds}^{[\alpha^{[n]},\gamma^{[n]},\theta^{[n]}]}$ and p_{ws}. Suppose $ds_i = m'$, with $m' \in \mathbb{N}^+$;

   (a) store m' zeroes;

3. repeat the previous steps until the preset number of generated rainfalls has been reached.

7.2.3 FPT Density Orthogonal Approximation

As explained in the beginning sections of Chapter 6, the outlined approximation method is a particular case of the more general procedure of orthogonal approximations of probability density functions. With this in mind, it should become clear that two of the presented applications could be extended to a more general setup, as in the following. Consider a rv X with unknown density g and known moments, such that a suitable reference density f is known, to which we can associate a complete family of orthonormal polynomials $\{P_j\}_{j\geq 0}$. Then, assume that $\frac{g}{f} \in L^2_f$ so that we have $\frac{g}{f} = \sum_{j=0}^{\infty} a_j P_j$. The approximant g_n for $n \in \mathbb{N}^+$ is defined as

$$g_n(x) = f(x) \sum_{j=0}^{n} a_j P_j(x), \quad x \in \text{supp}(X),$$

where we set $\text{supp}(X) = \{x \in \mathbb{R} : g(x) > 0\}$ and we recall that $a_j = \mathbb{E}[P_j(X)]$. The following two extensions should be considered. We shall report some more details for ease of explanation.

Approximated Acceptance Rejection Method Further suppose that the approximant $\hat{g}_n$ is positive. Recall that it clearly satisfies

$$\frac{\hat{g}_n(x)}{f(x)} = \sum_{j=0}^{n} a_j P_j(x) = p_n(x), \quad x \in \text{supp}(X). \tag{7.1}$$

Since p_n is a polynomial for any $n > 0$, the right hand side of (6.92) is unbounded if $\text{supp}(X)$ is. In the latter case, we have provided in Section 6.7 a suitable modification of the standard acceptance-rejection method with the aim of "approximately" sampling from X using (7.1). In the following, suppose $\hat{X}_n$ be the rv with pdf $\hat{g}_n$ over $\text{supp}(X)$. A key step in the procedure was to exploit the known behaviour for large t of the FPT pdf. Assuming a similar knowledge for X, by only slightly modifying the last step we can summarise a "more general" procedure in the following way:

i) find a constant C such that $\mathbb{P}(X > C) \leq \varepsilon$, for a fixed, small $\varepsilon > 0$;

ii) for $t \leq C$ apply the classical acceptance-rejection method using the ratio

$$\frac{\hat{f}_{\hat{X}_n|C}(x)}{\tilde{g}_{G|C}(x)} \leq M \quad \text{where} \quad M = \frac{\mathbb{P}(\hat{X}_n \leq C)}{\mathbb{P}(G \leq C)} \max_{x \in [0,C]} p_n(x) \tag{7.2}$$

and

$$\tilde{g}_{G|C}(x) = \frac{g(x)}{\mathbb{P}(G \leq C)} \mathbf{1}_{(0,C]}(x), \quad \hat{f}_{\hat{X}_n|C}(x) = \frac{\hat{g}_n(x)}{\mathbb{P}(\hat{X}_n \leq C)} \mathbf{1}_{(0,C]}(x); \quad (7.3)$$

iii) for $t > C$ sample from a truncated rv $\bar{X}$ whose behaviour should be arbitrarily close to the behaviour of X in its tails, as ε decreases and C increases.

Obviously, the first fundamental step in this extension would be a clarification and subsequent formalisation of the statement in step iii) above. Then, seemingly minor adjustments to the proofs at the end of Section 6.7 are needed.

Remark 7.2.1. Suppose $\mathrm{supp}(X)$ is bounded, without loss of generality we can take $\mathrm{supp}(X) = (a,b)$ with a and b in $\mathbb{R}$. In this case, we can just apply step ii) of the above procedure. That is, for $t \in (a,b)$, apply the classical acceptance-rejection method using the ratio

$$\frac{\hat{f}_{\hat{X}_n}(x)}{\tilde{g}(x)} \leq M, \quad x \in (a,b), \quad \text{where} \quad M = \max_{t \in (a,b)} p_n(t). \quad (7.4)$$

□

Approximated Maximum Likelihood Estimation In Section 6.7.3 we have presented simulation experiments of an approximated maximum likelihood procedure. In the following, we try to explain the steps needed to ground theoretically the proposed method, using the present broader setup. Suppose that the density of X has a parametric form $g(\cdot;\theta)$ (unknown as well), where $\theta \in \Theta$ and Θ is an open subset of $\mathbb{R}^m$ for a positive integer m. Let $\theta_0 \in \Theta$ be the true unknown parameter, that is such that $g(\cdot) = g(\cdot;\theta_0)$. Now, as usual in classical statistics, suppose to be in possession of n independent copies $X_1, X_2, \ldots, X_n$ of X. Notice that, in the following, n is now the sample size, unlike in the previous paragraph. Our goal is to estimate θ_0 through maximum likelihood, by overcoming the fact that g is not known in a closed form through the orthogonal series expansion we have described. From now on, we shall make a discussion taking inspiration from Aït-Sahalia (2002), which provided a similar construction in a more complex scenario. Define the classical log-likelihood as

$$\ell_n(\theta|X) = \sum_{i=1}^{n} \log g(X_i;\theta)$$

and the relative classical maximum likelihood estimator as

$$\hat{\theta}_n = \arg\sup_{\theta\in\Theta} \ell_n(\theta|X).$$

Since $g(\cdot;\theta)$ is not available, we consider its approximant g_J as in the following

$$g_J(x;\theta) = f(x)\sum_{j=0}^{J} \mathbb{E}[P_j(X)]_\theta P_j(x), \qquad x\in\operatorname{supp}(X),$$

where J in $\mathbb{N}^+$ is the truncation parameter. Note that we removed the hat from the approximation of g to avoid confusion with the maximum likelihood estimator and, for similar reasons, that we shall refer to the truncation parameter with J. The subscript θ in the writing $\mathbb{E}[P_j(X)]_\theta$ denotes that the dependence of g_J on θ is found in the known moments $\mathbb{E}[X^i]_\theta$ for $i\in\mathbb{N}$. Indeed, $\mathbb{E}[P_j(X)]_\theta = \sum_{i=0}^{j} b_{i,j}\mathbb{E}[X^i]_\theta$, where $b_{i,j}$ are the coefficients of P_j. Then, define the approximated log likelihood as

$$\ell_n^J(\theta|X) = \sum_{i=1}^{n} \log g_J(X_i;\theta)$$

and the relative approximated maximum likelihood estimator as

$$\hat{\theta}_n^J = \arg\sup_{\theta\in\Theta} \ell_n^J(\theta|X).$$

To proceed, it is necessary to prove that as J tends to infinity, $\hat{\theta}_n^J$ will converge (in a suitable way) to the classical estimator $\hat{\theta}_n$ (which is not computable) and then that it is possible to find a sequence J_n tending to infinity, as the sample size grows to infinity, such that $\hat{\theta}_n^{J_n}$ will converge (in a suitable way) to the true parameter θ_0. Let us make a series of (maybe too strict) assumptions, both on the convergence of the partial sums $g_J(\cdot;\theta)$ and on the regularity of $g(\cdot;\theta)$. They are based on a series of assumptions found in Aït-Sahalia (2002).

Assumption 1 (convergence of the partial sums)**.** *The convergence of* $g_J(x;\theta)$ *to* $g(x;\theta)$ *is uniform in* x *over* $\operatorname{supp}(X)$ *and uniform in* θ *over* Θ*.*

Assumption 2 (regularity)**.** *The true unknown density* g *satisfies the usual regularity conditions for the consistency and asymptotic normality of the classical MLE.*

Assumption 3 (positivity of the partial sums)**.** *The roots of the polynomials* $\{P_j\}$ *are such that, for any* $J>0$*, the approximant* g_J *is positive on at least a finite number of intervals contained in* $\mathrm{supp}(X)$*.*

Assumption 4 (tightness)**.** *For every* $\epsilon>0$ *there exist a compact* $K_\epsilon \subset \mathbb{R}$ *such that* $\mathbb{P}(X \in K_\epsilon) \geq 1-\epsilon$*.*

Then, we should be able to state the following theorem, whose proof is currently under writing. The arrow $\xrightarrow{p}$ indicates convergence in probability.

Theorem 7.2.1. *Let* P_{θ_0} *be the probability measure having density* $g(\cdot) = g(\cdot;\theta_0)$*. Then, under Assumptions 1, 2 and 4:*

1. *Fix the sample size* n*. Then, as* $J \to \infty$*, we have that* $\hat{\theta}_n^J \xrightarrow{p} \hat{\theta}_n$ *under* P_{θ_0}*.*
2. *As* $n \to \infty$*, a sequence* J_n *tending to infinity can be found such to deliver any convergence rate of* $\hat{\theta}_n^{J_n}$ *to* $\hat{\theta}_n$*. In particular, there exists* J_n *tending to infinity such that* $\hat{\theta}_n^{J_n}$ *and* $\hat{\theta}_n$ *will share the same asymptotic distribution.*

A next important step would need to provide a data-dependent choice of the truncation parameter J. Indeed, our stopping criteria are aimed at preserving numerical stability, and, in the current scenario where a sample is available, do not take into account the fact that a too large J would probably increase the variance of the approximated estimator $\hat{\theta}_n^J$. Finally, extensive numerical experiments would be needed to investigate the finite sample size behaviour of this procedure.

Bibliography

AghaKouchak, A. (2014). Entropy-copula in hydrology and climatology. *Journal of Hydrometeorology 15*(5), 2176–2189.

Agnese, C., G. Baiamonte, and C. Cammalleri (2014). Modelling the occurrence of rainy days under a typical mediterranean climate. *Advances in Water Resources 64*, 62–76.

Agnese, C., G. Baiamonte, C. Cammalleri, D. Cat Berro, S. Ferraris, and L. Mercalli (2012). Statistical analysis of inter-arrival times of rainfall events for italian sub-alpine and mediterranean areas. *Adv Sci Res 8*, 171–177.

Agnese, C., G. Baiamonte, E. Di Nardo, S. Ferraris, and T. Martini (2022). Modelling the frequency of interarrival times and rainfall depths with the poisson hurwitz-lerch zeta distribution. *Fractal and Fractional 6*(9).

Agresti, A. (2013). *Categorical Data Analysis* (Third ed.). Wiley.

Aït-Sahalia, Y. (2002). Maximum likelihood estimation of discretely sampled diffusions: A closed-form approximation approach. *Econometrica 70*(1), 223–262.

Aksenov, S. V. and M. A. Savageau (2005). Some properties of the lerch family of discrete distributions.

Ali, M. and D. Silvey (1966). A general class of coefficients of divergence of one distribution from another. *Journal of the Roy. Statist. Soc. Ser. B 28*, 131–140.

Alili, L., P. Patie, and J. L. Pedersen (2005). Representations of the first hitting time density of an Ornstein-Uhlenbeck process. *Stochastic Models 21*(4), 967–980.

Antunes, P., S. S. Ferreira, D. Ferreira, C. Nunes, and J. T. Mexia (2020). Estimation in additive models and anova-like applications. *Journal of Applied Statistics 0*(0), 1–10.

Asmussen, S., P.-O. Goffard, and P. J. Laub (2019). *Orthonormal Polynomial Expansions and Lognormal Sum Densities*, pp. 127–150. World Scientific.

Bacharach, M. (1965). Estimating nonnegative matrices from marginal data. *International Economic Review 6*(3), 294–310.

Badía, F. G., C. Sangüesa, and A. Federgruen (2021). Log-concavity of compound distributions with applications in operational and actuarial models. *Probab. Engrg. Inform. Sci. 35*(2), 210–235.

Baets, B. D., N. E. C. Verhoest, H. Vernieuwe, and S. Vandenberghe (2008). A continuous rainfall model based on vine copulas. *Nonlin. Processes Geophys. 15*, 761–772.

Baiamonte, G., C. Agnese, C. Cammalleri, E. Di Nardo, S. Ferraris, and T. Martini (2024). Applying different methods to model dry and wet spells at daily scale in a large range of rainfall regimes across europe. *Advances in Statistical Climatology, Meteorology and Oceanography 10*(1), 51–67.

Baldi, P., L. Caramellino, and M. Rossi (2020). Large deviations of conditioned diffusions and applications. *Stochastic Processes and their Applications 130*(3), 1289–1308.

Barthélemy, J. and T. Suesse (2018). mipfp: An R package for multidimensional array fitting and simulating multivariate Bernoulli distributions. *Journal of Statistical Software, Code Snippets 86*(2), 1–20.

Basu, A., H. Shioya, and C. Park (2011, 06). *Statistical Inference: The Minimum Distance Approach*. Chapman and Hall/CRC.

Bateman, H. and A. Erdélyi (1953). *Higher Transcendental Functions*, Volume 1. New York: McGraw-Hill.

Bélisle, C. J. P. (1992). Convergence theorems for a class of simulated annealing algorithms on $\mathbf{R}^d$. *J. Appl. Probab. 29*(4), 885–895.

Belt, H. J. and A. C. den Brinker (1997). Optimal parametrization of truncated generalized Laguerre series. In *1997 IEEE International Conference on Acoustics, Speech, and Signal Processing*, Volume 5, pp. 3805–3808. IEEE.

Ben-Israel, A. and T. N. E. Greville (2003). *Generalized Inverses: Theory and Applications* (2 ed.). CMS Books in Mathematics. New York, NY: Springer.

Bender, E. A. and E. R. Canfield (1996). Log-concavity and related properties of the cycle index polynomials. *J. Combin. Theory Ser. A 74*(1), 57–70.

Berger, A. and C. Goossens (1983, 01). Persistence of wet and dry days at uccle (belgium). *International Journal of Climatology - INT J CLIMATOL 3*, 21–34.

Bernardara, P., C. De Michele, and R. Rosso (2007). A simple model of rain in time: An alternating renewal process of wet and dry states with a fractional (non-gaussian) rain intensity. *Atmospheric research 84*(4), 291–301.

Berro, D. C., C. Agnese, G. Baiamonte, L. Mercalli, and S. Ferraris (2019). Modelling the frequency distribution of inter-arrival times from daily precipitation time-series in northwest italy. *Hydrology Research* (1), 339–357.

Bertini, L. and L. Passalacqua (2008). Modelling interest rates by correlated multi-factor CIR-like processes. *arXiv preprint arXiv:0807.3898*.

Borodin, A. N. and P. Salminen (2002). *Handbook of Brownian motion—facts and formulae* (Second ed.). Probability and its Applications. Birkhäuser Verlag, Basel.

Brewer, J. (1978). Kronecker products and matrix calculus in system theory. *IEEE Transactions on Circuits and Systems 25*(9), 772–781.

Broniatowski, M. and A. Keziou (2006). Minimization of ϕ-divergences on sets of signed measures. *Studia Scientiarum Mathematicarum Hungarica 43*(4), 403–442.

Broniatowski, M. and A. Keziou (2009). Parametric estimation and tests through divergences and the duality technique. *Journal of Multivariate Analysis 100*(1), 16–36.

Brossard, J. and C. Leuridan (2018). *Iterated Proportional Fitting Procedure and Infinite Products of Stochastic Matrices*, pp. 75–117. Cham: Springer International Publishing.

Brualdi, R. A. (2006). *Combinatorial Matrix Classes*. Encyclopedia of Mathematics and its Applications. Cambridge University Press.

Buishand, T. A. (1977). Stochastic modelling of daily rainfall sequences. *Communications Agricultural University Wageningen* (77-3), 212.

Buishand, T. A. (1978). Some remarks on the use of daily rainfall models. *J Hydrol 36*, 295–308.

Buonocore, A., L. Caputo, G. D'Onofrio, and E. Pirozzi (2015). Closed-form solutions for the first-passage-time problem and neuronal modeling. *Ricerche di Matematica 64*(2), 421–439.

Buonocore, A., A. G. Nobile, and L. M. Ricciardi (1987). A new integral equation for the evaluation of first-passage-time probability densities. *Advances in Applied Probability 19*(4), 784–800.

Caloiero, T. and R. Coscarelli (2020, 09). Analysis of the characteristics of dry and wet spells in a mediterranean region. *Environmental Processes 7*, 1–11.

Casella, G. and R. L. Berger (2002). *Statistical Inference*. Duxbury Press.

Casella, G., C. P. Robert, and M. T. Wells (2004). Generalized accept-reject sampling schemes. *Lecture Notes-Monograph Series*, 342–347.

Chakrabarty, D. and S. Khanna (2021). Better and simpler error analysis of the Sinkhorn–Knopp algorithm for matrix scaling. *Mathematical Programming 188*, 395–407.

Charalambides, C. A. (1977). On the generalized discrete distributions and the Bell polynomials. *Sankhyā Ser. B 39*(1), 36–44.

Charalambides, C. A. (2002). *Enumerative combinatorics*. CRC Press Series on Discrete Mathematics and its Applications. Chapman & Hall/CRC, Boca Raton, FL.

Chatfield, C. (1966). Wet and dry spells. *Weather 21*, 308–310.

Chib, S. and E. Greenberg (1995). Understanding the Metropolis-Hastings algorithm. *The american statistician 49*(4), 327–335.

Chin, E. H. and J. F. Miller (1980). On the conditional distribution of daily precipitation amounts. *Monthly Weather Review 108*, 1462–1464.

Chowdhury, R. K. and S. Beecham (2013). Characterization of rainfall spells for urban water management. *Int. J. Climatol. 33*, 959–967.

Comtet, L. (1970). *Analyse combinatoire. Tomes I, II*, Volume 5. Presses Universitaires de France, Paris.

Conceição, K. S., F. Louzada, M. G. Andrade, and E. S. Helou (2017). Zero-modified power series distribution and its hurdle distribution version. *Journal of Statistical Computation and Simulation 87*(9), 1842–1862.

Cox, D. (1970). *Renewal theory*. Methuen & Co.

Cox, D. and D. Hinkley (1974). *Theoretical Statistics*. London: Chapman and Hall/CRC.

Cox, J. C., J. E. Ingersoll, and S. A. Ross (1985). A theory of the term structure of interest rates. *Econometrica 53*(2), 385–407.

Csima, J. (1970). Multidimensional stochastic matrices and patterns. *Journal of Algebra 14*(2), 194–202.

Csiszár, I. (1967). Information-type measures of difference of probability distributions and indirect observations. *Studia Sci. Mathem. Hungarica 2*, 299–318.

Csiszár, I. (1975). I-divergence geometry of probability distributions and minimization problems. *The Annals of Probability 3*(1), 146 – 158.

Csiszár, I. and P. Shields (2004). *Information Theory and Statistics: A Tutorial*. Now publishers.

de Amo, E., M. Díaz Carrillo, F. Durante, and J. Fernández Sánchez (2017, August). Extensions of subcopulas. *Journal of Mathematical Analysis and Applications 452*(1), 1–15.

de Groen, M. M. and H. H. G. Savenije (2006). A monthly interception equation based on the statistical characteristics of daily rainfall. *Water Resour Res 42*, W12417.

Deni, M. S. and A. A. Jemain (2009). Mixed log series geometric distribution for sequences of dry days. *Atmospheric Research 92*, 236–243.

Deni, M. S., A. A. Jemain, and K. Ibrahim (2010). The best probability models for dry and wet spells in peninsular malaysia during monsoon seasons. *Int. J. Climatol. 30*, 1194–1205.

Di Nardo, E. and G. D'Onofrio (2021). A cumulant approach for the first-passage-time problem of the Feller square-root process. *Appl. Math. Comput. 391*, 125707, 13.

Di Nardo, E., G. D'Onofrio, and T. Martini (2023). Approximating the first passage time density from data using generalized Laguerre polynomials. *Communications in Nonlinear Science and Numerical Simulation 118*, 106991.

Di Nardo, E., G. D'Onofrio, and T. Martini (2024). Orthogonal gamma-based expansion for the cir's first passage time distribution. *Applied Mathematics and Computation 480*, 128911.

Di Nardo, E. and G. Guarino (2022). kstatistics: Unbiased estimates of joint cumulant products from the multivariate Faà Di Bruno's formula. *The R Journal 14*, 208–228.

Di Nardo, E., G. Guarino, and D. Senato (2008). A unifying framework for k-statistics, polykays and their multivariate generalizations. *Bernoulli 14*(2), 440–468.

Di Nardo, E., A. G. Nobile, E. Pirozzi, and L. M. Ricciardi (2001). A computational approach to first-passage-time problems for Gauss–Markov processes. *Advances in Applied Probability 33*(2), 453–482.

Di Nardo, E. and D. Senato (2006). An umbral setting for cumulants and factorial moments. *European Journal of Combinatorics 27*(3), 394 – 413.

Diggle, P. J. and P. Hall (1986). The selection of terms in an orthogonal series density estimator. *Journal of the American Statistical Association 81*(393), 230–233.

Ditlevsen, S. and P. Lansky (2006, Jun). Estimation of the input parameters in the feller neuronal model. *Phys. Rev. E 73*, 061910.

Dobi-Wantuch, I., J. Mika, and L. Szeidl (2000). Modelling wet and dry spells with mixture distributions. *Meteorol. Atm. Phys. 73*, 245–256.

D'Odorico, P., L. Ridolfi, A. Porporato, and I. Rodriguez-Iturbe. Preferential states of seasonal soil moisture: The impact of climate fluctuations. *Water Resources Research 36*(8), 2209–2219.

D'Onofrio, G., P. Lansky, and E. Pirozzi (2018). On two diffusion neuronal models with multiplicative noise: The mean first-passage time properties. *Chaos: An Interdisciplinary Journal of Nonlinear Science 28*(4).

D'Onofrio, G., C. Macci, and E. Pirozzi (2018). Asymptotic results for first-passage times of some exponential processes. *Methodology and Computing in Applied Probability 20*(4), 1453–1476.

Doob, J. L. (1949). Heuristic approach to the Kolmogorov-Smirnov theorems. *The Annals of Mathematical Statistics*, 393–403.

Durante, F. and E. Perrone (2020). Stochastic dependence with discrete copulas. In *Book of Short Papers SIS 2020*, pp. 1344–1349. Pearson.

Durante, F. and C. Sempi (2015). *Principles of Copula Theory*. Boca Raton, Florida: CRC Press.

Efromovich, S. (2010). Orthogonal series density estimation. *Wiley Interdisciplinary Reviews: Computational Statistics 2*(4), 467–476.

Eger, S. (2016). Identities for partial Bell polynomials derived from identities for weighted integer compositions. *Aequationes Math. 90*(2), 299–306.

El Hafyani, M. and K. El Himdi (2022). A comparative study of geometric and exponential laws in modelling the distribution of daily precipitation durations. In *12th International Conference on Environmental Science and Technology*, Volume 1006 of *IOP Conf. Ser.: Earth Environ. Sci.*, pp. 012005.

Ernst, O. G., A. Mugler, H.-J. Starkloff, and E. Ullmann (2012). On the convergence of generalized polynomial chaos expansions. *ESAIM Math. Model. Numer. Anal. 46*(2), 317–339.

Faugeras, O. (2017). Inference for copula modeling of discrete data: a cautionary tale and some facts. *Dependence Modeling 5*(1), 121–132.

Feller, W. (1951). Two Singular Diffusion Problems. *Annals of Mathematics 54*(1), 173–182.

Feller, W. (1952). The parabolic differential equations and the associated semi-groups of transformations. *Annals of Mathematics*, 468–519.

Fitzpatrick, P. (2009). *Advanced Calculus*. Pure and applied undergraduate texts. American Mathematical Society.

Forman, J. and M. Sørensen (2008). The Pearson diffusions: A class of statistically tractable diffusion processes. *Scandinavian Journal of Statistics 35*(3), 438–465.

Foufoula-Georgiou, E. and D. P. Lettenmaier (1986). Continuous-time versus discrete-time point process models for rainfall occurrence series. *Water Resources Research 22*(4), 531–542.

Foufoula-Georgiou, E. and D. P. Lettenmaier (1987). A markov renewal model for representing daily rainfall sequences. *Water Resources Research 23*, 875–884.

Funaro, D. (1992). *Polynomial approximation of differential equations*, Volume 8 of *Lecture Notes in Physics.* Springer-Verlag, Berlin.

Gabriel, K. R. and J. Neumann (1962). A markov chain model for daily rainfall occurrence at tel aviv. *Q J R Meteorol Soc 88*, 90–95.

Geenens, G. (2020). Copula modeling for discrete random vectors. *Dependence Modeling 8*(1), 417–440.

Geenens, G. (2023). Towards a universal representation of statistical dependence. Preprint.

Geenens, G. (2024). (re-)reading sklar (1959)—a personal view on sklar's theorem. *Mathematics 12*, 380.

Geenens, G., I. Kojadinovic, and T. Martini (2024). On the differentiability of ϕ-projections in the discrete finite case. Preprint.

Genest, C., K. Ghoudi, and L.-P. Rivest (1995). A semiparametric estimation procedure of dependence parameters in multivariate families of distributions. *Biometrika 82*, 543–552.

Genest, C. and J. Nešlehová (2007). A primer on copulas for count data. *The Astin Bulletin 37*, 475–515.

Genest, C. and B. Rémillard (2008). Validity of the parametric bootstrap for goodness-of-fit testing in semiparametric models. *Annales de l'Institut Henri Poincaré: Probabilités et Statistiques 44*, 1096–1127.

Gietl, C. and F. Reffel (2013). Continuity of f-projections on discrete spaces. In F. Nielsen and F. Barbaresco (Eds.), *Geometric Science of Information*, Berlin, Heidelberg, pp. 519–524. Springer Berlin Heidelberg.

Gietl, C. and F. Reffel (2017). Continuity of f-projections and applications to the iterative proportional fitting procedure. *Statistics 51*(3), 668–684.

Giorno, V., P. Lánskỳ, A. Nobile, and L. Ricciardi (1988). Diffusion approximation and first-passage-time problem for a model neuron: III. A birth-and-death process approach. *Biological cybernetics 58*, 387–404.

Giorno, V., A. G. Nobile, and L. M. Ricciardi (1990). On the asymptotic behaviour of first-passage-time densities for one-dimensional diffusion processes and varying boundaries. *Advances in Applied Probability 22*(4), 883–914.

Giorno, V., A. G. Nobile, L. M. Ricciardi, and L. Sacerdote (1986). Some remarks on the rayleigh process. *Journal of Applied Probability 23*(2), 398–408.

Giorno, V., A. G. Nobile, L. M. Ricciardi, and S. Sato (1989). On the evaluation of first-passage-time probability densities via non-singular integral equations. *Advances in Applied Probability 21*(1), 20–36.

Goffard, P.-O. and P. J. Laub (2020). Orthogonal polynomial expansions to evaluate stop-loss premiums. *Journal of Computational and Applied Mathematics 370*, 112648.

Going-Jaeschke, A. and M. Yor (2003, 04). A survey and some generalizations of Bessel processes. *Bernoulli 9*(2), 313–349.

Golub, G. and C. V. Loan (2013). *Matrix Computations - 4th Edition*. Philadelphia, PA: Johns Hopkins University Press.

Goodman, L. and W. Kruskal (1954). Measures of association for cross classifications. *Journal of the American Statistical Association 49*(268), 732–764.

Goodman, L. A. (1979). Simple models for the analysis of association in cross-classifications having ordered categories. *Journal of the American Statistical Association 74*(367), 537–552.

Goodman, L. A. (1985). The Analysis of Cross-Classified Data Having Ordered and/or Unordered Categories: Association Models, Correlation Models, and Asymmetry Models for Contingency Tables With or Without Missing Entries. *The Annals of Statistics 13*(1), 10 – 69.

Gottlieb, D. and S. A. Orszag (1977). *Numerical analysis of spectral methods: theory and applications*. SIAM.

Gradshteyn, I. S. and I. M. Ryzhik (1994). *Table of integrals, series, and products* (Fifth ed.). Academic Press, Inc., Boston, MA.

Green, J. R. (1970). A generalized probability model for sequences of wet and dry days. *Monthly Weather Review 98*(3), 238–241.

Grimaldi and Serinaldi (2007). Fully nested 3-copula: Procedure and application on hydrological data. *JOURNAL OF HYDROLOGIC ENGINEERING 12*, 412–430.

Gupta, P. L., R. C. Gupta, S.-H. Ong, and H. Srivastava (2008). A class of hurwitz–lerch zeta distributions and their applications in reliability. *Applied Mathematics and Computation 196*(2), 521–531.

Gupta, P. L., R. C. Gupta, and R. C. Tripathi (1995). Inflated modified power series distributions with applications. *Comm. Statist. Theory Methods 24*(9), 2355–2374.

Gupta, P. L., R. C. Gupta, and R. C. Tripathi (1997). On the monotonic properties of discrete failure rates. *J Stat Planning Inference 65*, 255–268.

Gupta, R. C. (1974). Modified power series distribution and some of its applications. *Sankhyā Ser. B 36*(3), 288–298.

Gutiérrez-Jáimez, R., P. Román-Román, and F. Torres-Ruiz (1995). A note on the Volterra integral equation for the first-passage-time probability density. *Journal of Applied Probability 32*(3), 635–648.

Hall, D. (2023). *bignum: Arbitrary-Precision Integer and Floating-Point Mathematics*. R package version 0.3.2.

Hall, P. (1980). Estimating a density on the positive half line by the method of orthogonal series. *Annals of the Institute of Statistical Mathematics 32*, 351–362.

Hassairi, A. and M. Zarai (2004). Characterization of the cubic exponential families by orthogonality of polynomials. *The Annals of Probability 32*(3B), 2463–2476.

Herrmann, S. and C. Zucca (2019). Exact simulation of the first-passage time of diffusions. *Journal of Scientific Computing 79*, 1477–1504.

Heyde, C. C. (1963). On a property of the lognormal distribution. *J. Roy. Statist. Soc. Ser. B 25*, 392–393.

Hille, E. (1926). On laguerre's series. *Proceedings of the National Academy of Sciences 12*(12), 265–269.

Hofert, M., I. Kojadinovic, M. Mächler, and J. Yan (2018). *Elements of Copula Modeling with R*. Springer.

Hofert, M., I. Kojadinovic, M. Mächler, and J. Yan (2022). *copula: Multivariate dependence with copulas*. R package version 1.1-0.

Holland, P. W. and Y. J. Wang (1987). Regional dependence for continuous bivariate densities. *Communications in Statistics - Theory and Methods 16*(1), 193–206.

Hope, A. C. (1968). A simplified monte carlo significance test procedure. *Journal of the Royal Statistical Society: Series B (Methodological) 30*(3), 582–598.

Hughes, J. P. and P. Guttorp (1994). Incorporating spatial dependence and atmospheric data in a model of precipitation. *Journal of Applied Meteorology and Climatology 33*(12), 1503 – 1515.

Hui Wan, X. Z. and E. M. Barrow (2005). Stochastic modelling of daily precipitation for canada. *Atmosphere-Ocean 43*(1), 23–32.

Idel, M. (2016). A review of matrix scaling and Sinkhorn's normal form for matrices and positive maps.

Jaynes, E. T. (1957). Information theory and statistical mechanics. *Phys. Rev. 106*, 620–630.

Jiménez-Gamero, M., R. Pino-Mejías, V. Alba-Fernández, and J. Moreno-Rebollo (2011). Minimum ϕ-divergence estimation in misspecified multinomial models. *Computational Statistics and Data Analysis 55*(12), 3365–3378.

Johnson, J. R. (2024). *nloptr: R Interface to NLopt*. R package version 2.0.3.

Johnson, N. L., S. Kotz, and N. Balakrishnan (2000). *Continuous multivariate distributions*. Wiley.

Johnson, W. P. (2002). The curious history of faà di bruno's formula. *The American Mathematical Monthly 109*(3), 217–234.

Kalantari, B., I. Lari, F. Ricca, and B. Simeone (2008). On the complexity of general matrix scaling and entropy minimization via the RAS algorithm. *Mathematical Programming 112*, 371–401.

Karlin, S. and H. M. Taylor (1981). *A second course in stochastic processes*. Academic Press, Inc., New York-London.

Kassambara, A. (2023). *ggpubr: 'ggplot2' based publication ready plots*. R package version 0.6.0.

Kateri, M. (2014). *Contingency Table Analysis*. New York: Birkhäuser.

Katz, R. W. (1977). Precipitation as a chain-dependent process. *Journal of Applied Meteorology 16*, 671–676.

Katz, R. W. and M. B. Parlange (1995). Generalizations of chain-dependent processes: Application to hourly precipitation. *Water Resources Research 31*, 1331–1341.

Keilson, J. and H. Gerber (1971). Some results for discrete unimodality. *Journal of the American Statistical Association 66*(334), 386–389.

Kendall, M. and J. Gibbons (1990). *Rank correlation methods* (Fifth ed.). New York: Oxford University Press.

Kent, J. T. (1980). Eigenvalue expansions for diffusion hitting times. *Zeitschrift fur Wahrscheinlichkeitstheorie und Verwandte Gebiete 52*(3), 309–319.

Kent, J. T. (1982). The spectral decomposition of a diffusion hitting time. *The Annals of Probability 10*(1), 207–219.

Keziou, A. (2003). Dual representation of phi-divergences and applications. *Comptes Rendus Mathematique 336*(10), 857–862.

Kojadinovic, I. and T. Martini (2024). Copula-like inference for discrete bivariate distributions with rectangular supports. *Electronic Journal of Statistics 18*(1), 2571–2619.

Kojadinovic, I. and J. Yan (2011). A goodness-of-fit test for multivariate multiparameter copulas based on multiplier central limit theorems. *Statistics and Computing 21*(1), 17–30.

Kolesarova, A., R. Mesiar, J. Mordelova, and C. Sempi (2006). Discrete copulas. *IEEE Transactions on Fuzzy Systems 14*(5), 698–705.

Kostal, L., P. Lansky, and O. Pokora (2011). Variability measures of positive random variables. *PLoS One 6*(7), e21998.

Kostal, L. and O. Pokora (2012). Nonparametric estimation of information-based measures of statistical dispersion. *Journal of Mathematics and Computer Science 14*(7), 1221–1233.

Kottegoda, N. T. and R. Rosso (1997). *Statistics, probability, and reliability for civil and environmental engineers*. USA: McGraw-Hill.

Kruithof, J. (1937). Telefoonverkeers rekening (calculation of telephone traffic). *De Ingenieur 52*, E15–E25.

Laio, F., A. Porporato, L. Ridolfi, and I. Rodriguez-Iturbe (2001). Plants in water-controlled ecosystems: active role in hydrologic processes and response to water stress. ii: Probabilistic soil moisture dynamics. *Adv Water Resour 24*(7), 707–723.

Lansky, P., L. Sacerdote, and F. Tomassetti (1995, Oct). On the comparison of Feller and Ornstein-Uhlenbeck models for neural activity. *Biological Cybernetics 73*(5), 457–465.

Lefebvre, M. (2002, 05). Geometric brownian motion as a model for river flows. *Hydrological Processes 16*, 1373 – 1381.

Liese, F. and I. Vajda (1987). *Convex statistical distances*. Leipzig: Teubner.

Liew, K. and S.-H. Ong (2012, 12). Stochastic orderings and parameter estimation for the hurwitz-lerch zeta distribution. *Malaysian Journal of Science 31*, 148–158.

Liew, K. W., S. H. Ong, and K. K. Toh (2020). The poisson-stopped hurwitz–lerch zeta distribution. *Communications in Statistics - Theory and Methods 0*(0), 1–15.

Lin, W. and J. E. Zhang (2022). The valid regions of Gram–Charlier densities with high-order cumulants. *Journal of Computational and Applied Mathematics 407*, 113945.

Lindsay, B. (1994). Efficiency Versus Robustness: The Case for Minimum Hellinger Distance and Related Methods. *The Annals of Statistics 22*(2), 1081 – 1114.

Linetsky, V. (2004). Computing hitting time densities for CIR and OU diffusions: Applications to mean-reverting models. *Journal of Computational Finance 7*, 1–22.

Little, R. J. A. and M. Wu (1991). Models for contingency tables with known margins when target and sampled populations differ. *Journal of the American Statistical Association 86*(413), 87–95.

Lung, T.-H. (1998). *Approximations for skewed probability densities based on Laguerre series and biological applications*. North Carolina State University.

Martin, E., U. Behn, and G. Germano (2011). First-passage and first-exit times of a Bessel-like stochastic process. *Physical Review E 83*(5), 051115.

Martin, R. J., M. J. Kearney, and R. V. Craster (2019, mar). Long- and short-time asymptotics of the first-passage time of the Ornstein–Uhlenbeck and other mean-reverting processes. *Journal of Physics A: Mathematical and Theoretical 52*(13), 134001.

Masoliver, J. and J. Perelló (2012). First-passage and escape problems in the Feller process. *Physical review E 86*(4), 041116.

Masoliver, J. and J. Perelló (2014). First-passage and escape problems in the Feller process. *Phys. Rev. E 86(4-1)*, 041116.

McCullagh, P. (1987). *Tensor methods in statistics*. Monographs on Statistics and Applied Probability. Chapman & Hall, London.

McMahon, T. and R. Srikanthan (2001). Stochastic generation of annual, monthly and daily climate data: A review. *Hydrology and Earth System Sciences 5*, 653–670.

Melbourne, J. (2020). Strongly convex divergences. *Entropy 22*(11), 1–20.

Mercadier, M. and F. Strobel (2021, November). A one-sided vysochanskii-petunin inequality with financial applications. *European Journal of Operational Research 295*(1), 374–377.

Mijatović, A., M. R. Pistorius, and J. Stolte (2015). Randomisation and recursion methods for mixed-exponential Lévy models, with financial applications. *Journal of Applied Probability 52*(4), 1076–1096.

Miller, S. (1960). Comparative social mobility. *Current Sociology 9*(1), 1–61.

Morales, D., L. Pardo, and I. Vajda (1995). Asymptotic divergence of estimates of discrete distributions. *Journal of Statistical Planning and Inference 48*(3), 347–369.

Mosteller, F. (1968). Association and estimation in contingency tables. *Journal of the American Statistical Association 63*(321), 1–28.

Nelsen, R. (2006). *An Introduction to Copulas*. New-York: Springer. Second edition.

Neslehova, J. (2004, May). *Dependence of Non-Continuous Random Variables*. Phd thesis, Carl von Ossietzky University Oldenburg, Oldenburg, Germany.

Nesterov, Y. (2004). *Introductory lectures on convex optimization: A basic course*. Springer.

Nishii, R. (1996). Orthogonal functions of inverse gaussian distributions. In *Lifetime Data: Models in Reliability and Survival Analysis*, pp. 243–250. Springer.

Nobile, A. G., L. M. Ricciardi, and L. Sacerdote (1985). Exponential trends of first-passage-time densities for a class of diffusion processes with steady-state distribution. *Journal of Applied Probability 22*(3), 611–618.

Oakley, S. J. (1990). *Orthogonal polynomials in the approximation of probability distributions*. Phd thesis.

Øksendal, B. (1998). *Stochastic Differential Equations: An Introduction with Applications*. Hochschultext / Universitext. Springer.

Ong, S. H., K. K. Toh, and K. W. Liew (2020). The poisson-stopped hurwitz–lerch zeta distribution. *Communications in Statistics - Theory and Methods*.

Onof, C., B. D. Baet, S. Vandenberghe, and N. E. C. Verhoest (2011). A comparative copula-based bivariate frequency analysis of observed and simulated storm events: A case study on bartlett-lewis modeled rainfall. *WATER RESOURCES RESEARCH 47*.

P. L. Gupta, R. C. G. and R. C. Tripathi (1997). On the monotonic properties of discrete failure rates. *Journal of Statistical Planning and Inference 65*(2), 255–268.

Paek, J., M. Pollanen, and K. Abdella (2023). A stochastic weather model for drought derivatives in arid regions: A case study in qatar. *Mathematics 11*(7).

Paton, E. (2022). Intermittency analysis of dry spell magnitude and timing using different spell definitions. *Journal of Hydrology 608*, 127645.

Pauwels, E. (1987). Smooth first-passage densities for one-dimensional diffusions. *Journal of Applied Probability 24*(2), 370–377.

Perrone, E., R. Fontana, and F. Rapallo (2024). Multi-way contingency tables with uniform margins. In J. Ansari, S. Fuchs, W. Trutschnig, M. A. Lubiano, M. Á. Gil, P. Grzegorzewski, and O. Hryniewicz (Eds.), *Combining, Modelling and Analyzing Imprecision, Randomness and Dependence*, Cham, pp. 349–356. Springer Nature Switzerland.

Perrone, E., L. Solus, and C. Uhler (2019). Geometry of discrete copulas. *Journal of Multivariate Analysis 172*, 162–179. Dependence Models.

Petrie, M. D. and N. A. Brunsell (2012). The role of precipitation variability on the ecohydrology of grasslands. *Ecohydrology 5*(3), 337–345.

Plackett, R. (1965). A class of bivariate distributions. *Journal of the American Statistical Association 60*(310), 516–522.

Platen, E. and P. E. Kloeden (1992). *Numerical solution of stochastic differential equations*. Springer-Verlag.

Porporato, A., G. Vico, and P. A. Fay (2006). Superstatistics of hydroclimatic fluctuations and interannual ecosystem productivity. *Geophysical Research Letters 33*(15).

Pretzel, O. (1980). Convergence of the Iterative Scaling Procedure for Non-Negative Matrices. *Journal of the London Mathematical Society s2-21*(2), 379–384.

Provost, S. B. and H.-T. Ha (2016). Distribution approximation and modelling via orthogonal polynomial sequences. *Statistics 50*(2), 454–470.

Pukelsheim, F. (2014). Biproportional scaling of matrices and the iterative proportional fitting procedure. *Annals of Operations Research 215*, 269–283.

R Core Team (2024). *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing.

Racsko, P., L. Szeidl, and M. Semenov (1991). A serial approach to local stochastic weather models. *Ecol. Modell. 57*, 27–41.

Ramos-Alarcón, F. and V. Kontorovich (2013). First-passage time statistics of Markov gamma processes. *J. Franklin Inst. 350*(7), 1686–1696.

Ratan, R. and V. Venugopal (2013). Wet and dry spell characteristics of global tropical rainfall. *Water Resources Research 49*(6), 3830–3841.

Raymond, F., A. Ullmann, P. Camberlin, P. Drobinski, and C. C. Smith (2016). Extreme dry spell detection and climatology over the mediterranean basin during the wet season. *Geophysical Research Letters 43*(13), 7196–7204.

Read, T. and N. Cressie (1988). *Goodness-of-Fit Statistics for Discrete Multivariate Data*. Springer Series in Statistics. New York, NY: Springer-Verlag.

Redner, S. (2001). *A Guide to First-Passage Processes*. Cambridge University Press.

Ricciardi, L. M., A. Di Crescenzo, V. Giorno, and A. G. Nobile (1999). An outline of theoretical and algorithmic approaches to first passage time problems with applications to biological modeling. *Math. Japon. 50*(2), 247–322.

Ricciardi, L. M., L. Sacerdote, and S. Sato (1984). On an integral equation for first-passage-time probability densities. *Journal of Applied Probability 21*(2), 302–314.

Ridolfi, L., P. D'Odorico, and F. Laio (2011). *Noise-Induced Phenomena in the Environmental Sciences*. Cambridge University Press.

Robert, C. P. and G. Casella (2004). *Monte Carlo Statistical Methods*. Springer.

Robertson, A. W., S. Kirshner, and P. Smyth (2004). Downscaling of daily rainfall occurrence over northeast brazil using a hidden markov model. *Journal of Climate 17*(22), 4407 – 4424.

Robertson, A. W., S. Kirshner, P. Smyth, S. P. Charles, and B. C. Bates (2006). Subseasonal-to-interdecadal variability of the australian monsoon over north queensland. *Quarterly Journal of the Royal Meteorological Society 132*(615), 519–542.

Rodriguez-Iturbe, I., D. R. Cox, and V. Isham (1987). Some models for rainfall based on stochastic point process. *Proc R Soc Lond A 410*, 269–288.

Rodriguez-Iturbe, I., D. R. Cox, and V. Isham (1988). A point model for rainfall: further developments. *Proc R Soc Lond A 417*, 283–298.

Román-Román, P., J. Serrano-Pérez, and F. Torres-Ruiz (2023). fptdapprox: Approximation of first-passage-time densities for diffusion processes, 2023. *R package version 2.5*.

Rosler, U. (1980). Unimodality of passage times for one-dimensional strong Markov processes. *The Annals of Probability 8*(4), 853.

Rothblum, U. G. and H. Schneider (1989). Scalings of matrices which have prespecified row sums and column sums via optimization. *Linear Algebra and its Applications 114-115*, 737–764. Special Issue Dedicated to Alan J. Hoffman.

Rudas, T. (2018). *Lectures on Categorical Data Analysis*. New York: Springer.

Rüschendorf, L. (1987). Projections of probability measures. *Statistics 18*(1), 123–129.

Sansone, G. (1991). *Orthogonal functions*. Dover Publications, Inc., New York.

Sato, K.-i. (2013). *Lévy processes and infinitely divisible distributions*, Volume 68 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge. Translated from the 1990 Japanese original, Revised edition of the 1999 English translation.

Semenov, M. A., R. J. Brooks, E. M. Barrow, and C. W. Richardson (1998). Comparison of the wgen and lars-wg stochastic weather generators for diverse climates. *Climate Research 10*(2), 95–107.

Seneta, E. (2004). Fitting the variance-gamma model to financial data. Volume 41A, pp. 177–187. Stochastic methods and their applications.

Serinaldi, F. (2009). A multisite daily rainfall generator driven by bivariate copula-based mixed distributions. *Journal of Geophysical Research: Atmospheres 114*(D10).

Shohat, J. A. (1935). On the development of functions in series of orthogonal polynomials. *Bulletin of the American Mathematical Society 41*(2), 49–82.

Siegert, A. J. F. (1951, Feb). On the first passage time probability problem. *Phys. Rev. 81*, 617–623.

Simonoff, J. (1995). Smoothing categorical data. *Journal of Statistical Planning and Inference 47*(1), 41–69.

Sinkhorn, R. (1974). Diagonal equivalence to matrices with prescribed row and column sums. ii. *Proceedings of the American Mathematical Society 45*(2), 195–198.

Sklar, A. (1959). Fonctions de répartition à n dimensions et leurs marges. *Publications de l'Institut de Statistique de l'Université de Paris 8*, 229–231.

Smith, C. E. (1991). A Laguerre series approximation for the probability density of the first passage time of the Ornstein-Uhlenbeck process. In *Noise in Physical Systems and 1/f Fluctuations*, pp. 389–392. Tokyo: Ohmsha.

Stechmann, S. N. and J. D. Neelin (2014). First-passage-time prototypes for precipitation statistics. *Journal of the Atmospheric Sciences 71*(9), 3269 – 3291.

Steutel, F. W. and K. van Harn (1979). Discrete analogues of self-decomposability and stability. *The Annals of Probability 7*(5), 893–899.

Stute, W., W. Gonzáles Manteiga, and M. Presedo Quindimil (1993). Bootstrap based goodness-of-fit tests. *Metrika 40*, 243–256.

Szegő, G. (1975). *Orthogonal polynomials* (Fourth ed.). American Mathematical Society, Providence, R.I.

Tagliani, A. (2000). Existence and stability of a discrete probability distribution by maximum entropy approach. *Applied Mathematics and Computation 110*(2), 105–114.

Tagliani, A. (2001). Discrete probability distributions and moment problem: numerical aspects. *Applied Mathematics and Computation 119*(1), 47–56.

Trefethen, L. and D. Bau III (1997). *Numerical Linear Algebra*. SIAM, Philadelphia.

Tsybakov, A. (2008). *Introduction to Nonparametric Estimation*. Springer Series in Statistics. Springer New York.

Tweedie, M. C. K. (1957). Statistical properties of inverse Gaussian distributions. I, II. *Ann. Math. Statist. 28*, 362–377, 696–705.

Vajda, I. and E. van der Meulen (2005). On minimum divergence adaptation of discrete bivariate distributions to given marginals. *IEEE Transactions on Information Theory 51*(1), 313–320.

Van Der Hoeven, A. (2005). A comparison of information criteria for model selection. *Journal of Statistical Computation and Simulation 75*, 37–45.

van der Vaart, A. (1998). *Asymptotic statistics*. Cambridge University Press.

White, H. (1982). Maximum likelihood estimation of misspecified models. *Econometrica 50*(1), 1–25.

Wilks, D. S. (1999a). Interannual variability and extreme-value characteristics of several stochastic daily precipitation models. *Agr. Forest Meteorol. 93*(3), 153–169.

Wilks, D. S. (1999b). The weather generation game: A review of stochastic weather models. *Progress in Physical Geography 23*, 329–357.

Wilks, S. S. (1938). The large-sample distribution of the likelihood ratio for testing composite hypotheses. *Ann. Math. Statist. 9*(1), 60–62.

Wilson, G. A. and A. Wragg (1973). Numerical methods for approximating continuous probability density functions, over $[0, \infty]$, using moments. *Journal of the Institute of Mathematics and its Applications 12*, 165–173.

Wiscombe, W., R. D. A. Cahalan, and A. Marshak (1994). Multifractal characterizations of nonstationarity and intermittency in geophysical fields: observed, retrieved, or simulated. *J Geophys Res 99*.

Wolfram Research, Inc. (2023). *Mathematica, Version 13.3*. Champaign, IL: Wolfram Research, Inc. Available at: https://www.wolfram.com/mathematica/.

Wolfram Research, Inc. (2024). General characteristics of laguerrel. Available at: https://functions.wolfram.com/HypergeometricFunctions/LaguerreL3General/21/01/02/03/ [Accessed: 2024-08-21].

Yang, L., C. L. E. Franzke, and Z. Fu (2020). Power-law behaviour of hourly precipitation intensity and dry spell duration over the united states. *International Journal of Climatology 40*(4), 2429–2444.

Yor, M. (1992). On some exponential functionals of Brownian motion. *Advances in Applied Probability 24*(3), 509–531.

Yu, Y. (2009). On the entropy of compound distributions on nonnegative integers. *IEEE Trans. Inform. Theory 55*(8), 3645–3650.

Yule, G. (1912). On the methods of measuring association between two attributes. *Journal of the Royal Statistical Society 75*(6), 579–652.

Zheng, Z., L. Wei, J. Hämäläinen, and O. Tirkkonen (2012). Approximation to distribution of product of random variables using orthogonal polynomials for lognormal density. *CoRR abs/1203.3288*.

Zolina, O., C. Simmer, K. Belyaev, S. K. Gulev, and P. Koltermann (2013). Changes in the duration of european wet and dry spells during the last 60 years. *J. Climate 26*(6), 2022–2047.