

AperTO - Archivio Istituzionale Open Access dell'Università di Torino

A "super" folk theorem for dynastic repeated games

This is the author's manuscript

Original Citation:

Availability:

This version is available <http://hdl.handle.net/2318/1723685> since 2020-01-16T19:37:50Z

Published version:

DOI:10.1007/s00199-007-0293-9

Terms of use:

Open Access

Anyone can freely access the full text of works made available as "Open Access". Works made available under a Creative Commons license can be used according to the terms and conditions of said license. Use of all other works requires consent of the right holder (author or publisher) if not exempted from copyright protection by the applicable law.

(Article begins on next page)

A “Super” Folk Theorem for Dynastic Repeated Games*

LUCA ANDERLINI
(Georgetown University)

DINO GERARDI
(Yale University)

ROGER LAGUNOFF
(Georgetown University)

January 2006

Abstract. We analyze “dynastic” repeated games. A stage game is repeatedly played by successive generations of finitely-lived players with dynastic preferences. Each individual has preferences that replicate those of the infinitely-lived players of a standard discounted infinitely-repeated game. When all players observe the past history of play, the standard repeated game and the dynastic game are equivalent

In our model all players live one period and *do not* observe the history of play that takes place before their birth, but instead receive a *private message* from their immediate predecessors.

Under very mild conditions, when players are sufficiently patient, *all feasible payoff vectors* (including those below the *minmax* of the stage game) can be sustained as a Sequential Equilibrium of the dynastic repeated game with private communication. The result applies to any stage game for which the standard Folk Theorem yields a payoff set with a non-empty interior.

We are also able to characterize entirely when a Sequential Equilibrium of the dynastic repeated game can yield a payoff vector not sustainable as a Subgame Perfect Equilibrium of the standard repeated game. For this to be the case it must be that the players’ equilibrium beliefs violate a condition that we term “Inter-Generational Agreement.”

JEL CLASSIFICATION: C72, C73, D82.

KEYWORDS: Dynastic Repeated Games, Private Communication, Folk Theorem.

ADDRESS FOR CORRESPONDENCE: Luca Anderlini, LA2@GEORGETOWN.EDU, Department of Economics, Georgetown University, 37th and O Streets, Washington DC 20057, USA.

*This paper was previously circulated with the title “The Folk Theorem in Dynastic Repeated Games.” We are grateful to Leonardo Felli, Navin Kartik, David Levine, Stephen Morris, Michele Piccione and Andrew Postlewaite, and to seminar audiences at Bocconi, Cambridge, Chicago, Columbia, Edinburgh, Georgetown, Leicester, LSE, Oxford, Rome (La Sapienza), Rutgers, SAET-Vigo, Stanford, SUNY-Albany, UCL, UC-San Diego, Venice and Yale for helpful feedback. All errors are our own.

1. Introduction

1.1. Motivation

The theory of infinitely repeated games is a cornerstone of modern economic analysis. Besides its tremendous effect on the theory of games itself, it has pervaded economic applications that stretch from industrial organization, to political economy, to development economics.

We depart from the standard repeated game model by considering strategic interactions among *dynasties*. This is appealing whenever the ongoing strategic interaction can be thought of as one between entities that outlive any individual player. For instance, firms, in some cases, may be appropriately modeled as “placeholders” for successive generations of decision-makers (owners or managers). The interaction between political parties also seems suited to this view, as are some instances of religious strife and tribal or family disputes.

We analyze “dynastic” repeated games. A stage game is repeatedly played by successive generations of finitely-lived players with dynastic preferences. Each individual has preferences that replicate those of the infinitely-lived players of a standard discounted infinitely repeated game: his payoff is a weighted sum of his current (stage) payoff and that of his successors in the same dynasty. The case of almost-infinite patience in the standard model is re-interpreted as a world with almost-complete altruism (towards successors in the same dynasty).

Suppose that all players observe the past history of play. An application of the one-shot deviation principle is enough to show that the standard repeated game and the dynastic game are equivalent, both in terms of Subgame-Perfect Equilibrium (henceforth SPE) strategies and consequently of equilibrium payoffs. This is so *regardless* of the “demographic structure” of the dynastic game. In other words, the equivalence holds regardless of the process governing birth and death of successive generations of players, provided that extinction of any dynasty is ruled out.

We focus on the case in which the players are *not* able to observe directly the history of play that took place before their birth (or more loosely, before they began participating in the game). Instead, our players rely on *messages* from their predecessors to fathom the past.

In fact, in many situations of economic interest, postulating that communication is the only way to convey information about the past seems a good assumption. For instance, if the dynasties are competing firms in an oligopolistic market, the current decision makers may of course have access to accounts and a host of other past records. But, in some cases at least, it seems compelling to say that these should be modeled as messages from past players, rather than direct observation of the past.

In this paper we characterize the set of Sequential Equilibria (henceforth SE) of dynastic repeated games in which all players live one period, and have to rely on *private* messages from their immediate predecessors about the previous history of play. Our results would remain intact if *public* messages were allowed alongside private ones.

Of course, one would like to know what happens when the demographics of the model are generalized, and what would be the effects of the presence of an imperfect but non-null

physical footprint of the past history of play, which we assume away. We return to these and other issues in Section 10 below, which concludes the paper.

1.2. Results

Our results indicate that there are stark differences between the dynastic repeated games with private communication that we analyze here and the standard repeated game model.

We show that, under very general conditions, in a dynastic repeated game with private communication there are equilibria that yield the dynasties (possibly even *all* dynasties) a payoff below their minmax value in the stage game. In short, we find that in a very broad class of games (which includes the $n \geq 3$ -player version of the Prisoners’ Dilemma), as the players become more and more altruistic (patient), *all* (interior) payoff vectors in the convex hull of the payoffs of the stage game can be sustained as an SE of the dynastic repeated game with private communication. For some stage games, the ability to sustain all (interior) payoffs also implies that, in some of the SE of the dynastic repeated game with private communication, one or more players will receive payoffs that are *higher* than the maximum they could receive in any SPE of the standard repeated game.

Each player in a dynasty controls directly (through his own actions in the stage game) only a fraction of his dynastic payoff; the rest is determined by the action profiles played by subsequent cohorts of players. As the players become more and more altruistic towards their descendants, the fraction of their dynastic payoff that they control directly shrinks to zero. This reassures us that it is *not impossible* that a player’s payoff in the dynastic repeated game be below his minmax level in the stage game. The fact that this can indeed happen in an SE of the dynastic repeated game, in which the interests of players in the same dynasty by assumption *are aligned*, is the subject of this study.

A closer look at some of the SE that we find in this paper reveals an interesting fact. There is an obvious sense in which, following some histories of play, in equilibrium, beliefs are “mismatched” across players of the same cohort. This phenomenon, in turn, can be traced back to the fact that following some histories of play, the structure of the equilibrium does not permit a player to communicate to his successor some information about the past that is relevant for future play. Below, we formalize this notion via a condition that we term “Inter-Generational Agreement.” The mismatch in beliefs is equivalent to saying that an SE violates Inter-Generational Agreement.

We are able to show that the mismatching phenomenon we have just described characterizes fully the difference between the set of SE of the dynastic repeated game and the set of SPE of the standard repeated game. Any SPE of the standard repeated game can be replicated as an SE of the dynastic repeated game that displays Inter-Generational Agreement. Conversely any SE of the dynastic repeated game that yields a payoff vector that is not sustainable as an SPE of the standard repeated game must violate Inter-Generational Agreement.

2. Outline

The outline of the rest of the paper is as follows. In Section 3 we lay down the notation and details of the Standard Repeated Game and of the Dynastic Repeated Game. In Section 4 we define what constitutes an SE for the Dynastic Repeated Game. In Section 5 we present our first result asserting that all payoffs that can be sustained as SPE of the Standard Repeated Game can be sustained as SE of the Dynastic Repeated Game. Section 6 is devoted to our first “extended” Folk Theorem for the Dynastic Repeated Game. This result applies to the case of three or more dynasties. In Section 7 we present our second extended Folk Theorem. The result we state there applies to the case of four dynasties or more. In Section 8 we report a result that completely characterizes the features of an SE of the Dynastic Repeated Game that make it possible to sustain payoff vectors that are *not* sustainable as SPE of the Standard Repeated Game. Section 9 reviews some related literature, and Section 10 concludes.

For ease of exposition, and for reasons of space, no formal proofs appear in the main body of the paper. The main ingredients (public randomization devices, strategies and trembles) for the proofs of our two extended Folk Theorems (Theorems 2 and 3) appear in Appendix A and Appendix B respectively. The complete proof of Theorem 4 appears in Appendix C. In the numbering of equations, Lemmas etc. a prefix of “A,” “B” or “C” or means that the item is located in the corresponding Appendix.

A technical addendum to the paper contains the rest of the formal proofs.¹ In particular, the technical addendum contains the proof of Theorem 1 and the analysis of “consistency” and “sequential rationality” that closes the proofs of Theorems 2 and 3. In the numbering of equations, Lemmas etc. a prefix of “T.X” means that the item is located in Section X of the technical addendum.²

3. The Model

3.1. A Standard Repeated Game

We first describe a standard, n -player repeated game. We will then augment this structure to describe the dynastic repeated game with communication from one cohort to the next. The standard repeated game structure is of course familiar. We set it up below simply to establish the basic notation.

Since in our main results below we will make explicit use of *public randomization* devices we build these in our notation right from that start.

The stage game is described by the array $G = (A, u; I)$ where $I = \{1, \dots, n\}$ is the set of players, indexed by i . The n -fold cartesian product $A = \times_{i \in I} A_i$ is the set of pure action profiles $a = (a_1, \dots, a_n) \in A$, assumed to be finite. Stage game payoffs are defined by $u = (u_1, \dots, u_n)$ where $u_i : A \rightarrow \mathbb{R}$ for each $i \in I$. Let $\sigma_i \in \Delta(A_i)$ denote a mixed strategy

¹The technical addendum is available via <http://www.georgetown.edu/faculty/1a2/Folktheorem.htm>.

²For ease of reference, pages and footnotes are also numbered separately in the technical addendum. Their numbers have a prefix “T.”

for i , with σ denoting the profile $(\sigma_1, \dots, \sigma_n)$.³ The symbol $\sigma_i(a_i)$ represents the probability of pure action a_i given by σ_i , so that for any $a = (a_1, \dots, a_n)$, with a minor abuse of notation, we can let $\sigma(a) = \prod_{i \in I} \sigma_i(a_i)$ denote the probability of pure action profile a that σ induces. The corresponding payoff to player i , is defined in the usual way: $u_i(\sigma) = \sum_{a \in A} \sigma(a) u_i(a)$. Dropping the i subscript and writing $u(\sigma)$ gives the entire profile of payoffs.

Throughout the rest of the paper, we denote by V the convex hull of the set of payoff vectors from pure strategy profiles in G . We let $\text{int}V$ denote the (relative) interior of V .⁴

We assume that all players observe the outcome of a public randomization device. We model this as a random variable \tilde{x} taking values in a *finite* set X , so that its distribution is a point in $\Delta(X)$. Each player can condition his choice of mixed strategy on the realization of \tilde{x} , denoted by x .⁵ Player i chooses a mixed strategy $\sigma_i \in \Delta(A_i)$ for each $x \in X$.

In the repeated game, time is discrete and indexed by $t = 0, 1, 2, \dots$. The pure (realized) action profile played at time t is denoted by a^t . In each period there is a new public randomization device \tilde{x}^t . These devices are i.i.d. across periods; we write \tilde{x} to indicate the random variable of which all the \tilde{x}^t s are independent “copies.” We refer to \tilde{x} as the action-stage randomization device.⁶ A history of length $t \geq 1$, denoted by h^t , is an object of the form $(x^0, a^0, \dots, x^{t-1}, a^{t-1})$. The “initial” history h^0 is of course the “null” history \emptyset . The set of all possible histories of length t is denoted by H^t , while $H = \cup_{t=0}^{\infty} H^t$ denotes the collection of all possible histories of play.

A strategy for player i in the repeated game is denoted by g_i and can be thought of as a collection $(g_i^0, g_i^1, \dots, g_i^t, \dots)$ with each g_i^t a function from $H^t \times X$ into $\Delta(A_i)$. The profile of repeated game strategies is $g = (g_1, \dots, g_n)$, while g^t indicates the time- t profile (g_1^t, \dots, g_n^t) .

Given a profile g , recursing forward as usual, we obtain a probability distribution over action profiles a^0 played in period 0, then a probability distribution over profiles a^1 to be played in period 1, then in period 2, and so on without bound, so that we have a distribution over the profile of actions to be played in every $t \geq 0$. Of course, this forward recursion yields a probability distribution $\mathcal{P}(g)$ over the set of all possible sequences (a^0, \dots, a^t, \dots) (which can be seen as a probability distribution over the set of possible actual histories of play H).

The players’ common discount factor is denoted by $\delta \in (0, 1)$. Given a profile g , player

³As standard, here and throughout the paper and the technical addendum, given any finite set Z , we let $\Delta(Z)$ be the set of all probability distributions over Z .

⁴Of course, when V has full dimension (n) the relative interior of V coincides with the interior of V . In general, $\text{int}V$ is simply the set of payoff vectors that can be achieved placing *strictly positive* weight on *all* payoff vectors obtained from pure action profiles in G . We use the qualifier “relative” since we will not be making *explicit* assumptions on the dimensionality of V , although the hypotheses of many of our results will imply that V should satisfy certain dimensionality conditions that we will point out below.

⁵It should be made clear at this stage that at no point in the paper do we assume that mixed strategies are observable.

⁶Throughout, we restrict attention to action-stage randomization devices that have *full support*. That is, we assume that \tilde{x} takes *all* values in the finite set X with strictly positive probability. This seems reasonable, and keeps us away of a host of unnecessary technicalities.

i 's expected payoff in the repeated game is denoted by $v_i(g)$ and is given by⁷

$$v_i(g) = E_{\mathcal{P}(g)} \left\{ (1 - \delta) \sum_{t=0}^{\infty} \delta^t u_i(a^t) \right\} \quad (1)$$

Given a profile g and a particular finite history h^t , we recurse forward as above to find the distribution over infinite sequences (a^0, \dots, a^t, \dots) , this time *conditional* on history h^t and the realization x^t having taken place. We denote this distribution by $\mathcal{P}(g, h^t, x^t)$. This allows us to define the *continuation* payoff to player i , conditional on the pair (h^t, x^t) . This will be denoted by $v_i(g|h^t, x^t)$ and is given by

$$v_i(g|h^t, x^t) = (1 - \delta) \left\{ u_i(g^t(h^t, x^t)) + E_{\mathcal{P}(g, h^t, x^t)} \left[\sum_{\tau=t+1}^{\infty} \delta^{\tau-t} u_i(a^\tau) \right] \right\} \quad (2)$$

As before, dropping the subscript i in either (1) or (2) will indicate the entire profile.

At this point we note that, as is standard, given (1) and (2) we know that the continuation payoffs in the repeated game follow a recursive relationship. In particular we have that

$$v_i(g|h^t, x^t) = (1 - \delta) u_i(g^t(h^t, x^t)) + \delta E_{g^t(h^t, x^t), \tilde{x}^{t+1}} \{ v_i(g|h^t, x^t, g^t(h^t, x^t), \tilde{x}^{t+1}) \} \quad (3)$$

where $(h^t, x^t, g^t(h^t, x^t), \tilde{x}^{t+1})$ is the (random) history generated by the concatenation of (h^t, x^t) with the realization of the mixed strategy profile $g^t(h^t, x^t)$ and the realization of \tilde{x}^{t+1} . The expectation is then taken with respect to the realization of the mixed strategy profile $g^t(h^t, x^t)$ and of \tilde{x}^{t+1} .

A *Subgame-Perfect Equilibrium*, g^* , for the repeated game is defined in the usual way: for each i , and each finite history (h^t, x^t) , and each strategy g_i for i , we require that $v_i(g^*|h^t, x^t) \geq v_i(g_i, g_{-i}^*|h^t, x^t)$.⁸

We denote by $\mathcal{G}^S(\delta, \tilde{x})$ the set of SPE strategy profiles, and by $\mathcal{E}^S(\delta, \tilde{x})$ the set of SPE payoff profiles of the repeated game when randomization device \tilde{x} is available and the common discount factor is δ . We also let $\mathcal{E}^S(\delta) = \bigcup_{\tilde{x}} \mathcal{E}^S(\delta, \tilde{x})$ with the union being taken over the set of all possible finite random variables that may serve as public randomization devices.

The standard model of repeated play we have just sketched out may be found in a myriad of sources. See, for example, Fudenberg and Maskin (1986) and the references contained therein. Hereafter, we refer to the standard model above as the *standard repeated game*.

We conclude our discussion of the standard repeated game noting that the one-shot devi-

⁷Clearly, $v_i(g)$ depends on δ as well. To lighten the notation, this will be omitted whenever doing so does not cause any ambiguity.

⁸As is standard, here, and throughout the rest of the paper and the technical addendum, a subscript or a superscript of $-i$ indicates an array with the i -th element taken out.

ation principle gives an immediate way to check, subgame by subgame, whether a profile g^* is an SPE. Since it is completely standard, we state the following without proof.

Remark 1. *One-Shot Deviation Principle:* A profile of strategies g^* is in $\mathcal{G}^S(\delta, \tilde{x})$ if and only if for every $i \in I$, $h^t \in H$, $x^t \in X$ and any $\sigma_i \in \Delta(A_i)$ we have that

$$v_i(g^*|h^t, x^t) \geq v_i(\sigma_i, g_i^{-t*}, g_{-i}^*|h^t, x^t) \quad (4)$$

where, in keeping with the notational convention we adopted above, g_i^{-t*} stands for the profile $(g_i^{0*}, \dots, g_i^{t-1*}, g_i^{t+1*}, \dots)$.

3.2. The Dynastic Repeated Game: Full Memory

The first dynastic repeated game that we describe is a straw-man. It turns out to be equivalent, in both payoffs and strategies, to the standard repeated game. It fulfills purely an expository function in our story.

Assume that each $i \in I$ indexes an entire progeny of individuals. We refer to each of these as a *dynasty*. Individuals in each dynasty are assumed to live one period. At the end of each period t (the beginning of period $t+1$), a new individual from each dynasty — the date $(t+1)$ -lived individual — is born and replaces the date t lived individual in the same dynasty. Hence, $u_i(a^t)$ now refers to the payoff directly received by the date t individual in dynasty i . Each date t individual is altruistic in the sense that his payoff includes the discounted sum of the direct payoffs of all future individuals in the same dynasty. Date t individual in dynasty i gives weight $1 - \delta$ to $u_i(a^t)$, and weight $(1 - \delta)\delta^\tau$ to $u_i(a^{t+\tau})$ for every $\tau \geq 1$. All cohorts in all dynasties can observe directly the past history of play. So all individuals in cohort t , can observe $h^t = (x^0, a^0, \dots, x^{t-1}, a^{t-1})$. It follows that g_i^t as above can now be interpreted as the strategy of individual i in cohort t (henceforth, simply “player $\langle i, t \rangle$ ”) in the full memory dynastic repeated game.

Therefore, in the full memory dynastic repeated game, given a profile of strategies g (now interpreted as an array giving a strategy for each player $\langle i, t \rangle$) the overall payoff to player $\langle i, t \rangle$ conditional on history (h^t, x^t) can be written exactly as in (2) above.

Denote by $\mathcal{G}^F(\delta, \tilde{x})$ the set of SPE strategy profiles, and by $\mathcal{E}^F(\delta, \tilde{x})$ the set of (dynastic) SPE payoff profiles for $t = 0$ players of the dynastic repeated game when the randomization device \tilde{x} is available and the common discount factor is δ .⁹ Now observe that from Remark 1 we know that a profile g constitutes an SPE of the standard repeated game if and only if when we interpret g as a strategy profile in the full memory dynastic repeated game, no player $\langle i, t \rangle$ can gain by unilaterally deviating. Hence the equilibrium sets of the standard repeated game and of the full memory dynastic repeated game must be the same. Purely for completeness we state the following formally, but we omit any further proof.

⁹Here for SPE, and later for SE, it is immediate to show that any dynastic equilibrium payoff vector for the $t = 0$ cohort is also a (dynastic) equilibrium payoff vector for players in any subsequent ($t \geq 1$) cohort.

Theorem 0. *Dynastic Interpretation of the Standard Repeated Game:* Let any stage game G , any $\delta \in (0, 1)$ and any public randomization device \tilde{x} be given. Then the standard repeated game and the full memory dynastic repeated game are equivalent in the sense that $\mathcal{G}^S(\delta, \tilde{x}) = \mathcal{G}^F(\delta, \tilde{x})$ and $\mathcal{E}^S(\delta, \tilde{x}) = \mathcal{E}^F(\delta, \tilde{x})$.

3.3. The Dynastic Repeated Game: Private Communication

We are now ready to drop the assumption that individuals in the t -th cohort observe the previous history of play. The players’ dynastic preferences are exactly as in the full memory dynastic repeated game described in Subsection 3.2 above.

Anderlini and Lagunoff (2005) analyzed the case in which each individual in the t -cohort has to rely on *publicly observable messages* from the previous cohort for any information about the previous history of play.¹⁰

In this paper, we assume that each player $\langle i, t \rangle$ receives a *private message* from player $\langle i, t - 1 \rangle$ about the previous history of play. It should be noted at this point that our results would survive intact if we allowed *public messages as well as private ones*. All the equilibria we construct below would still be viable, with the public messages ignored. We return to this issue in Section 10 below.¹¹

For simplicity, the set of messages M_i^{t+1} available to each player $\langle i, t \rangle$ to send to player $\langle i, t + 1 \rangle$ is the set H^{t+1} of finite histories of length $t + 1$ defined above.¹² The message sent by player $\langle i, t \rangle$ will be denoted by m_i^{t+1} , so that player $\langle i, t \rangle$ is the *recipient* of message m_i^t .

We also introduce a new public randomization device in every period, which is observable to all players in the t -th cohort at the stage in which they select the messages to be sent to the $t + 1$ -th cohort.¹³

¹⁰See Section 9 below for a brief discussion of the results in Anderlini and Lagunoff (2005) and other related papers.

¹¹Dealing explicitly with both private and public messages would be cumbersome, and make our results considerably less transparent. Analyzing the model with private messages only is the most economical way to put our main point across, and hence this is how we proceed.

¹²As will be apparent from our proofs, smaller messages spaces — even ones that stay bounded in size through time — would suffice for our purposes. Taking the message spaces to coincide with the sets of previous histories of play seems the natural canonical modeling choice.

¹³Three points are worth emphasizing at this stage. First of all, in general, we consider the availability of a public randomization device (a publicly observable finite random variable) a weak assumption in just about any game. To reiterate an obvious fact, the device simply has to be available for the players to observe; whether or not the players decide to condition their play on its realization is an attribute of the *equilibrium* of the model. Second, the assumption we are now making that a second public randomization device is available seems a natural one in the model we are setting up. The players in the t -th cohort each take two successive decisions: what to play and then what to say. Given that a public randomization device is available when play is decided, it seems natural, in fact compelling, that the circumstances should be the same when messages are sent to the next cohort of players. Third, and most importantly, all our theorems survive literally unchanged (with the proof of Theorem 2 being the only one that needs modification) if we assume that both randomization devices are simultaneously observable to all players in the t -th cohort at the time they are called upon to choose their actions in the stage game. In this case, clearly a *single* randomization device would suffice.

We let a finite random variable \tilde{y}^t taking values in a finite set Y be realized in every period. A typical realization of \tilde{y}^t is denoted by y^t . We take these random variables to be mutually independent across any two time periods. We write \tilde{y} to indicate the finite random variable of which all the \tilde{y}^t are independent “copies,” and we refer to \tilde{y} as the message-stage randomization device.¹⁴

To summarize, we take the time-line of events, observations and decisions within each time period t to be as follows. At the beginning of each time period t , each player $\langle i, t \rangle$ receives a private message $m_i^t \in H^t$ from player $\langle i, t - 1 \rangle$.¹⁵ Next, the random variable \tilde{x}^t is realized. Its realization x^t is observed by all players $\langle i, t \rangle$ in the t -th cohort. After observing x^t , each $\langle i, t \rangle$ in the t -th cohort selects a mixed strategy $\sigma_i^t \in \Delta(A_i)$ for the stage game G . These choices are simultaneous. Subsequently, all players $\langle i, t \rangle$ in the t -th cohort observe the *realized* action profile a^t , which is of course the realization of the mixed strategy profile σ^t . After the profile a^t is observed, the random variable \tilde{y}^t is realized and all players $\langle i, t \rangle$ in the t -th cohort observe its realization y^t . Finally, after observing y^t , each player $\langle i, t \rangle$ in the t -th cohort selects a probability distribution $\phi_i^t \in \Delta(H^{t+1})$ over messages m_i^{t+1} in the set H^{t+1} . The realized message from this distribution is then sent to player $\langle i, t + 1 \rangle$, who observes it.¹⁶

In terms of notation, we distinguish between the *action strategy* of player $\langle i, t \rangle$, denoted by g_i^t , and the *message strategy* of player $\langle i, t \rangle$, denoted by μ_i^t . Since we take the space of possible messages that player $\langle i, t \rangle$ may receive to be $M_i^t = H^t$, formally g_i^t is exactly the *same object* as in the standard repeated game of Subsection 3.1 and the full memory dynastic repeated game of Subsection 3.2 above. In particular g_i^t takes as input a message $m_i^t \in H^t$ and a value $x^t \in X$ and returns a mixed strategy $\sigma_i^t \in \Delta(A_i)$.

The message strategy μ_i^t of player $\langle i, t \rangle$ takes as inputs a message m_i^t , the realization x^t , the realized action profile a^t , the realized value y^t , and returns the probability distribution ϕ_i^t over messages $m_i^{t+1} \in H^{t+1}$. In what follows, we will often write $\mu_i^t(m_i^t, x^t, a^t, y^t)$ to indicate the (mixed) message $\phi_i^t \in \Delta(H^{t+1})$ that player $\langle i, t \rangle$ sends player $\langle i, t + 1 \rangle$ after observing the quadruple (m_i^t, x^t, a^t, y^t) .

In denoting profiles and sub-profiles of strategies we extend the notational conventions we established for the standard repeated game and the full memory dynastic repeated game. In other words we let g_i denote the i -th dynasty profile $(g_i^0, g_i^1, \dots, g_i^t, \dots)$, while g^t will indicate the time t profile (g_1^t, \dots, g_n^t) and g the entire profile of action strategies (g_1, \dots, g_n) . Similarly, we set $\mu_i = (\mu_i^0, \mu_i^1, \dots, \mu_i^t, \dots)$, as well as $\mu^t = (\mu_1^t, \dots, \mu_n^t)$ and $\mu = (\mu_1, \dots, \mu_n)$. The pair (g, μ) therefore entirely describes the behavior of all players in the dynastic repeated game

¹⁴Throughout, we restrict attention to message-stage randomization devices that have *full support*. That is, we assume that \tilde{y} takes *all* values in the finite set Y with strictly positive probability. See also footnote 6 above.

¹⁵When $t = 0$, we assume that player $\langle i, 0 \rangle$ receives the “null” message $m_i^0 = \emptyset$.

¹⁶Notice that we are excluding the realized value of y^t from the set of histories H^{t+1} , and hence from the message space of player $\langle i, t \rangle$. This is completely without loss of generality. All our results remain intact if the set of messages available to each player $\langle i, t \rangle$ is augmented to include the realized values $\{y^\tau\}_{\tau=0}^t$. A formal proof of this claim can be obtained as a minor adaptation of the proof of Lemma T.2.1 in the technical addendum.

with private communication. Since no ambiguity will ensue, from now on we refer to this game simply as the *dynastic repeated game*.

4. Sequential Equilibrium

It is reasonably clear what one would mean by the statement that a pair (g, μ) constitutes a Weak Perfect Bayesian Equilibrium (WPBE) of the dynastic repeated game described above.

However, it is also clear that the off-the-equilibrium-path beliefs of the players may be crucial in sustaining a given profile as a WPBE. To avoid “unreasonable” beliefs, the natural route to take is to restrict attention further, to the set of SE of the model. This is a widely accepted benchmark, in which beliefs are restricted so as to be consistent with a fully-fledged, common (across players) “theory of mistakes.” We return to a discussion of this point in Section 10 below.

The original definition of Sequential Equilibrium (Kreps and Wilson 1982), of course, does not readily apply to our dynastic repeated game since we have to deal with infinitely many players.

As it turns out only a minor adaptation of the SE concept is needed to apply it to our set-up. We spell this out in detail for the sake of completeness and because we believe it makes the proofs of many of our results below considerably more transparent.

We begin with the observation that the beliefs of player $\langle i, t \rangle$ can in fact be boiled down to a simpler object than one might expect at first sight, because of the structure of the dynastic repeated game. Upon receiving message m_i^t , in principle, we would have to define the beliefs of player $\langle i, t \rangle$ over the *entire set of possible past histories of play*. However, when player $\langle i, t \rangle$ is born, an entire cohort of new players replaces the $t - 1$ -th one, and hence the real history leaves no trace other than the messages (m_1^t, \dots, m_n^t) that have been sent to cohort t . It follows that, without loss of generality, after player $\langle i, t \rangle$ receives message m_i^t we can restrict attention to his beliefs over the $n - 1$ -tuple m_{-i}^t of messages received by other players in cohort t .¹⁷ This probability distribution, specifying the beginning-of-period beliefs of player $\langle i, t \rangle$, will be denoted by $\Phi_i^{tB}(m_i^t)$ throughout the rest of the paper and the technical addendum. When the dependence of this distribution on m_i^t can be omitted from the notation without causing any ambiguity we will write it as Φ_i^{tB} . The notation $\Phi_i^{tB}(\cdot)$ will indicate the entire array of possible probability distributions $\Phi_i^{tB}(m_i^t)$ as $m_i^t \in M_i^t$ varies.

Consider now the two classes of information sets at which player $\langle i, t \rangle$ is called upon to play: the first defined by a pair (m_i^t, x^t) when he has to select a mixed strategy σ_i^t , and the second defined by a quadruple (m_i^t, x^t, a^t, y^t) when he has to select a probability distribution ϕ_i^t over the set of messages H^{t+1} .

¹⁷It should be made clear that the beliefs of player $\langle i, t \rangle$ over m_{-i}^t will in fact depend on the relative likelihoods of the actual histories of play that could generate different $n - 1$ -tuples m_{-i}^t . What we are asserting here is simply that once we know the player’s beliefs over m_{-i}^t , we have all that is necessary to check that his behavior is optimal given his beliefs.

The same argument as above now suffices to show that at the (m_i^t, x^t) information set we can again restrict attention to the beliefs of player $\langle i, t \rangle$ over the $n - 1$ -tuple m_{-i}^t of messages received by other players in cohort t . Moreover, since all players observe the same x^t and this realization is independent of what happened in the past, player $\langle i, t \rangle$ beliefs over m_{-i}^t must be the same as when he originally received message m_i^t .

Finally, at the information set identified by the quadruple (m_i^t, x^t, a^t, y^t) , we can restrict attention to the beliefs of player $\langle i, t \rangle$ over the $n - 1$ -tuple m_{-i}^{t+1} of messages that the other players in cohort t are about to send to cohort $t + 1$. Just as before, since all players are replaced by a new cohort and time- t payoffs have already been realized, this is all that could ever matter for the payoff to player $\langle i, t \rangle$ from this point on. This probability distribution, specifying the end-of-period beliefs of player $\langle i, t \rangle$, will be denoted by $\Phi_i^{tE}(m_i^t, x^t, a^t, y^t)$ throughout the rest of the paper and the technical addendum. When the dependence of this distribution on (m_i^t, x^t, a^t, y^t) can be omitted from the notation without causing any ambiguity we will write it as Φ_i^{tE} . The notation $\Phi_i^{tE}(\cdot)$ will indicate the entire array of possible probability distributions $\Phi_i^{tE}(m_i^t, x^t, a^t, y^t)$ as the quadruple (m_i^t, x^t, a^t, y^t) varies.

During the proofs of our main results, we will also need to refer to the (revised) end-of-period beliefs of player $\langle i, t \rangle$ on the $n - 1$ -tuple of messages m_{-i}^t after he observes not only m_i^t , but also (x^t, a^t, y^t) . These will be indicated by $\Phi_i^{tR}(m_i^t, x^t, a^t, y^t)$, with the arguments omitted when this does not cause any ambiguity. The notation $\Phi_i^{tR}(\cdot)$ will indicate the entire array of possible probability distributions $\Phi_i^{tR}(m_i^t, x^t, a^t, y^t)$ as the quadruple (m_i^t, x^t, a^t, y^t) varies.

Throughout the rest of the paper we refer to the array $\Phi = \{\Phi_i^{tB}(\cdot), \Phi_i^{tE}(\cdot)\}_{t \geq 0, i \in I}$ as a *system of beliefs*. Following standard terminology we will also refer to a triple (g, μ, Φ) , a strategy profile and a system of beliefs, as an *assessment*. Also following standard terminology, we will say that an assessment (g, μ, Φ) is *consistent* if the system of beliefs Φ can be obtained (in the limit) using Bayes’ rule from a sequence of completely mixed strategies that converges pointwise to (g, μ) . Since this is completely standard, for reasons of space we do not specify any further details here.

Definition 1. *Sequential Equilibrium:* An assessment (g, μ, Φ) constitutes an SE for the dynastic repeated game if and only if (g, μ, Φ) is consistent, and for every $i \in I$ and $t \geq 0$ strategy g_i^t is optimal for player $\langle i, t \rangle$ given beliefs $\Phi_i^{tB}(\cdot)$, and strategy μ_i^t is optimal for the same player given beliefs $\Phi_i^{tE}(\cdot)$.

We denote by $\mathcal{G}^D(\delta, \tilde{x}, \tilde{y})$ the set of SE strategy profiles, and by $\mathcal{E}^D(\delta, \tilde{x}, \tilde{y})$ the set of (dynastic) SE payoff profiles for $t = 0$ players of the dynastic repeated game when the randomization devices \tilde{x} and \tilde{y} are available and the common discount factor is δ .¹⁸ As previously, we let $\mathcal{E}^D(\delta) = \bigcup_{\tilde{x}, \tilde{y}} \mathcal{E}^D(\delta, \tilde{x}, \tilde{y})$, with the union ranging over all possible pairs of finite random variables \tilde{x} and \tilde{y} .

¹⁸See footnote 9 above.

5. A Basic Inclusion

The first question we ask is whether all SPE of the standard repeated game survive as SE of the dynastic repeated game. The answer is affirmative.

We pursue this question separately from our limiting (as $\delta \rightarrow 1$) Folk Theorems for the dynastic repeated game presented below for several distinct reasons.

First of all Theorem 1 below does not require any assumptions on the stage game. Second, it asserts that the SPE of the standard repeated game survive as SE of the dynastic repeated game *regardless* of the discount factor δ . Third, Theorem 1 asserts that the actual strategies that form an SPE in the standard game will be (the action component of) some SE of the dynastic repeated game, thus going beyond any statement concerning equilibrium payoffs.

Theorem 1 also fulfils a useful expository function. Running through an intuitive outline of its proof helps an initial acquaintance with some of the mechanics of the SE of the dynastic repeated game. Before going any further, we proceed with a definition and a formal statement of the result.

Definition 2. *Truthful Message Strategies:* A communication strategy μ_i^t for player $\langle i, t \rangle$ in the dynastic repeated game is said to be truthful if and only if $\mu_i^t(m_i^t, x^t, a^t, y^t) = (m_i^t, x^t, a^t)$ for all m_i^t, x^t, a^t and y^t .¹⁹ The profile μ is called truthful if all its components μ_i^t are truthful.

Theorem 1. *Basic Inclusion:* Fix a δ , an \tilde{x} and any profile $g^* \in \mathcal{G}^S(\delta, \tilde{x})$. Then for every finite random variable \tilde{y} there exists a profile μ^* of truthful message strategies such that $(g^*, \mu^*) \in \mathcal{G}^D(\delta, \tilde{x}, \tilde{y})$.

It follows that $\mathcal{E}^S(\delta) \subseteq \mathcal{E}^D(\delta)$ for every $\delta \in (0, 1)$. In other words, the set of SE payoffs for the dynastic repeated game contains the set of SPE payoffs of the standard repeated game.

The argument we use to prove Theorem 1 in the technical addendum (Sections T.3 and T.4) is not hard to outline. Begin with an action-stage strategy profile $g^* \in \mathcal{G}^S(\delta)$. Now consider a message-stage strategy profile μ^* that is “truthful” in the sense of Definition 2.

Now suppose that each player $\langle i, t \rangle$, upon receiving any message m_i^t , on or off the equilibrium path, believes that all other time- t players have received exactly the *same* message as he has. Then it is not hard to see that since the action-stage strategy profile g^* is an SPE of the standard repeated game, it will not be profitable for any player $\langle i, t \rangle$ to deviate from the prescriptions of either g_i^{t*} or μ_i^{t*} in the dynastic repeated game.

So, why should player $\langle i, t \rangle$ hold such beliefs in an SE? Suppose for instance that player $\langle i, t \rangle$ receives a message that specifies a history of play that is “far away” from the equilibrium path, say a node that requires 1,000 past action-stage deviations to be reached. Of course

¹⁹ Notice that we are defining as truthful a message strategy that *ignores* the value of y^t . This is consistent with the fact that we are excluding the realizations of \tilde{y}^t from the set of possible messages. As we remarked before, all our results would be unaffected if these were included in the players’ message spaces. See also footnote 16 above.

he needs to weigh the possibility that he really is at such node against, for instance, the possibility that his immediate predecessor has deviated from μ_i^{t-1*} and has sent him the message he observes, but no action-stage deviation has ever occurred.

In an SE, these relative likelihood considerations are embodied in the sequence of completely mixed strategies that support the beliefs, in the limit, via Bayes’ rule. The core of the argument behind Theorem 1 is to show that the sequence of completely mixed strategies converging to the equilibrium strategies can be constructed in such a way that the likelihood of a *single* past deviation from equilibrium at the message stage compared to the likelihood of *all* players in every previous cohort deviating at the action stage shrinks to zero in the limit. Of course, there is more than one way to achieve this. Our formal argument in the technical addendum relies on “trembles” defining the completely mixed strategies with the following structure. The probability of deviating at the message stage stays constant (at $\varepsilon \rightarrow 0$) through time. On the other hand, the *order* of the infinitesimal of the probability of n deviations at the action stage decreases exponentially through time.²⁰ In this way it is possible to ensure that the probability that all players in every cohort deviate at the action stage shrinks to zero slower than does the probability of a single-player deviation at the message stage. Hence, the beliefs we have described above are consistent, and Theorem 1 follows.

6. Three Dynasties or More

In this Section we present our first result asserting that the set of payoffs that are possible in a SE of the dynastic game is larger than the set of SPE payoffs of the standard repeated game, in the limit as δ approaches 1. The increase is in fact quite dramatic.

We postpone any further discussion and proceed with a formal statement of our next result and a couple of remarks on its scope. The proof of Theorem 2 begins in Appendix A, and continues in the technical addendum to the paper (Sections T.5 through T.8).

Theorem 2. *Dynastic Folk Theorem: Three Dynasties or More: Let any stage game G with three or more players be given. Assume that G is such that we can find two pure action profiles a^* and a' in A with*

$$u_i(a^*) > u_i(a') > u_i(a_i^*, a'_{-i}) \quad \forall i \in I \quad (5)$$

Then for every $v \in \text{int}V$ there exists a $\underline{\delta} \in (0, 1)$ such that $\delta > \underline{\delta}$ implies $v \in \mathcal{E}^D(\delta)$.

Remark 2. *Boundary Points: The statement of Theorem 2 is framed in terms of interior payoff vectors $v \in \text{int}V$ mostly to facilitate the comparison with the statement of Theorem 3 below dealing with the case of four dynasties or more. However, from the proof of Theorem 2 it is immediately apparent that a stronger statement is in fact true. The condition that*

²⁰In particular, in equation T.3.3 in the proof of Theorem 1 below, we set the total probability of deviations from equilibrium at the action stage from the part of player $\langle i, t \rangle$ to be equal to $\varepsilon \frac{1}{(n+1)2^{t+1}}$.

$v \in \text{int}V$ can be replaced with the weaker requirement that v is any weighted average of payoff vectors that are feasible in G , which gives strictly positive weight to $u(a^*)$. Clearly, depending on the position of $u(a^*)$ within V this may include vectors that are on the boundary of V .

Remark 3. *Generalized Prisoners’ Dilemma:* Whenever the stage game G is a version of the n -player Prisoners Dilemma (with $n \geq 3$), Theorem 2 guarantees that all interior feasible payoff vectors can be sustained as an SE of the dynastic repeated game, provided that δ is near 1.

To see this, observe that if we label C_i the “cooperate” action and D_i the “defect” action for player i , in an n -person Prisoners’ Dilemma we obviously have that $u_i(C) > u_i(D) > u_i(C_i, D_{-i})$ for every $i \in I$. Hence identifying a^* and a' of the statement of Theorem 2 with the profiles C and D respectively, immediately yields the result.

Moreover, notice that Theorem 2 applies equally immediately to any stage game that is “locally” like an n -person Prisoners’ Dilemma. In other words, the result applies to any stage game G in which we can identify any pair of profiles C and D as above, regardless of how many other actions may be available to any player, and of which payoff vectors they may yield.

Before outlining the argument behind Theorem 2, it is necessary to clear-up a preliminary issue that concerns both the proof of Theorem 2 and the proof of Theorem 3 below.

For simplicity, in both cases, we work with message spaces that are smaller than the applicable set of finite histories H^t . As Lemma T.2.1 demonstrates, enlarging message spaces from the ones we use in these two proofs back to H^t will not shrink the set of SE payoffs. This is because we can “replicate” any SE of the dynastic repeated game with restricted message spaces as an SE of the dynastic repeated game with larger message spaces by mapping each message in the smaller set to a finite set of messages in the larger message space. A choice of message in the smaller message space corresponds to a (uniform) randomization over the entire corresponding set in the larger message space. A player receiving one of the randomized messages in the larger message space acts just like the corresponding player who receives the corresponding message in the smaller message set. It is then straightforward to check that the new strategies constitute an SE of the dynastic repeated game with larger message spaces, provided of course that we started off with an SE of the dynastic game with restricted message spaces in the first place.

In the SE that we construct to prove Theorems 2 and 3, player $\langle i, t - 1 \rangle$ may want to communicate to player $\langle i, t \rangle$ that dynasty i is being punished for having deviated, but will be unable to do so in an effective way. Given that in both cases we work with message spaces that are *smaller* than H^t , one may be tempted to conclude that this is due to the fact that player $\langle i, t - 1 \rangle$ “lacks the message” to communicate effectively to his successor $\langle i, t \rangle$ that he should respond in an appropriate way. This is misleading. When the message spaces coincide with H^t , it clearly cannot be that this inability to communicate is due to a shortage of possible messages. Given the argument (the proof of Lemma T.2.1) that we sketch out

above the correct interpretation is that *in equilibrium* there is *no message* (in H^t) that player $\langle i, t \rangle$ might possibly *interpret* in the way that $\langle i, t - 1 \rangle$ would like.²¹

We are now ready for an actual outline of the proof of Theorem 2. The intuition can be divided into two parts. First, we will argue that if δ is close to one, it is possible to sustain the payoff vector $u(a^*)$ as an SE of the dynastic repeated game using $u(a')$ as “punishment” payoffs. Notice that $u(a^*)$ could well already be *below the minmax* payoff for one or more players. We call this the “local” part of the argument. Second, we will argue that via a judicious use of the action-stage randomization device it is possible to go from the local argument to a “global” one and sustain every feasible payoff vector as required by the statement of the theorem.

Our construction relies on every player $\langle i, t \rangle$ having a message space with three elements. So, set $M_i^t = \{m^*, m^A, m^B\}$ for every i and t . Notice also that the assumptions made in Theorem 2 obviously guarantee that $a_i^* \neq a_i'$ for every $i \in I$.

We begin with the local part of the argument. So, suppose that δ is close to one and that we want to implement the payoff $v = u(a^*)$ as an SE of the dynastic repeated game. The strategy of every player $\langle i, 0 \rangle$ at the action stage is to play a_i^* . The strategy of each player $\langle i, t \rangle$ with $t \geq 1$ at the action stage is to play a_i^* after receiving message m^* and to play action a_i' after receiving message m^A or message m^B . If no player deviates from the prescriptions of his equilibrium strategy at the action and message stages, then play follows a path that involves the message profile (m^*, \dots, m^*) and the action profile a^* in every period. Moreover, even after deviations, if the message profile ever returns to being equal to (m^*, \dots, m^*) , then the path of play returns to being as above. We call this the *equilibrium phase*.

Play starts in the equilibrium phase. Suppose now that some player deviates to playing action a_i' (deviations by two or more players are ignored). This triggers what we call the *temporary punishment phase*. During the temporary punishment phase all players $\langle i \in I, t \rangle$ condition the messages they send to their successors on the realization y^t of the message-stage randomization device. The message-stage randomization device takes value $y(1)$ with probability $\gamma(\delta)$ and value $y(0)$ with probability $1 - \gamma(\delta)$. If $y^t = y(1)$ then all players send message m^* , and hence play returns to the equilibrium phase. If on the other hand $y^t = y(0)$ then all players send message m^A , and play remains in the temporary punishment phase. That is all players now play a_i' , and continue to coordinate their messages as we just described. Play remains in the temporary punishment phase until the realization of the message-stage randomization device is $y(1)$, at which point play goes back to the equilibrium phase.

Any other deviation, taking place in either the equilibrium phase or in the temporary punishment phase, triggers the *permanent punishment phase*. Suppose that any player $\langle i, t \rangle$ deviates to playing any action a_i different from both a_i^* and a_i' during the equilibrium phase or that he deviates from playing a_i' during the temporary punishment phase (deviations by

²¹At this point it is legitimate of course to wonder whether the concept of “neologism-proof” equilibrium (Farrell 1993) has any impact on what we are saying here. While neologism-proofness in its current form does not apply to our model, we return to this point at some length in Section 10 below.

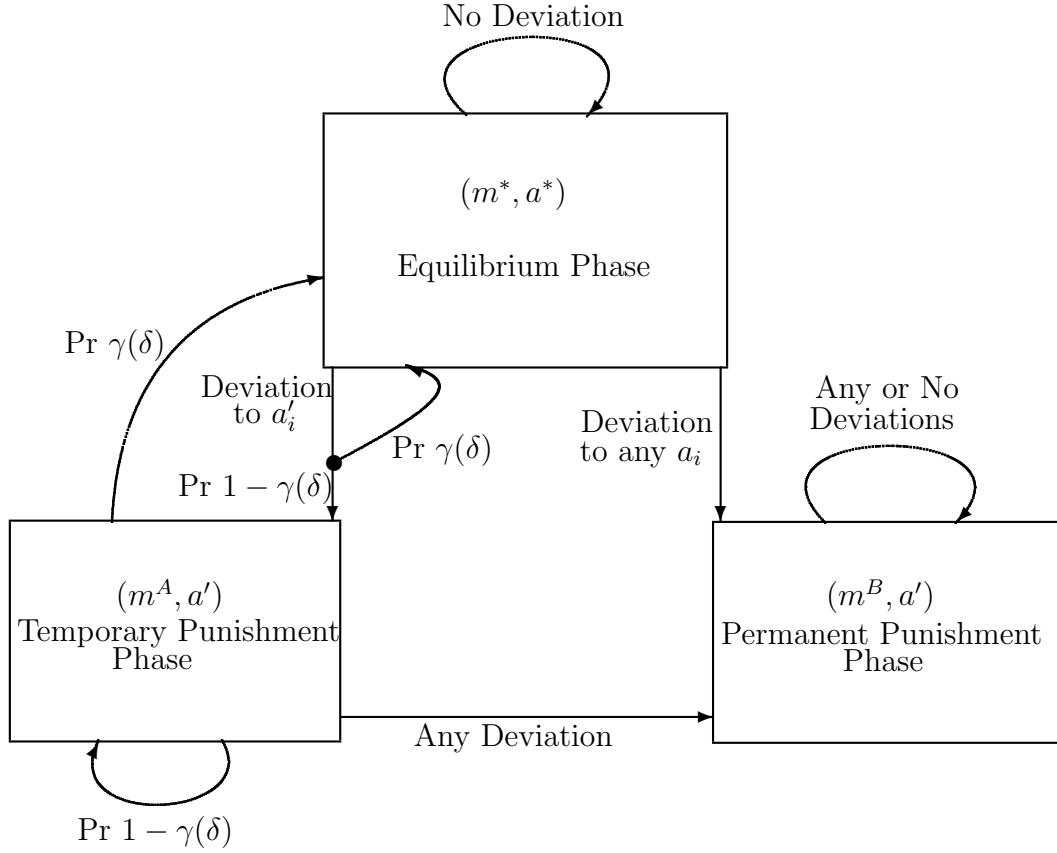


Figure 1: Three Phases of Play

two or more players are again ignored). Then all players $\langle i \in I, t \rangle$ send message m^B to their successors. From this point on the prescriptions of the equilibrium strategies are that all subsequent players should send message m^B , and that they should play action a'_i . Figure 1 is a schematic depiction of the three phases of play we have outlined.

To check that the strategies we have described constitute an SE of the dynastic repeated game, the players’ beliefs need to be specified. All players in the $t = 0$ cohort of course have correct beliefs. All other players, after receiving message m^* believe that all other players have also received message m^* with probability one. Similarly, after receiving message m^A all time- $t \geq 1$ players believe that all other players have received message m^A with probability one. So, when play is either in the equilibrium phase all players believe that this is indeed the case, and the same is true for the temporary punishment phase. Not surprisingly, it is possible to sustain these beliefs via a sequence of completely mixed strategies using Bayes’ rule, as required for an SE.

After receiving message m^B , all time- $t \geq 1$ players have non-degenerate beliefs as follows. With probability $\beta(\delta)$ they believe that play is in the equilibrium phase so that all other players have in fact received message m^* . With probability $1 - \beta(\delta)$ they believe that play is in the permanent punishment phase, so that all other players have received message m^B . So, when play is in the permanent punishment phase the players’ beliefs assign positive

probability to something that is in fact not taking place. Why is this possible in an SE? The intuition is not hard to outline. Upon receiving message m^B , player $\langle i, t \rangle$ needs to compare two possibilities: am I receiving message m^B because some action-deviations have occurred that triggered the permanent punishment path, or is it simply the case that play is in the equilibrium phase and my immediate predecessor has sent message m^B when he should have sent m^* ? Clearly the answer depends on the probability of a message-deviation relative to the probability of action-deviations that could have triggered the permanent punishment phase. By a careful use of infinitesimals of different orders in the sequence of completely mixed strategies, it is then possible to sustain the beliefs we have described via Bayes’ rule.

Given the beliefs we have specified, we can now argue that no player has an incentive to deviate, either at the action or at the message stage, from the strategies we have described. We distinguish again between the three phases of play identified above.

If play is in the equilibrium phase it is not hard to see that no player $\langle i, t \rangle$ has an incentive to deviate. If he adheres to the equilibrium strategy, player $\langle i, t \rangle$ gets a payoff of $u_i(a^*)$. Since δ is close to 1, and player $\langle i, t \rangle$ takes the strategies of all other players (including his successors) as given, any deviation may produce an instantaneous gain, but will decrease the overall payoff of player $\langle i, t \rangle$.

If play is in the temporary punishment phase and player $\langle i, t \rangle$ does not deviate then play will eventually go back to the equilibrium phase. Hence, since δ is close to one, the dynastic payoff to player $\langle i, t \rangle$ in this case is close to (but below) $u_i(a^*)$. Therefore, the same argument as for the equilibrium phase applies to show that he will not want to deviate during the temporary punishment phase.

Suppose now that play is in the permanent punishment phase. Begin with the action stage. Recall that at this point, that is after receiving message m^B , player $\langle i, t \rangle$ has the non-degenerate beliefs we described above giving positive probability to both the event that play is the equilibrium phase and to the event that play is in the permanent punishment phase.

If play were in the equilibrium phase, since δ is close to one, for an appropriate choice of $\gamma(\delta)$, player $\langle i, t \rangle$ would prefer taking action a_i^* to taking action a'_i , and would prefer the latter to taking any other action a_i different from both a_i^* and a'_i . This is because his continuation payoff (from the beginning of period $t + 1$ onwards) is higher in the equilibrium phase than in the temporary punishment phase, and lowest in the permanent punishment phase. Notice that if $\gamma(\delta)$ were too close to one, then the continuation payoff in the equilibrium phase and in the temporary punishment phase could be so close as to reverse (via an instantaneous gain in period t from playing a'_i) the preference of player $\langle i, t \rangle$ between a_i^* and a'_i . On the other hand if $\gamma(\delta)$ were too low, then the continuation payoff in the temporary punishment phase and in the permanent punishment phase could be so close as to reverse (via an instantaneous gain in period t from playing some action a_i different from both a_i^* and a'_i) the preference of player $\langle i, t \rangle$ between a'_i and some other action a_i .

If play were in the permanent punishment phase then clearly player $\langle i, t \rangle$ would prefer taking action a'_i to taking action a_i^* . This is simply because by assumption $u_i(a') > u_i(a_i^*, a'_{-i})$.

Of course, it is possible that some other action(s) a_i would be preferable to a'_i .

So, in one of two cases that player $\langle i, t \rangle$ entertains with positive probability after receiving message m^B action a_i^* is preferable to action a'_i which in turn is preferable to any other action, while in the other case action a'_i is preferable to a_i^* , but some other action(s) may be preferable to both. What is critical here is that there is *no* action (a_i^* or any other one) that is preferable to a'_i with probability one. As it turns out, this is sufficient to show that for some value of $\beta(\delta)$ action a'_i is in fact optimal for player $\langle i, t \rangle$. Therefore he does not want to deviate from the equilibrium strategy we have described at the action stage during the permanent punishment phase.

Finally, consider the choice of player $\langle i, t \rangle$ at the message stage when play is in the permanent punishment phase. Notice that after receiving message m^B , player $\langle i, t \rangle$ will discover that play is in the permanent punishment phase during the action stage of period t . This is so even if some other time- t player were to deviate from his prescribed action in period t . The reason is our assumption that $n \geq 3$. The fact that $n - 2$ or more players $\langle j, t \rangle$ play a'_j is sufficient to tell player $\langle i, t \rangle$ that play is in fact in the permanent punishment phase.²² Clearly, it is at this point that player $\langle i, t \rangle$ would like to “communicate effectively” to player $\langle i, t + 1 \rangle$ that play is in the permanent punishment phase but is unable to do so, as in our discussion above concerning message spaces. In fact a'_i could be very far from being a best response from a'_{-i} in the stage game.²³ Yet, given the strategies of all other players, inducing all his successors to play a myopic best-response to a'_{-i} is simply not an option that is open to player $\langle i, t \rangle$ even after he discovers that play is in the permanent punishment phase. Given that, by assumption, $u_i(a') > u_i(a_i^*, a'_{-i})$ the best that player $\langle i, t \rangle$ can do at this point is to send message m^B as required by his equilibrium strategy.

The argument we have just outlined suffices to show that the payoff vector $u(a^*)$ can be sustained as an SE of the dynastic repeated game. We now argue that this fact can be used as a “local anchor” for our global argument that shows that *any* interior payoff vector can be sustained as an SE. Fix any $v^* \in \text{int}V$ to be sustained in equilibrium. Since v^* is interior it is obvious that it can be expressed as $v^* = qu(a^*) + (1 - q)v$ for some $q \in (0, 1)$ and some $v \in V$.²⁴ The construction we put forward uses the payoffs $v' = qu(a') + (1 - q)v$ as “punishments.” Clearly, $v_i^* > v'_i$ for every $i \in I$.

We use the action stage randomization device to combine the local part of the argument with the global one. The possible realizations of \tilde{x}^t are $(x(1), \dots, x(\|A\|))$, with the probability

²²If we had $n = 2$ players the following problem would arise with our construction. Suppose that each player were to monitor the action of the other to decide whether play is in the equilibrium phase or in the permanent punishment phase. Then, after receiving m^B , say for instance player $\langle 1, 1 \rangle$, could find it advantageous to play a_1^* instead of a'_1 . This is because this deviation, together with player $\langle 1, 1 \rangle$ sending message m^* , would put the path of play back in the equilibrium phase with probability one.

²³Note that this would necessarily be the case if, for instance, $u_i(a^*)$ were below i 's minmax payoff in the stage game.

²⁴In fact, depending on the position of $u(a^*)$ within V , we may be able to express some v^* vectors that are on the *boundary* of V in this way as well. See Remark 2 above.

that $\tilde{x}^t = x(1)$ equal to q .²⁵ Whenever the realization of \tilde{x}^t is $x(1)$, the action and message strategies of all players are just as we described above. The action strategies of all players do not depend on the message received whenever $x^t \neq x(1)$. Moreover, the probability law governing \tilde{x}^t and the action strategies of all players for the case $x^t \neq x(1)$ are such that the per-period expected payoff (conditional on $x^t \neq x(1)$) is v_i for every $i \in I$. Whenever $x^t \neq x(1)$, *any* deviation from the prescription of the equilibrium strategies triggers the permanent punishment phase, as described above. It is not hard to verify that the latter is enough to keep all players from deviating at any point.

One point that is worth emphasizing here is that, in contrast to what happens in the local argument above, if play is in the permanent punishment phase, so that player $\langle i, t \rangle$ receives message m^B , and $x^t \neq x(1)$, then player $\langle i, t \rangle$ does *not* discover from a^t whether play is in the equilibrium phase or in the permanent punishment phase. Both at the beginning and at the end of period t he believes with probability $\beta(\delta)$ that play is in the equilibrium phase, and with probability $1 - \beta(\delta)$ that play is in the permanent punishment phase. However, he does know that any deviation to an action different from the prescribed one will trigger the permanent punishment phase for sure. Given that δ is close to one, he is then better-off not deviating. This ensures that the first of his successors who observes a realization of the action randomization device equal to $x(1)$ will play action a'_i , which is in fact optimal given the beliefs that player $\langle i, t \rangle$ has.

7. Four Dynasties or More

We now turn to the case in which the stage game G has four or more players. We need to introduce some further notation to work towards our next main result. We begin by constructing what we will refer to as the *restricted correlated minmax*.

Given a stage game $G = (A, u, I)$, we indicate by $\tilde{A} \subseteq A$ a typical set of pure action profiles with a *product structure*. In other words, we require that there exist an array $(\tilde{A}_1, \dots, \tilde{A}_n)$ with $\tilde{A} = \times_{i \in I} \tilde{A}_i$. Given a product set \tilde{A} , we let $V(\tilde{A})$ be the convex hull of the set of payoff vectors that can be achieved in G using pure action profiles in \tilde{A} . As before, $\text{int}V(\tilde{A})$ will denote the (relative) interior of $V(\tilde{A})$.

Definition 3. *Restricted Correlated Minmax:* Let a product set $\tilde{A} \subseteq A$ be given. Now let

$$\underline{\omega}_i(\tilde{A}) = \min_{z_{-i} \in \Delta(\tilde{A}_{-i})} \max_{a_i \in \tilde{A}_i} \sum_{a_{-i} \in \tilde{A}_{-i}} z_{-i}(a_{-i}) u_i(a_i, a_{-i}) \quad (6)$$

where z_{-i} is any probability distribution over the finite set \tilde{A}_{-i} (not necessarily the product of independent marginals), and $z_{-i}(a_{-i})$ denotes the probability that z_{-i} assign to the profile a_{-i} .

We then say that $\underline{\omega}_i(\tilde{A})$ is the restricted (to \tilde{A}) correlated minmax for i in G .

²⁵ Throughout the paper, we adopt the standard notational convention by which $\|\cdot\|$ denotes the cardinality of a set.

Roughly speaking, the restricted (to \tilde{A}) correlated minmax payoff for i is the best payoff that i can achieve when he is restricted to choosing an element of \tilde{A}_i , while all other players are choosing a profile of correlated mixed strategies with support at most \tilde{A}_{-i} .²⁶

We are now ready to state our next result. The proof of Theorem 3 begins in Appendix B, and continues in the technical addendum to the paper (Sections T.9 through T.12).

Theorem 3. *Dynastic Folk Theorem: Four Dynasties or More: Let a stage game G with four or more players be given. Assume that G is such that we can find a product set $\tilde{A} \subseteq A$ and an array of $n + 1$ payoffs vectors $\hat{v}, \bar{v}^1, \dots, \bar{v}^n$ for which the following conditions hold.*

- (i) *For every $i \in I$, the set \tilde{A}_i contains at least two elements.*
- (ii) *$\hat{v} \in \text{int}V(\tilde{A})$, and $\bar{v}^i \in V(\tilde{A})$ for every $i \in I$.*
- (iii) *$\underline{\omega}_i(\tilde{A}) < \bar{v}_i^i < \bar{v}_i^j$ and $\bar{v}_i^i < \hat{v}_i$ for every $i \in I$ and every $j \neq i$.*

Then for every $v \in \text{int}V$ there exists a $\underline{\delta} \in (0, 1)$ such that $\delta > \underline{\delta}$ implies $v \in \mathcal{E}^D(\delta)$.

Before any discussion of the argument behind Theorem 3 we proceed with a remark concerning its scope.

Remark 4. *Theorem 3 and the “Standard” Folk Theorem: Suppose that the stage game G has four or more players. Suppose further that G is such that the standard Folk Theorem (Fudenberg and Maskin 1986, Theorems 2 and 5) guarantees that, as δ approaches one, the standard repeated game has at least one SPE payoff vector that strictly dominates the standard minmax.²⁷ Then Theorem 3 guarantees that, as δ approaches one, all (interior) feasible payoff vectors can be sustained as an SE of the dynastic repeated game.*

To see this, assume that V contains a payoff vector that yields all players more than their (standard) minmax \underline{v}_i and that V has full dimension. Next, observe that (see footnote 26) if we set $\tilde{A} = A$ we know that $\underline{\omega}_i(A) \leq \underline{v}_i$ for every i . Therefore, using the full-dimensionality

²⁶Recall that the “standard minmax” for player i is defined as $\underline{v}_i = \min_{\sigma_{-i} \in \Pi_{j \neq i} \Delta(A_j)} \max_{a_i \in A_i} u_i(a_i, \sigma_{-i})$, where $\Pi_{j \neq i} \Delta(A_j)$ is the set of (independent) mixed strategy profiles for $-i$.

It is then straightforward to see that the restricted correlated minmax of Definition 3 is a “stronger” concept than the standard minmax in the following sense. Given any $G = (A, u, I)$, if we let the standard minmax payoff for i be denoted by \underline{v}_i we obviously have that $\underline{\omega}_i(A) \leq \underline{v}_i$ for every $i \in I$.

Of course, the relationship between the restricted correlated minmax and the standard minmax also depends on the “size” of product set \tilde{A} relative to A . In particular, let two product sets \tilde{A} and \tilde{A}' be given. It is then straightforward to check that if $\tilde{A}_{-i} \subseteq \tilde{A}'_{-i}$ and $\tilde{A}'_i \subseteq \tilde{A}_i$ then $\underline{\omega}_i(\tilde{A}') \leq \underline{\omega}_i(\tilde{A})$. Therefore, in general, depending on \tilde{A} , the value of $\underline{\omega}_i(\tilde{A})$ could be below, equal or above \underline{v}_i .

²⁷Notice that there is an open set of stage games for which the standard Folk Theorem does not imply any multiplicity of equilibrium payoffs. To see this, consider, for instance, a game G in which every player i has a strictly dominant strategy a_i^* and such that the standard minmax value is $u_i(a^*)$ for every $i \in I$. Clearly, these conditions can be satisfied even when $u(a^*)$ is not (weakly) Pareto-dominated by any other payoff vector in V . Moreover, if these conditions are satisfied, they are also satisfied for an entire open set of games around G .

of V , it is straightforward to see that setting $\tilde{A} = A$ all the hypotheses required by Theorem 3 are in fact satisfied.

It is important to emphasize at this point that Theorem 3 implies that all (interior) payoff vectors may be sustained as SE of the dynastic repeated game as δ goes to one, even when the standard Folk Theorem of Fudenberg and Maskin (1986) does *not* imply any multiplicity of payoffs. This is because of the two critical features that distinguish the Restricted Correlated Minmax of Definition 3 from the standard minmax: the correlation that it allows, and the restriction to the product set \tilde{A} .

To appreciate the role of correlation, consider a stage game G with four or more players, each with at least two pure actions, and with V of full dimension. Recall that (see footnote 26) we know that it is possible that $\underline{\omega}_i(A) < \underline{v}_i$. If the latter condition is satisfied for all players, it is clearly possible that no payoff vector in V gives all players more than \underline{v}_i , but at the same time all the hypotheses of Theorem 3 are in fact satisfied.²⁸

To see the role of the restriction to \tilde{A} , notice that, as in Theorem 2, the necessary conditions listed in the statement of Theorem 3 are “local.” In other words (see footnote 26), it is clearly possible that no payoff vector in V gives all players more than \underline{v}_i , but that by excluding one or more pure strategies for one or more players we obtain a product set \tilde{A} for which conditions (i), (ii) and (iii) of Theorem 3 are in fact satisfied.

Similarly to the case of Theorem 2, the construction we use to prove Theorem 3 relies on message spaces that are smaller than the appropriate set of finite histories H^t . Also similarly to Theorem 2, it is convenient to divide the argument behind Theorem 3 into a local one and a global one. As before, we begin with the local part of the argument.

Consider a product set \tilde{A} satisfying the hypotheses of Theorem 3. Let $V(\tilde{A})$ be the convex hull of payoff vectors generated with action profiles in \tilde{A} . For simplicity, assume that the payoff vectors $\bar{v}^1, \dots, \bar{v}^n$ can all be obtained from pure profiles of actions in \tilde{A} . Also for simplicity, assume that each of the payoffs $\underline{\omega}_i(\tilde{A})$ can be obtained from some (pure) profile of actions in \tilde{A} .²⁹ The local part of the argument behind Theorem 3 shows that, for δ close enough to one the payoff vector $\hat{v} \in V(\tilde{A})$ of the statement of the Theorem can be sustained as an SE of the dynastic repeated game.

²⁸The example we provided in footnote 27 may prompt the following question. Is it the case that whenever the standard Folk Theorem does not yield any payoff multiplicity it must necessarily be the case that every player has a dominant strategy in G ? The answer is no. This observation is relevant to the role of correlation we are discussing here. If all players have a dominant strategy in G , then the correlated minmax $\underline{\omega}_i(A)$ is in fact the *same* as the standard minmax \underline{v}_i for every i . Clearly, what we are saying here has real content only when G is such that the standard Folk Theorem does not yield any payoff multiplicity and it is *not* the case that all players have a dominant strategy.

²⁹Clearly, when $\underline{\omega}_i(\tilde{A})$ is achieved via a pure action profile in \tilde{A} the restricted correlated minmax is the same as the standard minmax restricted to \tilde{A} . It is still obviously the case that if \tilde{A} is a strict subset of A , then it is possible that $\underline{\omega}_i(\tilde{A})$ would be below the standard (unrestricted) minmax \underline{v}_i .

The equilibrium path generated by the strategies we construct consists of $n + 1$ phases.³⁰ We call the first one the *standard equilibrium phase*, the second one the *diversionary-1 equilibrium phase*, through to the *diversionary- n equilibrium phase*.

If all players $\langle i \in I, t \rangle$ receive message m^* then play is in the standard equilibrium phase. For simplicity again we proceed with our outline of the construction assuming that the equilibrium prescribes that the players $\langle i \in I, t \rangle$ play a pure action profile during the standard equilibrium phase, denoted by a' . The associated payoff vector is v' .

If all players $j \neq i$ in the t -th cohort receive message \check{m}^i , and player $\langle i, t \rangle$ receives any message $m^{i,\tau}$ in a finite set $\underline{M}(i, t) = \{m^{i,1}, \dots, m^{i,T}\} \subset M_i^t$, then play is in the diversionary- i phase.³¹ Let \underline{a}^i be the vector of actions (pure, for simplicity) for which i receives his restricted correlated minmax payoff $\omega_i(\tilde{A})$. During the diversionary- i equilibrium phase player $\langle i, t \rangle$ plays action \underline{a}_i^i . For all players $j \neq i$, let \check{a}_j^i be any action in \tilde{A}_j that is not equal to \underline{a}_j^i . Such action can always be found since, by assumption, each \tilde{A}_j contains at least two actions. During the diversionary- i equilibrium phase, any player $j \neq i$ plays action \check{a}_j^i . The (per-period) payoff vector associated with the diversionary- i equilibrium phase is denoted by \check{u}^i .

If in period t play is in any of the equilibrium phases we have just described, and no deviation occurs at the action stage, at the end of period t all players use the realization y^t of the message-stage randomization device to select the message to send to their successors. The possible realizations of \tilde{y}^t are $(y(0), y(1), \dots, y(n))$. The probability that $y^t = y(0)$ is $1 - \eta$ and the probability that $y^t = y(i)$ is η/n for every $i \in I$. Consider now the end of any period t in any equilibrium phase, and assume that no deviation has occurred. If $y^t = y(0)$ then all players $\langle i \in I, t \rangle$ send message m^* to their successors, so play in period $t + 1$ is in the standard equilibrium phase. If $y^t = y(i)$, then all players $j \neq i$ send message \check{m}^i to their successors and player $\langle i, t \rangle$ sends a (randomly selected) message $m^{i,\tau} \in \underline{M}(i, t)$ to player $\langle i, t + 1 \rangle$. So, in this case in period $t + 1$ play is in the diversionary- i equilibrium phase.

The profiles to be played in each diversionary- i equilibrium phase may of course be entirely determined by the (need to differ from the) correlated minmax action profiles, so we have no degrees of freedom there. However, we are free to choose the profile a' in constructing the standard equilibrium phase. Recall that we take a' to be pure solely for expositional simplicity. Using the action-stage randomization device, clearly we could select correlated actions for the standard equilibrium phase that yield any payoff vector v' in $V(\tilde{A})$. Since \hat{v}

³⁰The formal proof of Theorem 3 (see Appendix B and the technical addendum to the paper) actually treats period 0 differently from all other periods. In period 0, in equilibrium, the players use the action-stage randomization device to play actions that yield exactly \hat{v} . This payoff vector is achieved in *expected* terms (across different equilibrium phases) on the equilibrium path by the construction we outline here.

³¹In the formal proof of Theorem 3, presented in Appendix B and the technical addendum, the set of messages $\underline{M}(i, t)$ actually does depend on the time index t because not all messages are available for the first T periods of play. This is so in order to avoid any message space M_i^t having a cardinality that *exceeds* that of H^t .

$\in \text{int}V(\tilde{A})$ we can be sure that for some $v' \in V(\tilde{A})$ and some $\eta \in (0, 1)$

$$\hat{v} = (1 - \eta)v' + \frac{\eta}{n} \sum_{i=1}^n \check{u}_i \quad (7)$$

So that (modulo our expositional assumption that $v' = u(a')$) when play is in any equilibrium phase the expected (across all possible realizations of \tilde{y}^t) continuation payoff to any player $\langle i, t \rangle$ from the beginning of period $t + 1$ onward is \hat{v}_i .

The strategies we construct also define (off the equilibrium path) n *punishment phases*, one for each $i \in I$ and n *terminal phases*, again one for each dynasty i . In the punishment- i phase, in every period player i receives his restricted correlated minmax payoff $\underline{\omega}_i(\bar{a})$ payoff, and the phase lasts T periods. In the terminal- i phase in every period the players receive the vector of payoffs \bar{v}^i . The transition between any of the equilibrium phases and any of the punishment or terminal phases is akin to the benchmark construction in Fudenberg and Maskin (1986). In other words, a deviation by dynasty i during any of the equilibrium, any of the punishment phases or any of the terminal phases triggers the start (or re-start) of the punishment- i phase (deviations by two players or more are ignored). The terminal- i phase begins after play has been, without subsequent deviations, in the punishment- i phase for T periods. For an appropriately chosen (large) T , as δ approaches 1, with one critical exception, the inequalities in (iii) of the statement of Theorem 3 are used in much the same way as in Fudenberg and Maskin (1986) to guarantee that no player deviates from the prescriptions of the equilibrium strategies.

In Fudenberg and Maskin (1986), during the punishment- i phase player i plays a *myopic best response* to the actions of other players. Critically, this is not the case here. During the punishment- i phase dynasty i plays a best response to the (correlated) strategy of others *restricted to the set of pure actions in \tilde{A}_i* . Clearly this could be very far away (in per-period payoff terms) from an unconstrained best reply chosen at will within A_i . To understand how this can happen in an SE we need to specify what message profiles mark the beginning of the punishment- i phase and what the players’ beliefs are.

Suppose that player $\langle i, t \rangle$ deviates from the prescriptions of the equilibrium strategy and triggers the start of the punishment- i phase as of the beginning of period $t + 1$. Then all players $\langle j \in I, t \rangle$ send message $m^{i,T}$ to their successors. These messages are interpreted as telling all players $\langle j \neq i, t + 1 \rangle$ that the punishment- i phase has begun, and that there are T periods remaining. We return to the beliefs of player $\langle i, t + 1 \rangle$ shortly. In the following period of the punishment- i phase all players $\langle j \in I, t + 1 \rangle$ send message $m^{i,T-1}$, then $m^{i,T-2}$ and so on, so that the the index τ in a message $m^{i,\tau}$ is effectively interpreted (by dynasties $j \neq i$) as a “punishment clock,” counting down how many periods remain in the punishment- i phase.

The players’ beliefs in the SE we construct are “correct” in all phases of play except for the beliefs of player $\langle i, t \rangle$ whenever play is in the punishment- i phase at time t . Upon receiving any message $m^{i,\tau} \in \underline{M}(i, t)$, player $\langle i, t \rangle$ believes that play is in the punishment- i phase with probability zero. Instead he believes that play is in the i -diversionary equilibrium phase with

probability one. This is possible in an SE because player $\langle i, t \rangle$ receives the same message in both cases and play is in the diversionary- i phase with positive probability *on the equilibrium path*, while the punishment- i phase is entirely off the equilibrium path. It follows easily that player $\langle i, t \rangle$ does not want to deviate from the equilibrium strategies we have described when play is in the punishment- i phase.

Notice moreover that if at time t play is in the punishment- i phase, after the profile a^t is observed, player $\langle i, t \rangle$ will discover that play is in fact in the punishment- i phase, contrary to his beginning-of-period beliefs, even if a new deviation occurs at the action stage of period t . This is because, by construction, all players $\langle j \neq i, t \rangle$ play an action (namely \check{a}_j^i) in the diversionary- i equilibrium phase that is different from what they play in the punishment- i phase (namely \underline{a}_j^i). This, coupled with the assumption assumption that $n \geq 4$ will ensure that $\langle i, t \rangle$ will discover the truth, and the identity of any deviator.³²

Therefore any player $\langle i, t \rangle$ who knows that in period $t + 1$ play will be in the punishment- i phase would like to “communicate effectively” to player $\langle i, t + 1 \rangle$ that play is in the punishment- i phase but is unable to do so, as in our discussion above concerning message spaces. After receiving $m^{i,\tau}$ and discovering that play is in the punishment- i phase, sending any message to player $\langle i, t + 1 \rangle$ that is not $m^{i,\tau-1}$ may cause him, or some of his successors, to deviate and hence to re-start the punishment- i phase.³³ Sending $m^{i,\tau-1}$ will cause player $\langle i, t + 1 \rangle$ and his successors to play a best response among those that can be induced by any available message. This is because \underline{a}_i^i is in fact a best response to the the correlated minmax of the other players within the set \tilde{A}_i . Therefore, given the inequalities in (iii) of the statement of Theorem 3, given that T is sufficiently large, and that δ is close enough to one, no profitable deviation is available to player $\langle i, t \rangle$.

The argument we have just outlined suffices to show that the payoff vector \hat{v} of the statement of the theorem can be sustained as an SE of the dynastic repeated game. Similarly to Theorem 2, we now argue that this fact can be used as a “local anchor” for an argument that shows that *any* interior payoff vector can be sustained as an SE.

Fix any $v^* \in \text{int}V$ to be sustained in equilibrium. Since v^* is interior it is obvious that it can be expressed as $v^* = q\hat{v} + (1 - q)z$ for some $q \in (0, 1)$ and some $z \in V$. The “global” argument then consists of using the action-stage randomization device so that in each period with probability q play proceeds as in the construction above, while with probability $1 - q$ the (expected) payoff vector is z . The latter is achieved with action-stage strategies that do not depend on the messages received. A deviation by i from the (correlated) actions needed

³²Clearly, if $\langle i, t \rangle$ could not be guaranteed to discover that play is in the punishment- i phase, or the identity of any deviator at time t , then we could not construct strategies that guarantee that if $\langle j \neq i, t \rangle$ deviates during the punishment- i phase then play switches to the punishment- j phase, as required.

³³It is worth emphasizing here that checking sequential rationality at the message stage is somewhat more involved than may appear from our intuitive outline of the argument given here. This is because a deviation at the message stage may trigger *multiple deviations*; that is deviations at the action and/or message stage by more than one successor of any given player. The core of the argument dealing with this case is Lemma T.10.4.

to implement z triggers the punishment- i phase. With one proviso to be discussed shortly, it is not hard to then verify that this is sufficient to keep all players from deviating at any point, and hence that v^* can be sustained as an SE payoff vector of the dynastic repeated game for δ sufficiently close to one.

The difficulty with the global argument we have outlined that needs some attention is easy to point out. The periods in which the action-stage randomization device tells the players to implement the payoff vector z cannot be counted as real punishment periods. They in fact stochastically interlace all phases of play, including any punishment- i phase. However, the length of effective punishment T has to be sufficiently large to deter deviations. The solution we adopt is to ensure that the punishment clock does not decrease in any period in which (with probability $1 - q$) the payoff vector z is implemented at the action stage. In effect, this makes the length of any punishment- i phase stochastic, governed by a punishment clock that counts down only with probability q in every period.

8. Inter-Generational (Dis)Agreement

Some of the SE of the dynastic repeated game we have identified in Theorems 2 and 3 above clearly do not correspond in any meaningful sense to any SPE of the standard repeated game. This is obvious if we consider an SE of the dynastic repeated game in which one or more players receive a payoff below their standard minmax value.

An obvious question to raise at this point is then the following. What is it that makes these SE viable? To put it another way, can we identify any properties of an SE of the dynastic repeated game which ensure that it must correspond in a meaningful sense to an SPE of the standard repeated game? The answer is yes, and this is what this section of the paper is devoted to.

The critical properties of an SE that we identify concern the players’ beliefs. These properties characterize entirely the set of SE yielding payoffs outside the set of SPE of the standard repeated game. Therefore, if one wanted to attempt to “refine away” the equilibria yielding payoffs outside the set of SPE, our results in this section pin down precisely which belief systems the proposed refinement would have to rule out.

The first of the properties we identify is that a player’s (revised) beliefs at the end of the period over the messages received by other players at the beginning of the period should be the *same* as at the beginning of the period. This is equation (8) below. The second is that the end-of-period beliefs of a player (over messages sent by his opponents) should be the *same* as the beginning-of-period beliefs of his successor (over messages received by his opponents). This is equation (9) below. In fact we will be able to show that if this property holds for all players (and all information sets) in an SE of the dynastic repeated game, then this SE must be payoff-equivalent to some SPE of the standard repeated game.³⁴ For want of

³⁴As will be clear from the proof of Theorem 4 we are able to show that there is an, appropriately defined, strategic equivalence between such SE of the dynastic repeated game and the SPE of the standard repeated game.

a better term, when an SE of the dynastic repeated game has the two properties (of beliefs) that we just described informally, we will say that it displays *Inter-Generational Agreement*.

Definition 4. *Inter-Generational Agreement*: Let an SE triple (g, μ, Φ) of the dynastic repeated game be given.

We say that this SE displays *Inter-Generational Agreement* if and only if for every $i \in I$, $t \geq 0$, $m_i^t \in H^t$, $x^t \in X$, $a^t \in A$ and $y^t \in Y$ we have that

$$\Phi_i^{tR}(m_i^t, x^t, a^t, y^t) = \Phi_i^{tB}(m_i^t) \quad (8)$$

and for every m_i^{t+1} in the support of $\mu_i^t(m_i^t, x^t, a^t, y^t)$

$$\Phi_i^{tE}(m_i^t, x^t, a^t, y^t) = \Phi_i^{t+1B}(m_i^{t+1}) \quad (9)$$

We are now ready to state our last result.

Theorem 4. *SE of the Dynastic and SPE of the Standard Repeated Game*: Fix a stage game G , any $\delta \in (0, 1)$, and any \tilde{x} and \tilde{y} . Let (g, μ, Φ) be an SE of the dynastic repeated game. Assume that this SE displays *Inter-Generational Agreement* as in Definition 4. Let v^* be the vector of (dynastic) payoffs for $t = 0$ players in this SE.

Then $v^* \in \mathcal{E}^S(\delta)$.

The proof of Theorem 4 is in Appendix C. Before proceeding with an intuitive outline of the proof, we state a remark on the implications of the theorem.

Clearly, Theorem 4 implies that if a payoff vector $v \in \mathcal{E}^D(\delta)$ is not sustainable as an SPE so that $v \notin \mathcal{E}^S(\delta)$, then it must be the case that no SE which sustains v in the dynastic repeated game displays *Inter-Generational Agreement*.

Remark 5. *SE and SPE Payoffs*: Observe that the SE that we construct in the proof of Theorem 1 all in fact satisfy *Inter-Generational Agreement*. In other words, we know that if a payoff vector v is in $\mathcal{E}^S(\delta)$ then there is an SE of the dynastic repeated game that displays *Inter-Generational Agreement* which sustains v .

Together with Theorem 4, this gives us a complete characterization in payoff terms of the relationship between the SE of the dynastic and the SPE of the standard repeated game as follows.

Fix any stage game G and any $\delta \in (0, 1)$. Then a payoff vector $v \in \mathcal{E}^D(\delta)$ is in $\mathcal{E}^D(\delta)/\mathcal{E}^S(\delta)$ if and only if no SE which sustains v in the dynastic repeated game displays *Inter-Generational Agreement*.

To streamline the exposition of the outline of the argument behind Theorem 4, make the following two simplifying assumptions. First, assume that at the message stage the players do

not condition their behavior at the message-stage on the randomization device. The simplest way to fix ideas here is to consider a message-stage randomization device with a singleton Y (the set of possible realizations). Second, assume that all message-strategies μ_i^t are *pure*. In other words, even though they may randomize at the action stage, all players at the message stage send a single message, denoted $\mu_i^t(m_i^t, x^t, a^t)$, with probability one.³⁵

Now consider an SE (g, μ, Φ) of the dynastic repeated game that satisfies Inter-Generational Agreement as in Definition 4. Fix any history of play $h^t = (x^0, a^0, \dots, x^{t-1}, a^{t-1})$. For each dynasty i , using the message strategies of players $\langle i, 0 \rangle$ through to $\langle i, t-1 \rangle$, we can now determine the message m_i^t that player $\langle i, t-1 \rangle$ will send to his successor, player $\langle i, t \rangle$. Denote this message by $m_i^t(h^t)$. Notice that $m_i^t(h^t)$ can be determined simply by recursing forward from period 0. Recall that at the beginning of period 0 all players $\langle i \in I, 0 \rangle$ receive message $m_i^0 = \emptyset$. Therefore, given $h^1 = (x^0, a^0)$, using μ_i^0 , we know $m_i^1(h^1)$. Now using μ_i^1 we can compute $m_i^2(h^2) = \mu_i^1(m_i^1(h^1), x^1, a^1)$, and so recursing forward the value of $m_i^t(h^t)$ can be worked out.

Because the SE (g, μ, Φ) satisfies Inter-Generational Agreement it must be the case that, after any actual history of play (on or off the equilibrium path) h^t , and therefore after receiving message $m_i^t(h^t)$, player $\langle i, t \rangle$ believes that his opponents have received messages $(m_1^t(h^t), \dots, m_{i-1}^t(h^t), m_{i+1}^t(h^t), \dots, m_n^t(h^t))$ with probability one.

To see why this is the case, we can recurse forward from period 0 again. Consider the end of period 0. Since all players in the $t = 0$ cohort receive message \emptyset , after observing (x^0, a^0) , player $\langle i, 0 \rangle$ *knows* that any player $\langle j \neq i, 0 \rangle$ is sending message $m_j^1(x^0, a^0) = \mu_j^0(\emptyset, x^0, a^0)$ to his successor player $\langle j, 1 \rangle$.

Equation (9) of Definition 4 guarantees that the beginning-of-period beliefs of player $\langle i, 1 \rangle$ must be the same as the end-of-period beliefs of player $\langle i, 0 \rangle$. So, at the beginning of period 1, player $\langle i, 1 \rangle$ believes with probability one that every player $\langle j \neq i, 1 \rangle$ has received message $m_j^1(x^0, a^0)$ as above.

Equation (8) of Definition 4 guarantees that player $\langle i, 1 \rangle$ will not revise his beginning-of-period beliefs during period 1. Therefore, after observing any (x^1, a^1) , player $\langle i, 1 \rangle$ still believes that every player $\langle j \neq i, 1 \rangle$ has received message $m_j^1(x^0, a^0)$ as above. But this, via the message strategies μ_j^1 implies that player $\langle i, 1 \rangle$ must believe with probability one that every player $\langle j, 1 \rangle$ sends message $m_j^2(h^2) = m_j^2(x^0, a^0, x^1, a^1) = \mu_j^1(m_j^1(h^1), x^1, a^1)$ to his successor $\langle j, 2 \rangle$. Continuing forward in this way until period t we can then see that the beginning-of-period beliefs of player $\langle i, t \rangle$ are as we claimed above.

Before we proceed to close the argument for Theorem 4, notice that *both* conditions

³⁵Note that we are somewhat abusing notation here. Recall that for every player $\langle i, t \rangle$, both $g_i^t(m_i^t, x^t)$ and $\mu_i^t(m_i^t, x^t, a^t, y^t)$ denote probability distributions (over actions and messages respectively). Here and throughout the rest of the paper and the technical addendum, we will sometimes construct strategies in which a single action is played with probability one and/or a single message is sent with probability one. Abusing notation slightly, we will write $g_i^t(m_i^t, x^t) = a_i$ to mean that the distribution $g_i^t(m_i^t, x^t)$ assigns probability one to a_i . Similarly, we will write $\mu_i^t(m_i^t, x^t, a^t, y^t) = m_i^{t+1}$ to mean that the distribution $\mu_i^t(m_i^t, x^t, a^t, y^t)$ assigns probability one to m_i^{t+1} .

of Definition 4 are necessary for our argument so far to be valid. Intuitively, the forward recursion argument we have outlined essentially ensures that the “correct” (because all its members receive a *given* message $m_i^0 = \emptyset$) beliefs of the first cohort “propagate forward” as follows. At the beginning of period $t = 1$ each player $\langle i \in I, 1 \rangle$ must have correct beliefs since the end-of-period beliefs of all players in period 0 are trivially correct, and equation (9) of Definition 4 tells us that the end-of-period beliefs must be the same as the beginning-of-periods beliefs of the next cohort. Now some pair (x^1, a^1) is observed by all players $\langle i \in I, 1 \rangle$. If this pair is consistent with their beginning-of-period beliefs, then clearly no player $\langle i \in I, 1 \rangle$ can possibly revise his beliefs on the messages received by others at the beginning of period 1. However, if (x^1, a^1) is *not* consistent with the beliefs of players $\langle i \in I, 1 \rangle$ and their action strategies, some players in the $t = 1$ cohort may be “tempted” to revise their beginning-of-period beliefs. This is because an observed “deviation” from what they expect to observe in period 1 can always be attributed to two distinct sources. It could be generated by an actual deviation at action stage of period 1, or it could be the result of one (or more) players in the $t = 0$ cohort having deviated at the message stage of period 0. Equation (8) of Definition 4 essentially requires that the $t = 1$ players should always interpret an “unexpected” pair (x^1, a^1) as an actual deviation at the action stage. The same applies to all subsequent periods. So, while equation (9) of Definition 4 ensures that the initially correct beliefs are passed on from one generation to the next, equation (8) of Definition 4 guarantees that actual deviations will be treated as such in the beliefs of players who observe them. The beliefs of players $\langle i \in I, 0 \rangle$ are correct and the end-of-period beliefs of any cohort are guaranteed to be the same as the beginning-of-period beliefs of the next cohort by equation (9). However, without equation (8) following an action deviation from the equilibrium path the end-of-periods beliefs of some players $\langle i, t \rangle$ could be incorrect, and be passed on to the next cohort intact.

Now recall that the punch-line of the forward recursion argument we have outlined is that if the SE (g, μ, Φ) satisfies Inter-Generational Agreement then we know that after any actual history of play h^t , player $\langle i, t \rangle$ believes that his opponents have received messages $(m_1^t(h^t), \dots, m_{i-1}^t(h^t), m_{i+1}^t(h^t), \dots, m_n^t(h^t))$ with probability one. To see how we can construct an SPE of the standard repeated game that is equivalent to the given SE, consider the strategies g_i^{t*} for the standard repeated game defined as $g_i^{t*}(h^t, x^t) = g_i^t(m_i(h^t), x^t)$. Clearly, these strategies implement the same payoff vector that is obtained in the given SE of the dynastic repeated game. Now suppose that the strategy profile g^* we have just constructed is not an SPE of the standard repeated game. Then, by the one-shot deviation principle (Remark 1 above) we know that some player i in the standard repeated game would have an incentive to deviate in a single period t after some history of play h^t . However, given the property of beliefs in the SE (g, μ, Φ) with Inter-Generational Agreement that we have shown above, this implies that player $\langle i, t \rangle$ would have an incentive to deviate at the action stage in the dynastic repeated game. This of course contradicts the fact that (g, μ, Φ) is an SE of the dynastic repeated game. Hence the argument is complete. The proof of Theorem 4 that appears in Appendix C of course does not rely on the two simplifying assumptions we made here. However, modulo some additional notation and technical issues, the argument presented there runs along the

same lines as the sketch we have given here.

9. Relation to the Literature

The infinitely repeated game model is a widely used construct in economics. Fudenberg and Maskin (1986), Aumann (1981), and Pearce (1992) are all standard references for the benchmark model.³⁶ When it is assumed that all individuals have full and direct knowledge of the past history of play, the standard model and the one we have studied here are equivalent.

Some recent papers have asked what happens to the equilibrium set when the full memory assumption is relaxed. The question posed then is: to what extent can intergenerational communication substitute for memory? Recent papers by Anderlini and Lagunoff (2005), Kobayashi (2003), and Lagunoff and Matsui (2004) all posit dynastic game models to address this question.

Among these, Anderlini and Lagunoff (2005) is the closest and, in many ways, the most direct predecessor of the current paper. Anderlini and Lagunoff (2005) examines this dynastic model when each player $\langle i, t \rangle$ receives a *public messages* from the player $\langle j \in I, t - 1 \rangle$ about the previous history of play. If the public messages from all player in the previous cohort are *simultaneous*, then a *Folk Theorem* in the sense of Fudenberg and Maskin (1986) can be obtained. Namely, if there are three or more players, all individually rational feasible payoffs can be sustained as an SE. Intuitively, this is because a version of a well known “cross-checking” argument that goes back to Maskin (1999) can be applied in this case.³⁷ By contrast, the present paper studies the model in which private communication (within each dynasty) may occur. We show that the difference between purely public and possible private communication is potentially large. Equilibria that sustain payoffs below some dynasty’s minmax exist, but they require inter-generational “disagreement.”

Our model is substantially different from the OLG games analyzed by Kandori (1992a), Salant (1991) and Smith (1992) among others (all three prove versions of the standard Folk Theorem).³⁸ In these papers there is no dynastic component to the players’ payoffs and full memory is assumed.³⁹

Kobayashi (2003) and Lagunoff and Matsui (2004) examine OLG games with a dynastic payoff component. As in Anderlini and Lagunoff (2005), these models assume entrants have no prior memory, and they also allow for communication across generations. Though substantive differences exist between each of the models, they both prove standard (for OLG games) Folk

³⁶Pearce (1992) is also an excellent source for further references.

³⁷Baliga, Corchon, and Sjostrom (1997) use a similar type of mechanism in another model of communication when there are three or more players.

³⁸See also the early contribution by Cremer (1986). Notice also that this literature has focused, among other things, on the role of dimensionality assumptions on the stage game, which can be dispensed with in some cases. This is not the case in our model (see footnote 4 above).

³⁹As he points out, the results in Kandori (1992a) generalize to less than full memory of the past, although some direct memory is required. Bhaskar (1998) examines a related OLG model with no dynastic payoffs and very little, albeit some, direct memory by entrants. He shows that very limited memory is enough to sustain optimal transfers in a 2-period consumption-loan smoothing OLG game.

Theorems. Interestingly, both Folk Theorems make use of *intra*-generational disagreement of beliefs in the equilibrium continuations following deviations. Nevertheless, the constructed equilibria in both papers leave no room for *inter*-generational disagreement at the message stage.

The role of public messages has been studied in other repeated game contexts such as in games with private monitoring. Models of Ben-Porath and Kahneman (1996), Compte (1998), and Kandori and Matsushima (1998) all examine communication in repeated games when players receive private signals of others’ past behavior. As in Anderlini and Lagunoff (2005), these papers exploit cross-checking arguments to sustain the truthful revelation of one’s private signal in each stage of the repeated game.

Recent works by Schotter and Sopher (2001a), Schotter and Sopher (2001b), and Chaudhuri, Schotter, and Sopher (2001) examine the role of communication in an experimental dynastic environment. These papers set up laboratory experiments designed to mimic the dynastic game. The general conclusion seems to be that the presence of private communication has a significant (if puzzling) effect, even in the full memory game.

It is also worth noting the similarity between the present model and games with imperfect recall.⁴⁰ Each dynastic player could be viewed as an infinitely lived player with imperfect recall (e.g., the “absent-minded driver” with “multiple selves” in Piccione and Rubinstein (1997)) who can write messages to his future self at the end of each period.

By contrast, the present model is distinguishable from dynamic models that create memory from a tangible “piece” of history. For instance, Anderlini and Lagunoff (2005) extend the analysis to the case where history may leave “footprint,” i.e, hard evidence of the past history of play. Incomplete but hard evidence of the past history of play is also present in Johnson, Levine, and Pesendorfer (2001) and Kandori (1992b). In another instance, memory may be created from a tangible but intrinsically worthless asset such as fiat money. A number of contributions in monetary theory (e.g., Kocherlakota (1998), Kocherlakota and Wallace (1998), Wallace (2001), and Corbae, Temzelides, and Wright (2001)) have all shown, to varying degrees, the substitutability of money for memory.

10. Concluding Remarks

We posit a dynastic repeated game populated by one-period-lived individuals who rely on private messages from their predecessors to fathom the past. The set of equilibrium payoffs expands dramatically relative to the corresponding standard repeated game. Under extremely mild conditions, as the dynastic players care more and more about the payoffs of their successors, *all* interior payoff vectors that are *feasible* in the stage game are sustainable in an SE of the dynastic repeated game.

We are able to characterize entirely, via a property of the players’ beliefs, when an SE of the dynastic repeated game can yield a payoff vector not sustainable as an SPE of the

⁴⁰See the Special Issue of Games and Economic Behavior (1997) on Games with Imperfect Recall for extensive references.

standard repeated game: the SE in question must display a failure of what we have termed Inter-Generational Agreement.⁴¹

It is natural to ask why we focus on SE as opposed to other solution concepts. The answer is two-fold. First the concept of SE is an extremely widely accepted and used benchmark in the literature. Hence it seems an appropriate point to start. Second, the surprising features of some of the equilibria of our dynastic repeated game arise from a failure of Inter-Generational Agreement. Broadly speaking, failures of Inter-Generational Agreement can be traced back to the fact that, upon receiving an unexpected private message from his predecessor, a player always has to weigh two sets of possibilities. The message he observes could have been generated by a deviation at the action stage, or it could have been generated by a deviation at the message stage. The concept of SE requires all players to have a complete (and common across players) theory of the mistakes that might have caused deviations from equilibrium. Therefore, it seems like a prime candidate to require that the players’ beliefs should be plausible in an intuitive sense in the dynastic repeated game we analyze.

It is important to comment further on the relationship between our results and the idea of “neologism-proofness” that has been put forth in the literature (Farrell 1993, Mattheus, Okuno-Fujiwara, and Postlewaite 1991, among others). As we mentioned above, at least in its current form, neologism-proofness simply does not apply to our framework. The reason is simple. Roughly speaking, neologism-proofness builds into the solution concept the idea that in a sender-receiver game, provided the appropriate incentive-compatibility constraints are satisfied, a player’s *exogenous type* (in the standard sense of a “payoff type”) will be able to create a “neologism” (use an hitherto unused message) to distinguish himself from other types. The point is that in our dynastic game there are *no payoff types* for any of the players. It would therefore be impossible to satisfy any standard form of incentive-compatibility constraints. The different “types” of each player in our dynastic repeated game are only distinguished by their beliefs, which in turn are determined by equilibrium strategies together with a complete theory of mistakes as in any SE. To see that the logic of neologism-proofness can be conceptually troublesome in our context, consider for instance the construction we use to prove Theorem 3 above. Suppose that some player $\langle i, t \rangle$ deviates so as to trigger the punishment- i phase. At the end of period t player $\langle i, t \rangle$ may want to communicate to player $\langle i, t + 1 \rangle$ that play is in the punishment- i phase so that he can play a best response to the actions of others in period $t + 1$. For a “neologism” to work at this point player $\langle i, t + 1 \rangle$ would have to believe it. He would have to believe what player $\langle i, t \rangle$ is saying: I have made a mistake, therefore respond appropriately to the punishment that follows (there are no exogenous types to which $\langle i, t \rangle$ can appeal in his “speech”). However, as always in a dynastic game, whether $\langle i, t + 1 \rangle$ believes or not what he is told, depends on

⁴¹The logic of our Theorems 1 and 4 extends easily to *finite* games played by dynastic players who rely on private messages. Thus the appearance of “new” equilibria as a result of the presence of dynastic players is not uniquely a feature of *infinitely* repeated dynastic games. In all cases the new equilibria can be traced to the mis-match in beliefs that we have analyzed here. The fact that *all* feasible payoff vectors can be sustained in equilibrium is obviously not true in general for finite games.

the relative likelihood that he assigns to mistaken actions and mistaken messages; both are possible after all.⁴² So, for the neologism to work it would have to be “trusted” more as a message than the one prescribed in equilibrium. But, since there are no exogenous types to which to appeal, there do not seem to be compelling reasons for this to be the case. As with other possible refinements, our characterization of the new equilibria that appear in the dynastic repeated game in terms of Inter-Generational Agreement also seals the question of what bite a possible adaptation of neologism-proofness could have at a more general level. The new equilibria of the dynastic repeated game can be ruled out (without ruling out any of the traditional ones) if and only if beliefs that violate Inter-Generational Agreement can be ruled out. Whether this is the case or not is largely a matter of intuitive appeal.

While our results apply only to the actual formal model we have set forth, it is natural to ask which ones are essential and which ones are not. We have several remarks to make.

As we noted already, the absence of public messages alongside private ones is completely inessential to what we do here. Public messages could be added to our model without altering our results. It is always possible to replicate the SE of this model in another model with public messages as well; the public messages would be ignored by the players’ equilibrium strategies and beliefs.

We make explicit use of public randomization devices both at the action and at the message stage of the dynastic repeated game. While the use of two separate devices is not essential for our results (see footnote 13 above), the use of some public randomization device probably is.^{43,44} In our constructions it is essential that the players should be able to correlate the messages they send to the next cohort. Without this it is hard to see how play could switch between the different phases on which our constructions depend.

We have stipulated a very specific set of “demographics” for our dynastic repeated game: all players live one period and are replaced by their successor at the end of their lives. We believe that this is not essential to the qualitative nature of our results, although some features of the demographic structure of the model play an important role. We conjecture that results similar to ours would hold in a dynastic repeated game in which all players are simultaneously replaced by their successors in the same dynasty with positive probability, thus ensuring that looking forward from any point in time total replacement happens with probability one at some future date.

The actual history of play leaves no visible trace in our model. To what extent would our results change if a (noisy or incomplete, or both) “footprint” of the past history of play were available to the players? Intuitively this would make failures of Inter-Generational Agreement

⁴²By contrast, in a sender-receiver game the sender communicates to the receiver his *exogenous type*, which is chosen by Nature, and not by the sender himself.

⁴³This is in contrast with Fudenberg and Maskin (1986) in which “time averaging” across different pure action profiles in each period can (approximately) substitute for a randomization device.

⁴⁴We have examples showing that even without any randomization devices it is possible to display SE that push one or more dynasties below their minmax in the stage game. Whether a “Super” Folk Theorem is available in this case is an open question at this point.

harder to generate in equilibrium. Does the set of equilibria of the dynastic repeated game shrink as the direct information about the past history of play that is available to the players becomes more and more precise? While it seems plausible to conjecture that our results would change across this range of possibilities, we have no option but to say that the topic is worthy of future research.

Our Folk Theorems for the dynastic repeated game assume three or more dynasties in one case, and four or more in another. Is this essential to our results? (It clearly is essential to our construction in both cases.) It is not hard to construct examples of dynastic repeated games with two dynasties that admit SE in which the players’ payoffs are below their minmax in the stage game. Thus, it seems that there is no definite need to have more than two dynasties to generate SE payoffs in the dynastic repeated game that are not sustainable as SPE of the corresponding standard repeated game. Whether and under what conditions a Folk Theorem like the ones presented here is available for the case of two dynasties is an open question. We leave the characterization of the equilibrium set in this case for future work.

Lastly, our Folk Theorems for the dynastic repeated game show that, as δ approaches one, the set of SE payoffs includes “worse” vectors that push some (or even all) players below their minmax payoffs in the stage game. We do not have a full characterization of the SE payoffs for the dynastic repeated game when δ is bounded away from one. However, it is possible to construct examples showing that the set of SE payoffs includes vectors that Pareto-dominate those on the Pareto-frontier of the set of SPE payoffs in the standard repeated game when δ is bounded away from one. Intuitively, this is because some “bad” payoff vectors (pushing some players below the minmax) are sustainable in an SE when δ is bounded away from one. Thus, “harsher” punishments are available as continuation payoffs in the dynastic repeated game than in the standard repeated game. Using these punishments, higher payoffs are sustainable in equilibrium in the dynastic repeated game.

Appendix A: The Basics of the Proof of Theorem 2

A.1. Outline

Our proof is constructive. It runs along the following lines. Given a $v^* \in \text{int}(V)$ and a $\delta \in (0, 1)$ we construct a randomization device \tilde{x} , a randomization device $\tilde{y}(\delta)$, and an assessment (g, μ, Φ^δ) , which implements v^* , and which for δ sufficiently large constitutes an SE of the dynastic repeated game.

Notice that the profile (g, μ) is defined independently of δ , as is \tilde{x} . On the other hand, the probability distribution of the message-stage randomization device $\tilde{y}(\delta)$ and the system of beliefs Φ^δ are defined using δ as a parameter. This is simply a property of our construction. Given that our task is to show that there exists a $\underline{\delta} \in (0, 1)$ such that $\delta > \underline{\delta}$ implies $v^* \in \mathcal{E}^D(\delta) = \bigcup_{\tilde{x}, \tilde{y}} \mathcal{E}^D(\delta, \tilde{x}, \tilde{y})$, if (g, μ) and \tilde{x} were also dependent on δ this would not matter at all. Sometimes, when it does not cause any ambiguity, δ will be dropped from the notation for the message-stage randomization device and/or the system of beliefs.

Throughout the argument, we assume that the message space for any player $\langle i, t \rangle$ consists of three elements only. Formally, we let $M_i^t = \{m^*, m^A, m^B\}$. Since we assume that $n \geq 3$ and that $\|A_i\| \geq 2$ for every i , it is obvious that $\|H^t\| > 3$ for every $t \geq 1$. Therefore we can think of each M_i^t as a “restricted message space” in the sense of Definition T.2.1. It then follows from Lemma T.2.1 that proving the result for these restricted message spaces is sufficient to prove our claim for the case $M_i^t = H^t$, as required by the statement of the theorem.

In Section A.2 we define formally the randomization devices \tilde{x} and $\tilde{y}(\delta)$ and the profile (g, μ) . In Section A.3 we define the completely mixed strategies that generate the SE beliefs Φ^δ . In Section T.5 we define formally the system of beliefs Φ^δ . In Section T.6 we check that the assessment (g, μ, Φ^δ) satisfies sequential rationality when δ is close to one. Finally in Section T.7 we verify the consistency of the equilibrium beliefs.

Throughout the argument, $v^* \in \text{int}(V)$ is the vector of payoffs to be sustained as an SE, and a^* and a' are two action profiles as in the statement of the theorem.

A.2. Strategies and Randomization Devices

Definition A.1: Let $(a(1), \dots, a(\ell), \dots, a(\|A\|))$ be a list of all possible outcomes in G . Without loss of generality assume that the first element of this list is a^* so that $a(1) = a^*$. We can then define a set of weights $\{p(a(\ell))\}_{\ell=1}^{\|A\|}$ adding up to one and such that

$$v^* = \sum_{\ell=1}^{\|A\|} p(a(\ell))u(a(\ell)) \quad (\text{A.1})$$

Notice that since $v^* \in \text{int}(V)$, we can assume that $p(a(\ell)) > 0$ for all $\ell = 1, \dots, \|A\|$.

Since it will play a special role, we assign a separate symbol to $p(a(1))$ (the weight of a^* in equation (A.1)). In what follows we set $p(a(1)) = q$.

Definition A.2. *Action-Stage Randomization Device:* Given v^* , the action-stage randomization device \tilde{x} is defined as follows. The set X consists of $\|A\|$ elements denoted by $(x(1), \dots, x(\ell), \dots, x(\|A\|))$. We then set

$$\Pr(\tilde{x} = x(\ell)) = p(a(\ell)) \quad (\text{A.2})$$

where each $p(a(\ell))$ is as in Definition A.1.

Definition A.3: Let $\underline{u}_i = \min_{a \in A} u_i(a)$ and $\bar{u}_i = \max_{a \in A} u_i(a)$.⁴⁵ Let also α be a real number in $(0, 1)$ such that

$$\alpha < \frac{u_i(a') - u_i(a_i^*, a'_{-i})}{u_i(a') - u_i(a_i^*, a'_{-i}) + \bar{u}_i - \underline{u}_i} \quad \forall i \in I \quad (\text{A.3})$$

Definition A.4. *Message-Stage Randomization Device:* Given δ the message-stage randomization device $\tilde{y}(\delta)$ is defined as follows. The set Y consists of two elements, $y(0)$ and $y(1)$.

Let α be as in Definition A.3. Now define

$$\gamma(\delta) = \frac{(1 - \delta)(1 - \alpha)}{\alpha \delta} \quad (\text{A.4})$$

Notice that, given α , as δ increases towards one, clearly $\gamma(\delta)$ is between zero and one. We then take it to be the case that

$$\begin{aligned} \Pr[\tilde{y}(\delta) = y(0)] &= 1 - \gamma(\delta) \\ \Pr[\tilde{y}(\delta) = y(1)] &= \gamma(\delta) \end{aligned} \quad (\text{A.5})$$

⁴⁵Throughout the rest of the paper and the technical addendum, we will use the symbols \underline{u}_i and \bar{u}_i with the meaning we have just defined.

Definition A.5. *Action-Stage Strategies:* Let $\{a(\ell)\}_{\ell=1}^{\|A\|}$ be an enumeration of the strategy profiles in G as in Definition A.1, and let \tilde{x} be as in Definition A.2. Recall that at the beginning of period 0 all players receive the null message m_i^0 . For all players $\langle i, 0 \rangle$ define

$$g_i^0(m_i^0, x^0) = a_i(\ell) \quad \text{if } x^0 = x(\ell) \quad (\text{A.6})$$

and for any player $\langle i, t \rangle$ with $t \geq 1$ define

$$g_i^t(m_i^t, x^t) = \begin{cases} a_i^* & \text{if } m_i^t = m^* \text{ and } x^t = x(1) \\ a_i' & \text{if } m_i^t \in \{m^A, m^B\} \text{ and } x^t = x(1) \\ a_i(\ell) & \text{if } x^t = x(\ell) \text{ with } \ell \geq 2 \end{cases} \quad (\text{A.7})$$

Definition A.6. *Message-Stage Strategies:* Let $\{a(\ell)\}_{\ell=1}^{\|A\|}$ be an enumeration of the strategy profiles in G as in Definition A.1. Let \tilde{x} be as in Definition A.2. Let \tilde{y} be as in Definition A.4, where the dependence on δ has been suppressed for notational convenience.

To describe the message-stage strategies it is convenient to distinguish between two cases.⁴⁶

For any player $\langle i, t \rangle$, whenever $x^t = x(1)$ let⁴⁷

$$\mu_i^t(m_i^t, x(1), a^t, y^t) = \begin{cases} m^A & \text{if } y^t = y(0) \text{ and } a^t = (a'_j, a^*_{-j}) \text{ for some } j \in I \\ m^A & \text{if } m_i^t = m^A, y^t = y(0) \text{ and } a^t = a' \\ m^A & \text{if } m_i^t = m^A, y^t = y(0) \text{ and } a^t = (a^*_j, a'_{-j}) \text{ for some } j \in I \\ m^B & \text{if } m_i^t = m^B \text{ and } a^t = a' \\ m^B & \text{if } m_i^t = m^B, \text{ and } a^t = (a^*_j, a'_{-j}) \text{ for some } j \in I \\ m^B & \text{if } a^t_j \notin \{a^*_j, a'_j\} \text{ for some } j \in I \\ m^* & \text{otherwise} \end{cases} \quad (\text{A.8})$$

For any player $\langle i, t \rangle$, whenever $x^t = x(\ell)$ with $\ell \geq 2$ let

$$\mu_i^t(m_i^t, x(\ell), a^t, y^t) = \begin{cases} m^* & \text{if } m_i^t = m^*, \text{ and } a^t = a(\ell) \\ m^* & \text{if } m_i^t = m^A, y^t = y(1) \text{ and } a^t = a(\ell) \\ m^A & \text{if } m_i^t = m^A, y^t = y(0) \text{ and } a^t = a(\ell) \\ m^B & \text{otherwise} \end{cases} \quad (\text{A.9})$$

A.3. Completely Mixed Strategies

Definition A.7: Throughout this section we let ε denote a small positive number, which parameterizes the completely mixed strategies that we construct. It should be understood that our construction of beliefs involves the limit $\varepsilon \rightarrow 0$.

⁴⁶In the interest of brevity, we avoid writing down the message-stage strategies for players $\langle i \in I, 0 \rangle$ separately. Equations (A.8) and (A.9) that follow can be interpreted as defining the profile μ^0 by re-defining m_i^0 to be equal to m^* for all players $\langle i \in I, 0 \rangle$.

⁴⁷Notice that to distinguish between the first and the fifth case, and between the first and the last case, on the right-hand side of (A.8), player $\langle i, t \rangle$ must be able to distinguish between an action profile of the type (a'_j, a^*_{-j}) and an action profile of the type (a^*_j, a'_{-j}) . This is always possible because of our assumption that $n \geq 3$, and because our assumptions about a^* and a' obviously imply that $a^*_j \neq a'_j$ for every $j \in I$.

For every $t \geq 0$, given ε , define

$$\varepsilon_t = \frac{1}{\left[\max_{i \in I} \|A_i\| \right]^2} \varepsilon \frac{1}{n^{2t}} \quad (\text{A.10})$$

Notice also that in Definition A.9 below we make use of the quantity $\psi_i(\delta)$ that is defined in equation (T.7.1). What needs to be specified at this point is that, for δ sufficiently close to one, $\psi_i(\delta)$ is a real number in $(0, 1)$ for every $i \in I$. (See Remark T.7.1.)

Definition A.8. *Completely Mixed Action Strategies:* Given ε , and ε_t as in (A.10), the completely mixed strategies for all players $\langle i, t \rangle$ at the action stage are denoted by $g_{i,\varepsilon}^t$ and are defined as follows.

After receiving message m_i^0 and observing the realization x^t of the action-stage randomization device any player $\langle i, 0 \rangle$ plays the action prescribed by the proposed equilibrium strategy described in (A.6) with probability $1 - \varepsilon_t(\|A_i\| - 1)$ and plays all other actions in A_i with probability ε_t each.

After receiving message m_i^* and observing the realization x^t of the action-stage randomization device any player $\langle i, t \rangle$ plays the action prescribed by the proposed equilibrium strategy described in (A.7) with probability $1 - \varepsilon_t(\|A_i\| - 1)$ and plays all other actions in A_i with probability ε_t each.

After receiving any message $m_i^t \neq m_i^*$ and observing the realization x^t of the action-stage randomization device any player $\langle i, t \rangle$ plays the action prescribed by the proposed equilibrium strategy described in (A.7) with probability $1 - \sqrt{\varepsilon_t}(\|A_i\| - 1)$ and plays all other actions in A_i with probability $\sqrt{\varepsilon_t}$ each.

Definition A.9. *Completely Mixed Message Strategies:* Given ε , the completely mixed strategies for all players $\langle i, t \rangle$ at the message stage are denoted by $\mu_{i,\varepsilon}^t$ and are defined as follows.⁴⁸

Suppose that player $\langle i, t \rangle$ receives message m_i^* and that after observing (x^t, a^t, y^t) the proposed equilibrium strategy described in Definition A.6 prescribes that he should send message m^* . Then player $\langle i, t \rangle$ sends message m^* with probability $1 - \varepsilon_t^2 - \psi_i(\delta)\varepsilon_t$, sends message m^A with probability ε_t^2 , and sends message m^B with probability $\psi_i(\delta)\varepsilon_t$.

Suppose that player $\langle i, t \rangle$ receives message m_i^* and that after observing (x^t, a^t, y^t) the proposed equilibrium strategy described in Definition A.6 prescribes that he should send message m^A . Then player $\langle i, t \rangle$ sends message m^A with probability $1 - \varepsilon_t - \varepsilon_t^2$, sends message m^* with probability ε_t , and sends message m^B with probability ε_t^2 .

Suppose that player $\langle i, t \rangle$ receives message m_i^* and that after observing (x^t, a^t, y^t) the proposed equilibrium strategy described in Definition A.6 prescribes that he should send message m^B . Then player $\langle i, t \rangle$ sends message m^A with probability $1 - \varepsilon_t - \varepsilon_t^2$, sends message m^* with probability ε_t , and sends message m^A with probability ε_t^2 .

Finally, after receiving any message $m_i^t \neq m_i^*$ and observing (x^t, a^t, y^t) , player $\langle i, t \rangle$ sends the message prescribed by the equilibrium strategy described in Definition A.6 with probability $1 - 2\varepsilon_t$, and sends each of the other two messages with probability ε_t .

⁴⁸In the interest of brevity, we avoid an explicit distinction between the $t = 0$ players and all others. What follows can be interpreted as applying to all players re-defining m_i^0 to be equal to m^* for players $\langle i \in I, 0 \rangle$.

Appendix B: The Basics of the Proof of Theorem 3

B.1. Outline

Once again, our proof is constructive. It runs along the following lines. Given a $v^* \in \text{int}(V)$, we construct a randomization device \tilde{x} , a randomization device \tilde{y} , and an assessment (g, μ, Φ) , which implements the vector of payoffs v^* , and which for δ sufficiently large constitutes an SE of the dynastic repeated game.

Notice that all the elements of our construction are defined independently of δ . The sequential rationality of the strategy profile given the postulated beliefs holds when δ is sufficiently close to one.

Throughout the argument, we assume that the message space for any player $(i, t - 1)$ consists of a set smaller than the set H^t . Therefore we can think of each M_i^t as a “restricted message space” in the sense of Definition T.2.1. It then follows from Lemma T.2.1 that proving the result for these restricted message spaces is sufficient to prove our claim for the case $M_i^t = H^t$, as required by the statement of the theorem.

In Section B.2 we define formally the randomization devices \tilde{x} and \tilde{y} , the players’ message spaces M_i^t and the profile (g, μ) . In Section B.3 we define the completely mixed strategies that generate the SE beliefs Φ . In Section T.9 we define formally the system of beliefs Φ . In Section T.10 we check that the assessment (g, μ, Φ) satisfies sequential rationality when δ is close to one. Finally in Section T.11 we verify the consistency of the equilibrium beliefs.

In what follows, $v^* \in \text{int}(V)$ is the vector of payoffs to be sustained as an SE as in the statement of the theorem. Throughout the argument, \tilde{A} is a product set and \hat{v} and \bar{v}^1 through \bar{v}^n are vectors of payoffs as in the statement of the theorem. Of course, these are fixed throughout the proof.

B.2. Strategies, Randomization Devices and Message Spaces

Definition B.1: As in Definition A.1, let $(a(1), \dots, a(\ell), \dots, a(\|A\|))$ be a list of all possible outcomes in G . Without loss of generality, assume that the first $\|\tilde{A}\| \leq \|A\|$ elements in this enumeration are the strategy profiles in the product set \tilde{A} . This enumeration will be taken as fixed throughout the rest of the argument.

Remark B.1: From Definition 3, for each $i \in I$ we can find a set of profiles of actions $\tilde{A}^i \subset \tilde{A}$ corresponding to a set of indices $(i_1, \dots, i_\ell, \dots, i_{\|\tilde{A}^i\|})$ in the enumeration of Definition B.1, and a set of positive weights $\{p^i(a(i_\ell))\}_{\ell=1}^{\|\tilde{A}^i\|}$ adding up to one and such that

$$\underline{\omega}_i(\tilde{A}) = \max_{a_i \in \tilde{A}^i} \sum_{\ell=1}^{\|\tilde{A}^i\|} p^i(a(i_\ell)) u_i(a_i, a_{-i}(i_\ell)) = \sum_{\ell=1}^{\|\tilde{A}^i\|} p^i(a(i_\ell)) u_i(a(i_\ell)) \quad (\text{B.1})$$

Without loss of generality, we can take it to be the case that for every i_ℓ , $a_i(i_\ell)$ is the same action in \tilde{A}^i . We denote this by a_i^i so that $a_i(i_\ell) = a_i^i$ for $\ell = 1, \dots, \|\tilde{A}^i\|$.

For convenience, since \tilde{A} is fixed throughout the argument, in what follows we will use the following notation for the payoffs of each i corresponding to the weights $\{p^j(a(j_\ell))\}_{\ell=1}^{\|\tilde{A}^j\|}$.

$$\underline{\omega}_i^j = \sum_{\ell=1}^{\|\tilde{A}^j\|} p^j(a(j_\ell)) u_i(a(j_\ell)) \quad (\text{B.2})$$

Of course, we have that $\underline{\omega}_i^i = \underline{\omega}_i(\tilde{A})$.

Definition B.2: Let \tilde{A}^i be as in Remark B.1. For each $i \in I$ and for each element $a(i_\ell)$ of \tilde{A}^i , construct a new action profile $\check{a}^i(i_\ell)$ as follows. For all $j \neq i$, set $\check{a}_j^i(i_\ell)$ to satisfy $\check{a}_j^i(i_\ell) \neq a_j(i_\ell)$ and $\check{a}_j^i(i_\ell) \in \tilde{A}_j$. Notice that this is always possible since, by assumption, \tilde{A}_j contains at least two elements for every $j \in I$. Finally, set $\check{a}_i^i(i_\ell) = a_i^i(i_\ell) = a_i^i$.

In what follows, for every i and j in I we will let

$$\check{u}_i^j = \sum_{\ell=1}^{\|\tilde{A}^j\|} p^j(a(j_\ell)) u_i(\check{a}^j(j_\ell)) \quad (\text{B.3})$$

Remark B.2: Since each of the payoff vectors \bar{v}^j must only satisfy strong inequalities (see (iii) of the statement of the theorem), without loss of generality we can take it to be the case that $\bar{v}^j \in \text{int}(V(\tilde{A}))$, for each $j \in I$. It then follows that for every $j \in I$ we can find a set of positive weights $\{\bar{p}^j(a(\ell))\}_{\ell=1}^{\|\tilde{A}\|}$ adding up to one and such that for every $i \in I$

$$\bar{v}_i^j = \sum_{\ell=1}^{\|\tilde{A}\|} \bar{p}^j(a(\ell)) u_i(a(\ell)) \quad (\text{B.4})$$

Remark B.3: Since the payoff vector \hat{v} is in $\text{int}(V(\tilde{A}))$, we can find an $\eta \in (0, 1)$ and a set of positive weights $\{\hat{p}(a(\ell))\}_{\ell=1}^{\|\tilde{A}\|}$ adding up to one and such that for every $i \in I$

$$\hat{v}_i = (1 - \eta) \sum_{\ell=1}^{\|\tilde{A}\|} \hat{p}(a(\ell)) u_i(a(\ell)) + \frac{\eta}{n} \sum_{j=1}^n \check{u}_i^j \quad (\text{B.5})$$

Definition B.3: Throughout the rest of the argument we let

$$\hat{v}_i = \sum_{\ell=1}^{\|\tilde{A}\|} \hat{p}(a(\ell)) u_i(a(\ell)) \quad (\text{B.6})$$

where the weights $\{\hat{p}(a(\ell))\}_{\ell=1}^{\|\tilde{A}\|}$ are as in Remark B.3.

Remark B.4: Since the payoff vector v^* is in $\text{int}(V)$, we can find a $q \in (0, 1)$ and a set of positive weights $\{p^*(a(\ell))\}_{\ell=1}^{\|A\|}$ adding up to one and such that for every $i \in I$

$$v_i^* = (1 - q) \sum_{\ell=1}^{\|A\|} p^*(a(\ell)) u_i(a(\ell)) + q \hat{v}_i \quad (\text{B.7})$$

In what follows we will let $z_i = \sum_{\ell=1}^{\|A\|} p^*(a(\ell)) u_i(a(\ell))$.

Remark B.5: Recall that by assumption the payoff vector v^* is in $\text{int}(V)$. Therefore, we can find a set of positive weights $\{p^0(a(\ell))\}_{\ell=1}^{\|A\|}$ adding up to one and such that for every $i \in I$

$$v_i^* = \sum_{\ell=1}^{\|A\|} p^0(a(\ell)) u_i(a(\ell)) \quad (\text{B.8})$$

Definition B.4. *Action-stage Randomization Device:* The action stage randomization device \tilde{x} is defined as follows. The set X consists of $\|A\| \|\tilde{A}\|^{n+1} \prod_{i \in I} \|\tilde{A}^i\| + \|A\|^2$ elements.

Let $(x(1), \dots, x(\kappa), \dots, x(\|X\|))$ be an enumeration of the elements of X , and let $\bar{\kappa} = \|A\| \|\tilde{A}\|^{n+1} \prod_{i \in I} \|\tilde{A}^i\|$.

Each of the first $\bar{\kappa}$ elements of X can be identified by a string of $1 + (n + 1) + n = 2n + 2$ indices as follows. With a slight abuse of notation, for $\kappa \leq \bar{\kappa}$, we will write $x(\kappa) = x(\ell_0, \hat{\ell}, \ell_1, \dots, \ell_n, 1_\ell, \dots, i_\ell, \dots, n_\ell)$

with ℓ_0 running from 1 to $\|A\|$, $\hat{\ell}$ and each of the indices ℓ_1 through ℓ_n running from 1 to $\|\tilde{A}\|$, and each of the n indices i_ℓ , each with ℓ running from 1 to $\|\tilde{A}^i\|$. Obviously, the last $\|X\| - \bar{\kappa}$ elements of X can be identified by a pair of indices ℓ_{00} and ℓ^* both running from 1 to $\|A\|$. In this case, with a slight abuse of notation again, we will write $x(\kappa) = x(\ell_{00}, \ell^*)$.

We are now ready to define the probability distribution governing the realization of \tilde{x} . For $\kappa \leq \bar{\kappa}$ let

$$\Pr(\tilde{x} = x(\ell_0, \hat{\ell}, \ell_1, \dots, \ell_n, 1_\ell, \dots, i_\ell, \dots, n_\ell)) = q \left[p^0(a(\ell_0)) \hat{p}(a(\hat{\ell})) \bar{p}^1(a(\ell_1)) \cdots \bar{p}^n(a(\ell_n)) p^1(a(1_\ell)) \cdots p^i(a(i_\ell)) \cdots p^n(a(n_\ell)) \right] \quad (\text{B.9})$$

and for $\kappa = \bar{\kappa} + 1, \dots, \|X\|$ let

$$\Pr(\tilde{x} = x(\ell_{00}, \ell^*)) = (1 - q) [p^0(a(\ell_{00})) p^*(a(\ell^*))] \quad (\text{B.10})$$

Definition B.5. *Message-Stage Randomization Device:* The message stage randomization device \tilde{y} is defined as follows. The set Y consists of $n + 1$ elements, which we denote $(y(0), y(1), \dots, y(n))$. The random variable \tilde{y} takes value $y(0)$ with probability $1 - \eta$, and each of the other possible values with probability η/n .

Definition B.6: Throughout the rest of the argument, we let T be an integer sufficiently large so as to guarantee that the following inequality is satisfied for all $i \in I$.

$$T(\bar{v}_i^i - \underline{\omega}_i^i) > \bar{u}_i - \underline{u}_i \quad (\text{B.11})$$

Definition B.7. *Message Spaces:* For each $j \in I$ and each $t = 1, \dots, T - 1$ let

$$\underline{M}(j, t) = \{\underline{m}^{j, T-t+1}, \underline{m}^{j, T-t+2}, \dots, \underline{m}^{j, T}\} \quad (\text{B.12})$$

and for every $t \geq T$ let

$$\underline{M}(j, t) = \underline{M}(j, T) = \{\underline{m}^{j, 1}, \dots, \underline{m}^{j, T}\} \quad (\text{B.13})$$

Define also $\bar{M} = \{\bar{m}^1, \dots, \bar{m}^n\}$, and $\check{M} = \{\check{m}^1, \dots, \check{m}^n\}$, and recall that according to our notational convention when we write \check{M}_{-i} we mean $\{\check{m}^1, \dots, \check{m}^{i-1}, \check{m}^{i+1}, \check{m}^n\}$.

Recall that M_i^t denotes the set of messages that a player $\langle i, t - 1 \rangle$ can send to player $\langle i, t \rangle$. For any $t = 1, \dots, T$ let

$$M_i^t = \{m^*\} \cup \check{M}_{-i} \cup \underline{M}(1, t) \cup \dots \cup \underline{M}(n, t) \quad (\text{B.14})$$

For any $t \geq T + 1$ let

$$M_i^t = \{m^*\} \cup \check{M}_{-i} \cup \underline{M}(1, t) \cup \dots \cup \underline{M}(n, t) \cup \bar{M} \quad (\text{B.15})$$

Definition B.8. *Action-Stage Strategies:* Let k be an element of I , and j be an element of I not equal to i .

Let $(a(1), \dots, a(\|A\|))$ be the enumeration of the elements of A of Definition B.1 and consider the indexing of the elements of X in Definition B.4, according to whether $x(\kappa)$ has $\kappa \leq \bar{\kappa}$ or not.

Recall that at the beginning of period 0 all players $\langle i \in I, 0 \rangle$ receive message $m_i^0 = \emptyset$. For all players $\langle i \in I, 0 \rangle$ then define

$$g_i^0(m_i^0, x^0) = \begin{cases} a_i(\ell_0) & \text{if } x^0 = x(\ell_0, \dots) \\ a_i(\ell_{00}) & \text{if } x^0 = x(\ell_{00}, \cdot) \end{cases} \quad (\text{B.16})$$

Now consider any player $\langle i, t \rangle$ with $t \geq 1$. It is convenient to distinguish between the two cases $x^t = x(\kappa)$ with $\kappa \leq \bar{\kappa}$ and with $\kappa > \bar{\kappa}$.

For any $i \in I$ and $t \geq 1$ whenever $x^t = x(\kappa)$ with $\kappa \leq \bar{\kappa}$ define⁴⁹

$$g_i^t(m_i^t, x^t) = \begin{cases} a_i(\hat{\ell}) & \text{if } m_i^t = m^* & \text{and } x^t = x(\cdot, \hat{\ell}, \dots) \\ \check{a}_i^j(j_\ell) & \text{if } m_i^t = \check{m}_j & \text{and } x^t = x(\dots, j_\ell, \dots) \\ a_i(\ell_k) & \text{if } m_i^t = \bar{m}^k & \text{and } x^t = x(\dots, \ell_k, \dots) \\ a_i(k_\ell) & \text{if } m_i^t \in \underline{M}(k, t) & \text{and } x^t = x(\dots, k_\ell, \dots) \end{cases} \quad (\text{B.17})$$

For any $i \in I$, $t \geq 1$ and m_i^t , whenever $x^t = x(\kappa)$ with $\kappa > \bar{\kappa}$ define

$$g_i^t(m_i^t, x^t) = a_i(\ell^*) \text{ if } x^t = x(\cdot, \ell^*) \quad (\text{B.18})$$

Definition B.9. Message-Stage Strategies: Let k be any element of I , and j be any element of I not equal to i .

We begin with period $t = 0$. Recall that $m_i^0 = \emptyset$ for all $i \in I$. Let also $g^0(m^0, x^0) = (g_1^0(m_1^0, x^0), \dots, g_n^0(m_n^0, x^0))$, and define $g_{-k}^0(m^0, x^0)$ in the obvious way.

We let

$$\mu_i^0(m_i^0, x^0, a^0, y^0) = \begin{cases} \check{m}^j & \text{if } a^0 = g^0(m^0, x^0) & \text{and } y^0 = y(j) \\ \underline{m}^{i,T} & \text{if } a^0 = g^0(m^0, x^0) & \text{and } y^0 = y(i) \\ \underline{m}^{k,T} & \text{if } a_{-k}^0 = g_{-k}^0(m^0, x^0) & \text{and } a_k^0 \neq g_k^0(m^0, x^0) \\ m^* & \text{otherwise} \end{cases} \quad (\text{B.19})$$

For the periods $t \geq 1$ it is convenient to distinguish between several cases. Assume first that $x^t = x^t(\kappa)$ with $\kappa > \bar{\kappa}$. Let

$$\mu_i^t(m_i^t, x^t, a^t, y^t) = \begin{cases} m_i^t & \text{if } x^t = x(\cdot, \ell^*) & \text{and } a^t = a(\ell^*) \\ \underline{m}^{k,T} & \text{if } x^t = x(\cdot, \ell^*), a_{-k}^t = a_{-k}(\ell^*) & \text{and } a_k^t \neq a_k(\ell^*) \\ m^* & \text{otherwise} \end{cases} \quad (\text{B.20})$$

Now consider the case $x^t = x^t(\kappa)$ with $\kappa \leq \bar{\kappa}$. We divide this case into several subcases, according to which message player $\langle i, t \rangle$ has received. We begin with $m_i^t = m^*$. Let⁵⁰

$$\mu_i^t(m^*, x^t, a^t, y^t) = \begin{cases} \check{m}^j & \text{if } x^t = x(\cdot, \hat{\ell}, \dots), a^t = a(\hat{\ell}) & \text{and } y^t = y(j) \\ \nu(\underline{M}(i, t+1)) & \text{if } x^t = x(\cdot, \hat{\ell}, \dots), a^t = a(\hat{\ell}) & \text{and } y^t = y(i) \\ \underline{m}^{k,T} & \text{if } x^t = x(\cdot, \hat{\ell}, \dots), a_{-k}^t = a_{-k}(\hat{\ell}) & \text{and } a_k^t \neq a_k(\hat{\ell}) \\ m^* & \text{otherwise} \end{cases} \quad (\text{B.21})$$

Our next subcase of $\kappa \leq \bar{\kappa}$ is that of $m_i^t = \check{m}^j$. With the understanding that j' is any element of I not equal to i , we let

$$\mu_i^t(\check{m}^j, x^t, a^t, y^t) = \begin{cases} \check{m}^{j'} & \text{if } x^t = x(\dots, j_\ell, \dots), a^t = \check{a}^j(j_\ell) & \text{and } y^t = y(j') \\ \nu(\underline{M}(i, t+1)) & \text{if } x^t = x(\dots, j_\ell, \dots), a^t = \check{a}^j(j_\ell) & \text{and } y^t = y(i) \\ \underline{m}^{k,T} & \text{if } x^t = x(\dots, j_\ell, \dots), a_{-k}^t = \check{a}_{-k}^j(j_\ell) & \text{and } a_k^t \neq \check{a}_k^j(j_\ell) \\ m^* & \text{otherwise} \end{cases} \quad (\text{B.22})$$

⁴⁹Notice that the third case in (B.17) can only possibly apply when $t \geq T+1$.

⁵⁰Throughout the paper we adopt the following notational convention. Given any finite set, we denote by $\nu(\cdot)$ the uniform probability distribution over the set. So, if Z is a finite set, $\nu(Z)$ assigns probability $1/|Z|$ to every element of Z .

Still assuming $\kappa \leq \bar{\kappa}$ we now deal with the subcase $m_i^t \in \underline{M}(i, t)$. For any $\underline{m}^{i, \tau} \in \underline{M}(i, t)$, we let

$$\mu_i^t(\underline{m}^{i, \tau}, x^t, a^t, y^t) = \begin{cases} \check{m}^j & \text{if } x^t = x(\cdots, j_\ell, \cdots), a^t = \check{a}^i(i_\ell) & \text{and } y^t = y(j) \\ \nu(\underline{M}(i, t+1)) & \text{if } x^t = x(\cdots, j_\ell, \cdots), a^t = \check{a}^i(i_\ell) & \text{and } y^t = y(i) \\ \underline{m}^{k, T} & \text{if } x^t = x(\cdots, j_\ell, \cdots), a_{-k}^t = \check{a}_{-k}^i(i_\ell) & \text{and } a_k^t \neq \check{a}_k^i(i_\ell) \\ \underline{m}^{k, T} & \text{if } x^t = x(\cdots, j_\ell, \cdots), a_{-k}^t = a_{-k}^i(i_\ell) & \text{and } a_k^t \neq a_k^i(i_\ell) \\ \underline{m}^{i, \tau-1} & \text{if } x^t = x(\cdots, j_\ell, \cdots) & \text{and } a^t = a(i_\ell) \\ m^* & \text{otherwise} \end{cases} \quad (\text{B.23})$$

where we set $\underline{m}^{i, 0} = \bar{m}^i$. Notice that player $\langle i, t \rangle$ may need to distinguish between the third and fourth cases of (B.23) since clearly they may be generated by different values of the index $k \in I$. To verify that this distinction is always feasible, recall that, by construction (see Definition B.2), $\check{a}_{-i}(i_\ell)$ differs from $a_{-i}(i_\ell)$ in every component, and that of course $n \geq 4$.

The next subcase of $\kappa \leq \bar{\kappa}$ we consider is that of $m_i^t \in \underline{M}(j, t)$. For any $\underline{m}^{j, \tau} \in \underline{M}(j, t)$, we let

$$\mu_i^t(\underline{m}^{j, \tau}, x^t, a^t, y^t) = \begin{cases} \underline{m}^{j, \tau-1} & \text{if } x^t = x(\cdots, j_\ell, \cdots) & \text{and } a^t = a(j_\ell) \\ \underline{m}^{k, T} & \text{if } x^t = x(\cdots, j_\ell, \cdots), a_{-k}^t = a_{-k}(j_\ell) & \text{and } a_k^t \neq a_k(j_\ell) \\ m^* & \text{otherwise} \end{cases} \quad (\text{B.24})$$

where we set $\underline{m}^{j, 0} = \bar{m}^j$.

Finally, still assuming that $\kappa \leq \bar{\kappa}$, we consider the case in which $m_i^t = \bar{m}^{k'}$ for some $k' \in I$. We let

$$\mu_i^t(\bar{m}^{k'}, x^t, a^t, y^t) = \begin{cases} \bar{m}^{k'} & \text{if } x^t = x(\cdots, \ell_{k'}, \cdots) & \text{and } a^t = a(\ell_{k'}) \\ \underline{m}^{k, T} & \text{if } x^t = x(\cdots, \ell_{k'}, \cdots), a_{-k}^t = a_{-k}(\ell_{k'}) & \text{and } a_k^t \neq a_k(\ell_{k'}) \\ m^* & \text{otherwise} \end{cases} \quad (\text{B.25})$$

B.3. Completely Mixed Strategies

Definition B.10: Throughout this section we let ε denote a small positive number, which parameterizes the completely mixed strategies that we construct. It should be understood that our construction of beliefs involves the limit $\varepsilon \rightarrow 0$.

Definition B.11. *Completely Mixed Action Strategies:* Given ε , the completely mixed strategies for all players $\langle i, t \rangle$ at the action stage are denoted by $g_{i, \varepsilon}^t$ and are defined as follows.⁵¹

After receiving a message $m \in \{m^*\} \cup \check{M}_{-i} \cup \underline{M}(i, t)$ and observing the realization x^t of the action-stage randomization device, any player $\langle i, t \rangle$ plays the action prescribed by the action-stage strategy described in Definition B.8 with probability $1 - \varepsilon^2(\|A\|_i - 1)$ and plays all other actions in A_i with probability ε^2 each.

After receiving any message $m \notin \{m^*\} \cup \check{M}_{-i} \cup \underline{M}(i, t)$ and observing the realization x^t of the action-stage randomization device, any player $\langle i, t \rangle$ plays the action prescribed by the action-stage strategy described in Definition B.8 with probability $1 - \varepsilon(\|A\|_i - 1)$ and plays all other actions in A_i with probability ε each.

Definition B.12. *Completely Mixed Message Strategies:* Given ε , the completely mixed strategies for all players $\langle i, t \rangle$ at the message stage are denoted by $\mu_{i, \varepsilon}^t$ and are defined as follows.

Player $\langle i, t \rangle$ sends the message prescribed by the message-stage strategy described in Definition B.9 with probability $1 - \varepsilon^{2n+1}(\|M_i^{t+1}\| - 1)$ and sends all other messages in M_i^{t+1} with probability ε^{2n+1} each.

⁵¹In the interest of brevity, we avoid an explicit distinction between the $t = 0$ players and all others. What follows can be interpreted as applying to all players re-defining m_i^0 to be equal to m^* for players $\langle i \in I, 0 \rangle$.

Appendix C: The Proof of Theorem 4

Definition C.1: Let a profile of message strategies μ be given. Fix an ‘‘augmented history’’ $\kappa^t = (x^0, a^0, y^0, \dots, x^{t-1}, a^{t-1}, y^{t-1})$. In other words, fix a history h^t , together with a sequence of realizations of the message-stage randomization device (y^0, \dots, y^{t-1}) . In what follows, κ^0 will denote the null history \emptyset , and for any $\tau \leq t$, κ^τ will denote the appropriate subset of κ^t .

For every $i \in I$ let $\mathcal{M}_i^0(m_i^0 | \kappa^0, \mu_i) = 1$. Then, recursively forward, define

$$\mathcal{M}_i^t(m_i^t | \kappa^t, \mu_i) = \sum_{m_i^{t-1} \in H^{t-1}} \mu_i^{t-1}(m_i^t | m_i^{t-1}, x^{t-1}, a^{t-1}, y^{t-1}) \mathcal{M}_i^{t-1}(m_i^{t-1} | \kappa^{t-1}, \mu_i) \quad (\text{C.1})$$

So that $\mathcal{M}_i^t(m_i^t | \kappa^t, \mu_i)$ is the probability that player $\langle i, t-1 \rangle$ sends message m_i^t given κ^t and the profile μ_i . We also let $\mathcal{M}_{-i}^t(m_{-i}^t | \kappa^t, \mu_{-i}) = \mathcal{M}_{-i}^t((m_i^t, \dots, m_{i-1}^t, m_{i+1}^t, \dots, m_n^t) | \kappa^t, \mu_{-i}) = \prod_{j \neq i} \mathcal{M}_j^t(m_j^t | \kappa^t, \mu_j)$.

Lemma C.1: Fix any $\delta \in (0, 1)$, any \tilde{x} and any \tilde{y} . Fix any SE of the dynastic repeated game (g, μ, Φ) . Assume that it displays Inter-Generational Agreement as in Definition 4.

Let any augmented history κ^t as in Definition C.1 be given. Let also any $i \in I$ and any m_i^t such that $\mathcal{M}_i^t(m_i^t | \kappa^t, \mu_i) > 0$ be given.

Then for any m_{-i}^t

$$\Phi_i^{tB}(m_{-i}^t | m_i^t) = \mathcal{M}_{-i}^t(m_{-i}^t | \kappa^t, \mu_{-i}) \quad (\text{C.2})$$

Proof: We proceed by induction. Given the fixed κ^t , let $\kappa^0 = \emptyset$ and κ^τ with $\tau = 1, \dots, t$ be the augmented histories comprising the first three components (x^0, a^0, y^0) of κ^t , the first six components $(x^0, a^0, y^0, x^1, a^1, y^1)$ of κ^t and so on. First of all notice that setting $\tau = 0$ yields

$$\Phi_i^{1B}(m_{-i}^0 | m_i^0) = \mathcal{M}_{-i}^1(m_{-i}^0 | \kappa^0, \mu_{-i}) = 1 \quad (\text{C.3})$$

which is trivially true given that all players $\langle i \in I, 0 \rangle$ receive the null message by construction.

Our working hypothesis is now that the claim is true for an arbitrary $\tau - 1 < t - 1$, and our task is to show that it holds for τ .

Consider any message m_i^τ in $\text{Supp}(\mathcal{M}_i^\tau(\cdot | \kappa^\tau, \mu_i))$.⁵² Then there must exist a message $m_i^{\tau-1}$ such that

$$\mu_i^{\tau-1}(m_i^\tau | m_i^{\tau-1}, x^{\tau-1}, a^{\tau-1}, y^{\tau-1}) \mathcal{M}_i^{\tau-1}(m_i^{\tau-1} | \kappa^{\tau-1}, \mu_i) > 0 \quad (\text{C.4})$$

Therefore, using (9) we can write

$$\Phi_i^{\tau B}(m_{-i}^\tau | m_i^\tau) = \Phi_i^{\tau-1 E}(m_{-i}^{\tau-1} | m_i^{\tau-1}, x^{\tau-1}, a^{\tau-1}, y^{\tau-1}) \quad (\text{C.5})$$

Notice that in any SE it must be the case that the right-hand side of (C.5) is equal to

$$\sum_{m_{-i}^{\tau-1}} \Phi_i^{\tau-1 R}(m_{-i}^{\tau-1} | m_i^{\tau-1}, x^{\tau-1}, a^{\tau-1}, y^{\tau-1}) \left[\prod_{j \neq i} \mu_j^{\tau-1}(m_j^{\tau-1} | m_j^{\tau-1}, x^{\tau-1}, a^{\tau-1}, y^{\tau-1}) \right] \quad (\text{C.6})$$

Using (8), we know that (C.6) is equal to

$$\sum_{m_{-i}^{\tau-1}} \Phi_i^{\tau-1 B}(m_{-i}^{\tau-1} | m_i^{\tau-1}) \left[\prod_{j \neq i} \mu_j^{\tau-1}(m_j^{\tau-1} | m_j^{\tau-1}, x^{\tau-1}, a^{\tau-1}, y^{\tau-1}) \right] \quad (\text{C.7})$$

⁵² $\text{Supp}(\cdot)$ denotes the support of a probability distribution.

Our working hypothesis can now be used to assert that (C.7) is in turn equal to

$$\sum_{m_{-i}^{\tau-1}} \mathcal{M}_{-i}^{\tau-1}(m_{-i}^{\tau-1} | \kappa^{\tau-1}, \mu_{-i}) \left[\prod_{j \neq i} \mu_j^{\tau-1}(m_j^{\tau-1} | m_j^{\tau-1}, x^{\tau-1}, a^{\tau-1}, y^{\tau-1}) \right] \quad (\text{C.8})$$

Rearranging terms (C.8) we find that it can also be written as

$$\prod_{j \neq i} \left[\sum_{m_j^{\tau-1}} \mathcal{M}_j^{\tau-1}(m_j^{\tau-1} | \kappa^{\tau-1}, \mu_j) \mu_j^{\tau-1}(m_j^{\tau-1} | m_j^{\tau-1}, x^{\tau-1}, a^{\tau-1}, y^{\tau-1}) \right] \quad (\text{C.9})$$

Using now (C.1), it is immediate that (C.9) is equal to

$$\prod_{j \neq i} \mathcal{M}_j^{\tau}(m_j^{\tau} | \kappa^{\tau}, \mu_j) = \mathcal{M}_{-i}^{\tau}(m_{-i}^{\tau} | \kappa^{\tau}, \mu_{-i}) \quad (\text{C.10})$$

and hence the claim is proved. ■

Definition C.2: Fix any $\delta \in (0, 1)$, any \tilde{x} and any \tilde{y} . Fix any strategy profile, (g, μ) , for the dynastic repeated game.

Consider the standard repeated game with the same common discount factor δ , and with the following action-stage randomization device \hat{x} . The random variable \hat{x} takes values in the finite set $Y \times X$ (the sets in which \tilde{y} and \tilde{x} take values respectively), and the probability that \hat{x} is equal to $\hat{x} = (y, x)$ is $\Pr(\tilde{y} = y) \Pr(\tilde{x} = x)$. For notational convenience we will denote the realization \hat{x}^t of \hat{x}^t by the pair (y^{t-1}, x^t) .

Recall that a history in the standard repeated game with randomization device \hat{x} is an object of the type $h^t = (\hat{x}^0, a^0, \dots, \hat{x}^{t-1}, a^{t-1})$. Therefore, using our notational convention about time superscripts of the realizations of \hat{x}^t we have that any pair (h^t, \hat{x}^t) can be written as a triple (y^{-1}, κ^t, x^t) , where κ^t corresponds to h^t in the obvious way.

We say that the strategy profile g^* for the standard repeated game with randomization device \hat{x} is derived from the dynastic repeated game profile (g, μ) as above if and only if it is defined as follows.

$$g_i^{t*}(h^t, \hat{x}^t) = g_i^{t*}(y^{-1}, \kappa^t, x^t) = \sum_{m_i^t} \mathcal{M}_i^t(m_i^t | \kappa^t, \mu_i) g_i^t(m_i^t, x^t) \quad (\text{C.11})$$

Lemma C.2: Fix any $\delta \in (0, 1)$, any \tilde{x} and any \tilde{y} . Consider any SE, (g, μ) , of the dynastic repeated game that displays Inter-Generational Agreement as in Definition 4.

Now consider the strategy profile g^* for the standard repeated game with randomization device \hat{x} that is derived from (g, μ) as in Definition C.2.

Given g^* , fix any pair (h^t, \hat{x}^t) representing a history and realized randomization device for the standard repeated game. For any $a_{-i}^t \in A_{-i}$, let $\mathcal{P}_{g^* | h^t, \hat{x}^t}(a_{-i}^t)$ be the probability that the realized action profile for all players but i at time t is a_{-i}^t .

Given the pair (h^t, \hat{x}^t) , consider the corresponding triple (y^{-1}, κ^t, x^t) as in Definition C.2. Then

$$\mathcal{P}_{g^* | h^t, \hat{x}^t}(a_{-i}^t) = \prod_{j \neq i} \left\{ \sum_{m_j^t} \mathcal{M}_j^t(m_j^t | \kappa^t, \mu_j) g_j^t(a_j^t | m_j^t, x^t) \right\} \quad (\text{C.12})$$

Proof: The claim is a direct consequence of (C.11) of Definition C.2. ■

Lemma C.3: Fix any $\delta \in (0, 1)$, any \tilde{x} and any \tilde{y} . Consider any SE (g, μ, Φ) of the dynastic repeated game that displays Inter-Generational Agreement as in Definition 4.

Fix any pair (h^t, \hat{x}^t) representing a history and realized randomization device for the standard repeated game. Given the pair (h^t, \hat{x}^t) , consider the corresponding triple (y^{-1}, κ^t, x^t) as in Definition C.2. Given the last two elements of this triple (κ^t, x^t) , now use (C.1) to find a message \bar{m}_i^t such that $\mathcal{M}_i^t(\bar{m}_i^t | \kappa^t, \mu_i) > 0$.

Finally, consider the following alternative action-stage and message-stage strategies $(\bar{g}_i^t, \bar{\mu}_i^t)$ for player $\langle i, t \rangle$. Whenever $m_i^t \neq \bar{m}_i^t$, set $\bar{g}_i^t = g_i^t$ and $\bar{\mu}_i^t = \mu_i^t$. Then define

$$\bar{g}_i^t(\bar{m}_i^t, x^t) = \sum_{m_i^t} \mathcal{M}_i^t(m_i^t | \kappa^t, \mu_i) g_i^t(m_i^t, x^t) \quad (\text{C.13})$$

and

$$\bar{\mu}_i^t(\bar{m}_i^t, x^t, a^t, y^t) = \sum_{m_i^t} \mathcal{M}_i^t(m_i^t | \kappa^t, \mu_i) \mu_i^t(m_i^t, x^t, a^t, y^t) \quad (\text{C.14})$$

Then⁵³

$$v_i^t(g, \mu | \bar{m}_i^t, x^t, \Phi_i^{tB}) = v_i^t(\bar{g}_i^t, g_i^{-t}, g_{-i}, \bar{\mu}_i^t, \mu_i^{-t}, \mu_{-i} | \bar{m}_i^t, x^t, \Phi_i^{tB}) \quad (\text{C.15})$$

Proof: The claim is a direct consequence of Lemma C.1 and of (8) of Definition 4. The details are omitted for the sake of brevity. ■

C.1. Proof of the Theorem

Fix any $\delta \in (0, 1)$, any \tilde{x} and any \tilde{y} . Consider any SE triple (g, μ, Φ) for the dynastic repeated game. Assume that this SE displays Inter-Generational Agreement as in Definition 4.

Now consider the strategy profile g^* for the standard repeated game with common discount δ and randomization device \hat{x} that is derived from (g, μ) as in Definition C.2.

Since (g, μ) and g^* obviously give rise to the same payoff vector, to prove the claim it is enough to show that $g^* \in \mathcal{G}^S(\delta, \hat{x})$. By way of contradiction, suppose that $g^* \notin \mathcal{G}^S(\delta, \hat{x})$.

By Remark 1 (One-Shot Deviation Principle) this implies that there exist an i , an h^t , an \hat{x}^t and a σ_i such that

$$v_i(\sigma_i, g_i^{-t*}, g_{-i}^* | h^t, \hat{x}^t) > v_i(g^* | h^t, \hat{x}^t) \quad (\text{C.16})$$

Given the pair (h^t, \hat{x}^t) , consider the corresponding triple (y^{-1}, κ^t, x^t) as in Definition C.2. Given the last two elements of this triple (κ^t, x^t) , now use (C.1) to find a message \bar{m}_i^t such that $\mathcal{M}_i^t(\bar{m}_i^t | \kappa^t, \mu_i) > 0$.

Using Lemmas C.1 and C.2 we can now conclude that (C.16) implies that

$$v_i^t(\sigma_i, g_i^{-t}, g_{-i}, \bar{\mu}_i^t, \mu_i^{-t}, \mu_{-i} | \bar{m}_i^t, x^t, \Phi_i^{tB}) > v_i^t(\bar{g}_i^t, g_i^{-t}, g_{-i}, \bar{\mu}_i^t, \mu_i^{-t}, \mu_{-i} | \bar{m}_i^t, x^t, \Phi_i^{tB}) \quad (\text{C.17})$$

where σ_i is the profitable deviation identified in (C.16) and \bar{g}_i^t and $\bar{\mu}_i^t$ are the alternative action-stage and message-stage strategies of Lemma C.3.

⁵³Here and throughout the rest of the paper and the technical addendum, we denote by $v_i^t(g, \mu | m_i^t, x^t, \Phi_i^{tB})$ the continuation payoff to player $\langle i, t \rangle$ given the profile (g, μ) , after he has received message m_i^t has observed the realization x^t , and given that his beliefs over the $n - 1$ -tuple m_{-i}^t are Φ_i^{tB} . See also our Point of Notation T.1.1.

However, using (C.15) of Lemma C.3, the inequality in (C.17) clearly implies that

$$v_i^t(\sigma_i, g_i^{-t}, g_{-i}, \bar{\mu}_i^t, \mu_i^{-t}, \mu_{-i} | \bar{m}_i^t, x^t, \Phi_i^{tB}) > v_i^t(g, \mu | \bar{m}_i^t, x^t, \Phi_i^{tB}) \quad (\text{C.18})$$

But since (C.18) contradicts the fact that (g, μ, Φ) is an SE of the dynastic repeated game, the proof is now complete. ■

References

- ANDERLINI, L., AND R. LAGUNOFF (2005): “Communication in Dynastic Repeated Games: ‘Whitewashes’ and ‘Coverups,’” *Economic Theory*, 26, 265–299.
- AUMANN, R. J. (1981): “Survey of Repeated Games,” in *Essays in Game Theory and Mathematical Economics, in Honor of Oskar Morgenstern*, pp. 11–42. Mannheim: Wissenschaftsverlag Bibliographisches Institut.
- BALIGA, S., L. CORCHON, AND T. SJOSTROM (1997): “The Theory of Implementation When the Planner is a Player,” *Journal of Economic Theory*, 77, 15–33.
- BEN-PORATH, E., AND M. KAHNEMAN (1996): “Communication in Repeated Games with Private Monitoring,” *Journal of Economic Theory*, 70, 281–297.
- BHASKAR, V. (1998): “Informational Constraints and the Overlapping Generations Model: Folk and Anti-Folk Theorems,” *Review of Economic Studies*, 65, 135–149.
- CHAUDHURI, A., A. SCHOTTER, AND B. SOPHER (2001): “Talking Ourselves to Efficiency: Coordination in Inter-Generational Minimum Games with Private, Almost Common and Common Knowledge of Advice,” New York University, mimeo.
- COMPTE, O. (1998): “Communication in Repeated Games with Imperfect Private Monitoring,” *Econometrica*, 66, 597–626.
- CORBAE, D., T. TEMZELIDES, AND R. WRIGHT (2001): “Endogenous Matching and Money,” Univeristy of Pennsylvania, mimeo.
- CREMER, J. (1986): “Cooperation in Ongoing Organizations,” *Quarterly Journal of Economics*, 101, 33–49.
- FARRELL, J. (1993): “Meaning and Credibility in Cheap Talk Games,” *Games and Economic Behavior*, 5, 514–531.
- FUDENBERG, D., AND E. S. MASKIN (1986): “The Folk Theorem in Repeated Games with Discounting or with Incomplete Information,” *Econometrica*, 54, 533–556.
- GAMES AND ECONOMIC BEHAVIOR (1997): *Special Issue on Imperfect Recall*, vol. 20, no. 1. New-York: Academic Press.
- JOHNSON, P., D. K. LEVINE, AND W. PESENDORFER (2001): “Evolution and Information in a Gift Giving Game,” *Journal of Economic Theory*, 100, 1–22.
- KANDORI, M. (1992a): “Repeated Games Played by Overlapping Generations of Players,” *Review of Economic Studies*, 59, 81–92.
- (1992b): “Social Norms and Community Enforcement,” *Review of Economic Studies*, 59, 63–80.
- KANDORI, M., AND H. MATSUSHIMA (1998): “Private Observation, Communication and Collusion,” *Econometrica*, 66, 627–652.

- KOBAYASHI, H. (2003): “Folk Theorem for Infinitely Repeated Games Played by Organizations with Short-Lived Members,” Osaka Prefecture University, mimeo.
- KOCHERLAKOTA, N. (1998): “Money is Memory,” *Journal of Economic Theory*, 81, 232–251.
- KOCHERLAKOTA, N., AND N. WALLACE (1998): “Incomplete Record Keeping and Optimal Payout Arrangements,” *Journal of Economic Theory*, 81, 272–289.
- KREPS, D. M., AND R. WILSON (1982): “Sequential Equilibria,” *Econometrica*, 50, 863–894.
- LAGUNOFF, R., AND A. MATSUI (2004): “Organizations and Overlapping Generations Games: Memory, Communication, and Altruism,” *Review of Economic Design*, 8, 383–411.
- MASKIN, E. S. (1999): “Nash Equilibrium and Welfare Optimality,” *Review of Economic Studies*, 66, 23–38.
- MATTEHWS, S., M. OKUNO-FUJIWARA, AND A. POSTLEWAITE (1991): “Refining Cheap-Talk Equilibria,” *Journal of Economic Theory*, 55, 247–273.
- PEARCE, D. (1992): “Repeated Games: Cooperation and Rationality,” in *Advances in Economic Theory: Sixth World Congress*, ed. by J.-J. Laffont. Cambridge: Cambridge University Press.
- PICCIONE, M., AND A. RUBINSTEIN (1997): “On the Interpretation of Decision Problems with Imperfect Recall,” *Games and Economic Behavior*, 20, 3–24.
- SALANT, D. (1991): “A Repeated Game with Finitely Lived Overlapping Generations of Players,” *Games and Economic Behavior*, 3, 244–259.
- SCHOTTER, A., AND B. SOPHER (2001a): “Advice and Behavior in Intergenerational Ultimatum Games: An Experimental Approach,” New York University, mimeo.
- (2001b): “Social Learning and Coordination Conventions in Inter-Generational Games: An Experimental Study,” New York University, mimeo.
- SMITH, L. (1992): “Folk Theorems in Overlapping Generations Games,” *Games and Economic Behavior*, 4, 426–449.
- WALLACE, N. (2001): “Whither Monetary Economics?,” *International Economic Review*, 42, 847–869.

A “Super” Folk Theorem for Dynastic Repeated Games: Technical Addendum

T.1. Notation

Point of Notation T.1.1: Abusing the notation we established for the standard repeated game, we adopt the following notation for continuation payoffs in the dynastic repeated game. Let an assessment (g, μ, Φ) be given.

Recall that we denote by $v_i^t(g, \mu | m_i^t, x^t, \Phi_i^{tB})$ the continuation payoff to player $\langle i, t \rangle$ given the profile (g, μ) , after he has received message m_i^t , has observed the realization x^t , and given that his beliefs over the $n - 1$ -tuple m_{-i}^t are Φ_i^{tB} (see footnote 53). In view of our discussion at the beginning of Section 4, it is clear that the only component of the system of beliefs Φ that is relevant to define this continuation payoff is in fact Φ_i^{tB} . Our discussion there also implies that the argument m_i^t is redundant once Φ_i^{tB} has been specified. We keep it in our notation since it helps streamline some of the arguments below.

We let $v_i^t(g, \mu | m_i^t, x^t, a^t, y^t, \Phi_i^{tE})$ denote the continuation payoff (viewed from the beginning of period $t + 1$) to player $\langle i, t \rangle$ given the profile (g, μ) , after he has received message m_i^t , has observed the triple (x^t, a^t, y^t) , and given that his beliefs over the $n - 1$ -tuple m_{-i}^{t+1} are given by Φ_i^{tE} . In view of our discussion at the beginning of Section 4, it is clear that once Φ_i^{tE} has been specified, the arguments (m_i^t, x^t, a^t, y^t) are redundant in determining the end-of-period continuation payoff to player $\langle i, t \rangle$. Whenever this does not cause any ambiguity (about Φ_i^{tE}) we will write $v_i^t(g, \mu | \Phi_i^{tE})$ instead of $v_i^t(g, \mu | m_i^t, x^t, a^t, y^t, \Phi_i^{tE})$.

As we noted in the text all continuation payoffs clearly depend on δ as well. To keep notation down this dependence will be omitted whenever possible.

Point of Notation T.1.2: We will abuse our notation for $\Phi_i^{tB}(\cdot)$, $\Phi_i^{tE}(\cdot)$ and $\Phi_i^{tR}(\cdot)$ slightly in the following way. We will allow events of interest and conditioning events to appear as arguments of Φ_i^{tB} , Φ_i^{tE} and Φ_i^{tR} , to indicate their probabilities under these distributions.

So, for instance when we write $\Phi_i^{tB}(m_{-i}^t = (z, \dots, z) | m_i^t) = c$ we mean that according to the beginning-of-period beliefs of player $\langle i, t \rangle$, after observing m_i^t , the probability that m_{-i}^t is equal to the $n - 1$ -tuple (z, \dots, z) is equal to c .

Point of Notation T.1.3: Whenever the profile (g, μ) is a profile of completely mixed strategies, the beliefs $\Phi_i^{tB}(\cdot)$, $\Phi_i^{tE}(\cdot)$ and $\Phi_i^{tR}(\cdot)$ are of course entirely determined by what player $\langle i, t \rangle$ observes and by (g, μ) using Bayes' rule. In this case, we will allow the pair (g, μ) to appear as a “conditioning event.”

So, for instance, $\Phi_i^{tB}(m_{-i}^t | m_i^t, g, \mu)$ is the probability of the $n - 1$ -tuple m_{-i}^t , after m_i^t has been received, obtained from the completely mixed profile (g, μ) via Bayes' rule. Events may appear as arguments in this case as well, consistently with our Point of Notation T.1.2 above.

Moreover, since the completely mixed pair (g, μ) determines the probabilities of all events, concerning for instance histories, messages of previous cohorts and the like, we will use the notation \Pr to indicate such probabilities, using the pair (g, μ) as a conditioning event.

So, given any two events L and J , the notation $\Pr(L | J, g, \mu)$ will indicate the probability of event L , conditional on event J , as determined by the completely mixed pair (g, μ) via Bayes' rule.

T.2. A Preliminary Result

Definition T.2.1: Consider the dynastic repeated game described in full in Section 3. Now consider the dynastic repeated game obtained from this when we restrict the message space of player $\langle i, t \rangle$ to be $M_i^{t+1} \subseteq H^{t+1}$, with all other details unchanged.

We call this the restricted dynastic repeated game with message spaces $\{M_i^t\}_{i \in I, t \geq 1}$. For any given $\delta \in (0, 1)$, \tilde{x} and \tilde{y} , we denote by $\mathcal{G}^D(\delta, \tilde{x}, \tilde{y}, \{M_i^t\}_{i \in I, t \geq 1})$ the set of SE strategy profiles, while we write $\mathcal{E}^D(\delta, \tilde{x}, \tilde{y}, \{M_i^t\}_{i \in I, t \geq 1})$ for the set of SE payoff profiles of this dynastic repeated game with restricted message spaces.

Lemma T.2.1: *Let any $\delta \in (0, 1)$, \tilde{x} and \tilde{y} be given. Consider now any restricted dynastic repeated game with message spaces $\{M_i^t\}_{i \in I, t \geq 1}$. Then $\mathcal{E}^D(\delta, \tilde{x}, \tilde{y}, \{M_i^t\}_{i \in I, t \geq 1}) \subseteq \mathcal{E}^D(\delta, \tilde{x}, \tilde{y})$.*

Proof: Let a profile $(g^*, \mu^*) \in \mathcal{G}^D(\delta, \tilde{x}, \tilde{y}, \{M_i^t\}_{i \in I, t \geq 1})$ with associated beliefs Φ^* be given. To prove the statement, we proceed to construct a new profile $(g^{**}, \mu^{**}) \in \mathcal{G}^D(\delta, \tilde{x}, \tilde{y})$ and associated beliefs Φ^{**} that are consistent with (g^{**}, μ^{**}) , and which gives every player the same payoff as (g^*, μ^*) .

Denote a generic element of M_i^t by z_i^t . Since $M_i^t \subseteq H^t$, we can partition H^t into $\|M_i^t\|$ non-empty mutually exclusive exhaustive subsets, and make each of these subsets correspond to an element z_i^t of M_i^t . In other words, we can find a map $\rho_i^t: M_i^t \rightarrow 2^{H^t}$ such that $\rho_i^t(z_i^t) \neq \emptyset$ for all $z_i^t \in M_i^t$, $\rho_i^t(z_i^t) \cap \rho_i^t(z_i^{t'}) = \emptyset$ whenever $z_i^t \neq z_i^{t'}$, and $\bigcup_{z_i^t \in M_i^t} \rho_i^t(z_i^t) = H^t$.

We can now describe how the profile (g^{**}, μ^{**}) is obtained from the given (g^*, μ^*) . We deal first with the action stage. For any player $\langle i, t \rangle$, and any $z_i^t \in M_i^t$, set

$$g_i^{t**}(m_i^t, x) = g_i^{t*}(z_i^t, x) \quad \forall m_i^t \in \rho_i^t(z_i^t) \quad (\text{T.2.1})$$

At the message stage, for any player $\langle i, t \rangle$, any (z_i^t, x^t, a^t, y^t) , any $m_i^t \in \rho_i^t(z_i^t)$, and any $z_i^{t+1} \in \text{Supp}(\mu_i^{t*}(z_i^t, x^t, a^t, y^t))$, set

$$\mu_i^{t**}(m_i^{t+1} | m_i^t, x^t, a^t, y^t) = \frac{1}{\|\rho_i^{t+1}(z_i^{t+1})\|} \mu_i^{t*}(z_i^{t+1} | z_i^t, x^t, a^t, y^t) \quad \forall m_i^{t+1} \in \rho_i^{t+1}(z_i^{t+1}) \quad (\text{T.2.2})$$

Next, we describe Φ^{**} , starting with the beginning-of-period beliefs. For any player $\langle i, t \rangle$, any $z_i^t \in M_i^t$ and any $z_{-i}^t \in M_{-i}^t$, set

$$\Phi_i^{tB**}(m_{-i}^t | m_i^t) = \frac{\Phi_i^{tB*}(z_{-i}^t | z_i^t)}{\prod_{j \neq i} \|\rho_j^t(z_j^t)\|} \quad \forall m_i^t \in \rho_i^t(z_i^t), \quad \forall m_{-i}^t \in \prod_{j \neq i} \rho_j^t(z_j^t) \quad (\text{T.2.3})$$

Similarly, concerning the end-of-period beliefs, for any player $\langle i, t \rangle$, any (z_i^t, x^t, a^t, y^t) and any $z_{-i}^{t+1} \in M_{-i}^{t+1}$, set

$$\begin{aligned} \Phi_i^{tE**}(m_{-i}^{t+1} | m_i^t, x^t, a^t, y^t) = \\ \frac{\Phi_i^{tE*}(z_{-i}^{t+1} | z_i^t, x^t, a^t, y^t)}{\prod_{j \neq i} \|\rho_j^{t+1}(z_j^{t+1})\|} \quad \forall m_i^t \in \rho_i^t(z_i^t), \quad \forall m_{-i}^{t+1} \in \prod_{j \neq i} \rho_j^{t+1}(z_j^{t+1}) \end{aligned} \quad (\text{T.2.4})$$

Since the profile (g^*, μ^*) is sequentially rational given Φ^* , it is immediate from (T.2.1), (T.2.2), (T.2.3) and (T.2.4) that the profile (g^{**}, μ^{**}) is sequentially rational given Φ^{**} , and we omit further details of the proof of this claim.

Of course, it remains to show that $(g^{**}, \mu^{**}, \Phi^{**})$ is a consistent assessment.

Let $(g_\varepsilon^*, \mu_\varepsilon^*)$ be parameterized completely mixed strategies which converge to (g^*, μ^*) and give rise, in the limit as $\varepsilon \rightarrow 0$, to beliefs Φ^* via Bayes’ rule.

Given any $\varepsilon > 0$, let $(g_\varepsilon^{**}, \mu_\varepsilon^{**})$ be a profile of completely mixed strategies obtained from $(g_\varepsilon^*, \mu_\varepsilon^*)$ exactly as in (T.2.1) and (T.2.2).

We start by verifying the consistency of the beginning-of-period beliefs. Observe that for any given $z^t = (z_i^t, z_{-i}^t)$, from (T.2.2) we know that whenever $m^t = (m_i^t, m_{-i}^t) \in \prod_{j \in I} \rho_j^t(z_j^t)$

$$\Pr(m_i^t, m_{-i}^t | g_\varepsilon^{**}, \mu_\varepsilon^{**}) = \frac{\Pr(z_i^t, z_{-i}^t | g_\varepsilon^*, \mu_\varepsilon^*)}{\prod_{j \in I} \|\rho_j^t(z_j^t)\|} \quad (\text{T.2.5})$$

Similarly, using (T.2.2) again we know that whenever $m_i^t \in \rho_i^t(z_i^t)$

$$\Pr(m_i^t | g_\varepsilon^{**}, \mu_\varepsilon^{**}) = \frac{\Pr(z_i^t | g_\varepsilon^*, \mu_\varepsilon^*)}{\|\rho_i^t(z_i^t)\|} \quad (\text{T.2.6})$$

Taking the ratio of (T.2.5) and (T.2.6) and taking the limit as $\varepsilon \rightarrow 0$ now yields that for any any $z_i^t \in M_i^t$ and any $z_{-i}^t \in M_{-i}^t$

$$\lim_{\varepsilon \rightarrow 0} \Phi_i^{tB**}(m_{-i}^t | m_i^t, g_\varepsilon^{**}, \mu_\varepsilon^{**}) = \frac{\Phi_i^{tB*}(z_{-i}^t | z_i^t)}{\prod_{j \neq i} \|\rho_j^t(z_j^t)\|} \quad \forall m_i^t \in \rho_i^t(z_i^t), \quad \forall m_{-i}^t \in \prod_{j \neq i} \rho_j^t(z_j^t) \quad (\text{T.2.7})$$

Hence we have shown that the beginning-of-period beliefs as in (T.2.3) are consistent with (g^{**}, μ^{**}) .

The proof that the end-of-period beliefs as in (T.2.4) are consistent with (g^{**}, μ^{**}) runs along exactly the same lines, and we omit the details. ■

T.3. Proof of Theorem 1: Preliminaries

Lemma T.3.1: *Let (g, μ) be a strategy profile in the dynastic repeated game and assume μ is truthful according to Definition 2.*

Then there exists a system of beliefs Φ that is consistent with (g, μ) and such that for every $i \in I, t \geq 0, m_i^t \in H^t, x^t \in X, a^t \in A$ and $y^t \in Y$ we have that

$$\Phi_i^{tB}[m_{-i}^t = (m_i^t, \dots, m_i^t) | m_i^t] = 1 \quad (\text{T.3.1})$$

and

$$\Phi_i^{tE}[m_{-i}^{t+1} = ((m_i^t, x^t, a^t), \dots, (m_i^t, x^t, a^t)) | (m_i^t, x^t, a^t, y^t)] = 1 \quad (\text{T.3.2})$$

In other words, Φ is such that every player $\langle i, t \rangle$ at the beginning of the period believes with probability one that all other players in his cohort have received the same message as he has, and at the end of the period believes that all other players in his cohort are sending the same (truthful and hence pure) message as he is.

Proof: We construct a sequence of completely mixed strategies in which deviations at the action stage are much more likely than deviations at the message stage. We parameterize the sequence of perturbations by a small positive number ε , which will eventually be shrunk to zero.

Given ε , the completely mixed strategy for player $\langle i, t \rangle$ at the action stage is denoted by $g_{i,\varepsilon}^t$. Recall that $g_i^t(m_i^t, x^t)$ is itself a mixed strategy in $\Delta(A_i)$. Then $g_{i,\varepsilon}^t$ is given by^{T.1}

$$g_{i,\varepsilon}^t(m_i^t, x^t) = (1 - \varepsilon^{\frac{1}{(n+1)2^{t+1}}}) g_i^t(m_i^t, x^t) + \varepsilon^{\frac{1}{(n+1)2^{t+1}}} \nu(A_i) \quad (\text{T.3.3})$$

Given ε , the completely mixed strategy for player $\langle i, t \rangle$ at the message stage is denoted by $\mu_{i,\varepsilon}^t$. Recall that $\mu_i^t(m_i^t, x^t, a^t, y^t)$ is itself a mixed strategy in $\Delta(H^{t+1})$. Then $\mu_{i,\varepsilon}^t$ is given by

$$\mu_{i,\varepsilon}^t(m_i^t, x^t, a^t, y^t) = (1 - \varepsilon) \mu_i^t(m_i^t, x^t, a^t, y^t) + \varepsilon \nu(H^{t+1}) \quad (\text{T.3.4})$$

In words, at the action stage, player $\langle i, t \rangle$ deviates from g_i^t with probability $\varepsilon^{\frac{1}{(n+1)2^{t+1}}}$ and all deviations are equally likely. At the message stage, player $\langle i, t \rangle$ deviates from μ_i^t with probability ε , again with all

^{T.1} $\nu(\cdot)$ stands for the the uniform distribution. See also footnote 50.

deviations equally likely. Denote by $(g_\varepsilon, \mu_\varepsilon)$ the profile of completely mixed strategies we have just described. Clearly, as $\varepsilon \rightarrow 0$ the profile $(g_\varepsilon, \mu_\varepsilon)$ converges pointwise to (g, μ) as required.

Of course, to prove (T.3.1) it is enough to show that

$$\lim_{\varepsilon \rightarrow 0} \Phi_i^{tB}[m_{-i}^t = (m_i^t, \dots, m_i^t) | m_i^t, g_\varepsilon, \mu_\varepsilon] = 1 \quad (\text{T.3.5})$$

Notice now that (T.3.5) follows almost immediately from the way we have defined the completely mixed profile $(g_\varepsilon, \mu_\varepsilon)$ in (T.3.3) and (T.3.4) above.

To see this, notice if $m_j^t \neq m_i^t$ for some $j \neq i$ then it must be that at least one player $\langle k \in I, \tau \rangle$ with $\tau = 0, 1, \dots, t-1$ has “lied” his successor in the same dynasty. Given (T.3.4) this happens with a probability that is an infinitesimal in ε of order 1 or higher.^{T.2} This needs to be compared with the overall probability of observing m_i^t . Clearly many paths of play could have generated this outcome. However, *one* way in which m_i^t can arise is certainly that the *true* history h^t is equal to m_i^t and that no player has ever deviated from truth-telling. In the worst case the true history being equal to m_i^t will involve *all* players $\langle j \in I, \tau \rangle$ deviating from g_j^τ in every $\tau \leq t-1$. Using (T.3.3) this is an infinitesimal in ε of order at most $n \sum_{\tau=0}^{t-1} 1/(n+1)2^{\tau+1} < n/(n+1)$. Hence, the overall probability of observing m_i^t cannot be lower than an infinitesimal of order $n/(n+1)$. Since $1 > n/(n+1)$, equation (T.3.5) now follows.^{T.3}

The proof of (T.3.2) is the exact analogue of the proof of (T.3.1) we have just given, and hence we omit the details. ■

T.4. Proof of Theorem 1

Fix a $\delta \in (0, 1)$, a \tilde{x} and any profile $g^* \in \mathcal{G}^S(\delta, \tilde{x})$. Then there exists a profile μ^* of message strategies which are truthful in the sense of Definition 2 and such that $(g^*, \mu^*) \in \mathcal{G}^D(\delta, \tilde{x}, \tilde{y})$ for every finite random variable \tilde{y} .

Since the strategies μ^* are truthful, we know from Lemma T.3.1 that there is a system of beliefs Φ^* that is consistent with (g^*, μ^*) for which (T.3.1) and (T.3.2) hold. We will show that beliefs Φ^* support the strategies (g^*, μ^*) as an SE of the dynastic repeated game regardless of \tilde{y} . Indeed, notice that since μ_i^{t*} is truthful for every player $\langle i, t \rangle$, we know from Definition 2 that all players in fact *ignore* the realization of the message-stage randomization device. Hence our argument below is trivially valid for any \tilde{y} .

We simply check that no player $\langle i, t \rangle$ has an incentive to deviate either at the action or at the communication stage.

Consider player $\langle i, t \rangle$ at the communication stage, after observing (m_i^t, x^t, a^t, y^t) . If he sends message $m_i^{t+1} = (m_i^t, x^t, a^t)$ as prescribed by μ_i^{t*} , given that his beliefs are Φ_i^{tE*} his expected continuation payoff is $E_{\tilde{x}^{t+1}} v_i(g^* | m_i^t, x^t, a^t, \tilde{x}^{t+1})$. Notice that by construction we know that this continuation payoff is equal to $E_{\tilde{x}^{t+1}} v_i(g^* | h^{t+1}, \tilde{x}^{t+1})$ when we set $h^{t+1} = (m_i^t, x^t, a^t)$, namely the expected continuation payoff to player i given g^* in the standard repeated game after history $h^{t+1} = (m_i^t, x^t, a^t)$ has taken place, and before the realization of \tilde{x}^{t+1} has been observed.

We now need to check that player $\langle i, t \rangle$ cannot gain by deviating and sending any other (mixed) message $\phi_i^t \in \Delta(H^{t+1})$. Given that the strategies μ^* are truthful and that his beliefs are Φ_i^{tE*} , his expected continuation payoff following such deviation clearly cannot be above $\max_{g_i} E_{\tilde{x}^{t+1}} v_i(g_i, g_{-i}^* | h^{t+1}, \tilde{x}^{t+1})$ when we set $h^{t+1} = (m_i^t, x^t, a^t)$. In other words it cannot exceed the maximum (by choice of g_i) expected continuation payoff that player i can achieve in the standard repeated game after history $h^{t+1} = (m_i^t, x^t, a^t)$ given that all

^{T.2}Throughout, we use the words “infinitesimal in ε of order z ” to indicate any quantity that can be written as a constant times ε^z . Similarly, we use the words “infinitesimal of order higher than z ” to mean any quantity that can be written as a constant times $\varepsilon^{z'}$, with $z' > z$.

^{T.3}It is instructive to notice essentially the same argument we are following here is enough to show that in fact, upon receiving m_i^t player $\langle i, t \rangle$ will assign probability one to the event that the true history of play is equal to m_i^t . Formally this would be expressed as $\lim_{\varepsilon \rightarrow 0} \Pr(h^t = m_i^t | m_i^t, g_\varepsilon, \mu_\varepsilon) = 1$.

other players are playing according to g_{-i}^* . However, since $g^* \in \mathcal{G}^S(\delta, \tilde{x})$ we know that $E_{\tilde{x}^{t+1}} v_i(g^* | h^{t+1}, \tilde{x}^{t+1}) = \max_{g_i} E_{\tilde{x}^{t+1}} v_i(g_i, g_{-i}^* | h^{t+1}, \tilde{x}^{t+1})$. Hence, we can conclude that no player $\langle i, t \rangle$ cannot gain by deviating in this way.

Now consider player $\langle i, t \rangle$ at the action stage, after observing (m_i^t, x^t) . If player $\langle i, t \rangle$ follows the prescription of g_i^{t*} given that his beliefs are Φ_i^{tE*} his expected continuation payoff is given by $v_i(g^* | m_i^t, x^t)$. If he deviates to playing any other $\sigma_i \in \Delta(A_i)$, given his beliefs, his expected continuation payoff is $v_i(\sigma_i, g_{-i}^*, g_{-i}^* | m_i^t, x^t)$. Since $g^* \in \mathcal{G}^S(\delta, \tilde{x})$, by (4) of Remark 1 we can then conclude that he cannot gain by deviating in this way. ■

T.5. Proof of Theorem 2: Beliefs

Definition T.5.1: Using the notation of Definition A.1, let

$$\hat{v} = \frac{1}{1-q} \sum_{\ell=2}^{\|A\|} p(a(\ell)) u(a(\ell)) \quad (\text{T.5.1})$$

and

$$v' = qu(a') + (1-q)\hat{v} \quad (\text{T.5.2})$$

Remark T.5.1: Let \hat{v} and v' be as in Lemma T.5.1, then

$$v^* = qu(a^*) + (1-q)\hat{v} \quad (\text{T.5.3})$$

and

$$v^* - v' = q[u(a^*) - u(a')] \quad (\text{T.5.4})$$

so that using our assumptions about a^* and a'

$$v_i^* - v_i' > 0 \quad \forall i \in I \quad (\text{T.5.5})$$

Remark T.5.2: Given that $\alpha \in (0, 1)$ is such that (A.3) of Definition A.3 holds, then simple algebra shows that the interval

$$\left(\frac{\bar{u}_i - \underline{u}_i}{(1-\alpha)(v_i^* - v_i')}, \frac{u_i(a') - u_i(a_i^*, a'_{-i})}{\alpha(v_i^* - v_i')} \right) \quad (\text{T.5.6})$$

is not empty for every $i \in I$.

Definition T.5.2: Using Remark T.5.2, for each $i \in I$, define r_i to be a number in the interval in (T.5.6). Moreover, for each $i \in I$ define $\beta_i(\delta) = (1-\delta)r_i$. Notice that, given r_i , as δ grows towards one, clearly $\beta_i(\delta) \in (0, 1)$.

Definition T.5.3. *Beginning-of-Period Beliefs:* The beginning-of-period beliefs of all players $\langle i \in I, 0 \rangle$ are trivial. Of course, all players believe that all other players have received the null message m_i^0 .

The beginning-of-period beliefs $\Phi_i^{tB}(m_i^t)$ of any other player $\langle i, t \rangle$, depending on the message he receives from player $\langle i, t-1 \rangle$ are as follows

$$\begin{aligned}
& \text{if } m_i^t = m^* \quad \text{then } m_{-i}^t = (m^*, \dots, m^*) \text{ with probability } 1 \\
& \text{if } m_i^t = m^A \quad \text{then } m_{-i}^t = (m^A, \dots, m^A) \text{ with probability } 1 \\
& \text{if } m_i^t = m^B \quad \text{then } \begin{cases} m_{-i}^t = (m^*, \dots, m^*) \text{ with probability } \beta_i(\delta) \\ m_{-i}^t = (m^B, \dots, m^B) \text{ with probability } 1 - \beta_i(\delta) \end{cases}
\end{aligned} \tag{T.5.7}$$

Definition T.5.4. *End-of-Period Beliefs:* For ease of notation, we divide our description of the end-of-period beliefs of player $\langle i, t \rangle$ into two cases: $x^t = x(1)$, and $x^t = x(\ell)$ with $\ell \geq 2$.^{T.4}

For any player $\langle i, t \rangle$, whenever $x^t = x(1)$, let $\Phi_i^{tE}(m_i^t, x(1), a^t, y^t)$ be as follows^{T.5}

$$\begin{aligned}
& \text{if } a^t = (a'_j, a_{-j}^*) \text{ and } y^t = y(0) \quad \text{then } m_{-i}^{t+1} = (m^A, \dots, m^A) \text{ with probability } 1 \\
& \text{if } m_i^t = m^A, a^t = (a_j^*, a'_{-j}) \text{ and } y^t = y(0) \quad \text{then } m_{-i}^{t+1} = (m^A, \dots, m^A) \text{ with probability } 1 \\
& \text{if } m_i^t = m^A, a^t = a' \text{ and } y^t = y(0) \quad \text{then } m_{-i}^{t+1} = (m^A, \dots, m^A) \text{ with probability } 1 \\
& \text{if } m_i^t = m^B \text{ and } a^t = a' \quad \text{then } m_{-i}^{t+1} = (m^B, \dots, m^B) \text{ with probability } 1 \\
& \text{if } m_i^t = m^B \text{ and } a^t = (a_j^*, a'_{-j}) \quad \text{then } m_{-i}^{t+1} = (m^B, \dots, m^B) \text{ with probability } 1 \\
& \text{if } a_j^t \notin \{a_j^*, a'_j\} \text{ for some } j \in I \quad \text{then } m_{-i}^{t+1} = (m^B, \dots, m^B) \text{ with probability } 1 \\
& \text{in all other cases} \quad \text{then } m_{-i}^{t+1} = (m^*, \dots, m^*) \text{ with probability } 1
\end{aligned} \tag{T.5.8}$$

For any player $\langle i, t \rangle$, whenever $x^t = x(\ell)$ with $\ell \geq 2$ let $\Phi_i^{tE}(m_i^t, x(\ell), a^t, y^t)$ be as follows

$$\begin{aligned}
& \text{if } a^t = x(\ell), m_i^t = m^A \text{ and } y^t = y(0) \quad \text{then } m_{-i}^{t+1} = (m^A, \dots, m^A) \text{ with probability } 1 \\
& \text{if } a^t \neq x(\ell) \quad \text{then } m_{-i}^{t+1} = (m^B, \dots, m^B) \text{ with probability } 1 \\
& \text{if } a^t = x(\ell) \text{ and } m_i^t = m^* \quad \text{then } m_{-i}^{t+1} = (m^*, \dots, m^*) \text{ with probability } 1 \\
& \text{if } a^t = x(\ell), m_i^t = m^A \text{ and } y^t = y(1) \quad \text{then } m_{-i}^{t+1} = (m^*, \dots, m^*) \text{ with probability } 1 \\
& \text{if } a^t = x(\ell) \text{ and } m_i^{t-1} = m^B \quad \text{then } \begin{cases} m_{-i}^{t+1} = (m^*, \dots, m^*) \text{ prob. } \beta_i(\delta) \\ m_{-i}^{t+1} = (m^B, \dots, m^B) \text{ prob. } 1 - \beta_i(\delta) \end{cases}
\end{aligned} \tag{T.5.9}$$

T.6. Proof of Theorem 2: Sequential Rationality

We begin by checking the sequential rationality of the message strategies we have defined.

^{T.4}In the interest of brevity, we avoid writing down the end-of-period beliefs for players $\langle i \in I, 0 \rangle$ separately. Equations (T.5.8) and (T.5.9) that follow can be interpreted as defining the end-of-period beliefs of the time 0 players by re-defining m_i^0 to be equal to m^* for all players $\langle i \in I, 0 \rangle$.

^{T.5}Using a short-hand version of notation established in Definition A.6, when, for instance, we write $a^t = (a_j^*, a'_{-j})$, we mean that this is the case for *some* $j \in I$.

Definition T.6.1: Let \mathcal{I}_i^{tE} denote the end-of-period- t collection of information sets that belong to player $\langle i, t \rangle$, with typical element \mathcal{I}_i^{tE} .

It is convenient to partition \mathcal{I}_i^{tE} into four mutually disjoint exhaustive subsets on the basis of the associated beliefs of player $\langle i, t \rangle$.

Let $\mathcal{I}_i^{tE}(\ast) \subset \mathcal{I}_i^{tE}$ be the collection of information sets in which player $\langle i, t \rangle$ believes that m_{-i}^{t+1} is equal to (m^\ast, \dots, m^\ast) with probability one. These beliefs will be denoted by $\Phi_i^{tE}(\ast)$. Notice that these information sets are those in the last case of (T.5.8) and the third and fourth case of (T.5.9).

Let $\mathcal{I}_i^{tE}(A) \subset \mathcal{I}_i^{tE}$ be the collection of information sets in which player $\langle i, t \rangle$ believes that m_{-i}^{t+1} is equal to (m^A, \dots, m^A) with probability one. These beliefs will be denoted by $\Phi_i^{tE}(A)$. Notice that these information sets are those in the first three cases of (T.5.8) and the first case of (T.5.9).

Let $\mathcal{I}_i^{tE}(B) \subset \mathcal{I}_i^{tE}$ be the collection of information sets in which player $\langle i, t \rangle$ believes that m_{-i}^{t+1} is equal to (m^B, \dots, m^B) with probability one. These beliefs will be denoted by $\Phi_i^{tE}(B)$. Notice that these information sets are those in fourth, fifth and sixth cases of (T.5.8), and the second case of (T.5.9).

Finally, let $\mathcal{I}_i^{tE}(\ast B) \subset \mathcal{I}_i^{tE}$ be the collection of information sets in which player $\langle i, t \rangle$ believes that m_{-i}^{t+1} is equal to (m^\ast, \dots, m^\ast) with probability $\beta_i(\delta)$ and to (m^B, \dots, m^B) with probability $1 - \beta_i(\delta)$. These beliefs will be denoted by $\Phi_i^{tE}(\ast B)$. Notice that these information sets are those in the last case of (T.5.9).

Definition T.6.2: Given the strategy profile (g, μ) that we defined in Section A.2 and given Definition T.6.1, we can appeal to the stationarity of the game and of (g, μ) to define the following.

With a slight abuse of notation, for any pair of messages m and \hat{m} both in $\{m^\ast, m^A, m^B\}$, we denote by $v_i(m, \hat{m}, \delta)$ the end-of-period- t (discounted as of the beginning of period $t + 1$) payoff to player $\langle i, t \rangle$, if he sends message m , and all other players send message \hat{m} .

Lemma T.6.1: Let the assessment (g, μ, Φ) described in Section A.2 be given. Then the end-of-period continuation payoffs for any player $\langle i, t \rangle$ at information sets $\mathcal{I}_i^t \in \{\mathcal{I}_i^{tE}(\ast) \cup \mathcal{I}_i^{tE}(A) \cup \mathcal{I}_i^{tE}(B)\}$ are as follows.^{T.6}

$$v_i^t(g, \mu | \Phi_i^{tE}(\ast)) = v_i(m^\ast, m^\ast, \delta) = v_i^\ast \quad (\text{T.6.1})$$

$$v_i^t(g, \mu | \Phi_i^{tE}(A)) = v_i(m^A, m^A, \delta) = \alpha v_i' + (1 - \alpha) v_i^\ast \quad (\text{T.6.2})$$

$$v_i^t(g, \mu | \Phi_i^{tE}(B)) = v_i(m^B, m^B, \delta) = v_i' \quad (\text{T.6.3})$$

Proof: The first equalities in equations (T.6.1), (T.6.2) and (T.6.3) are obvious from Definitions T.6.1 and T.6.2.

Equations (T.6.1), (T.6.3) are a direct consequence of the way we have defined strategies and beliefs in Section A.2, and we omit the details. To see that (T.6.2) holds notice that we can write this continuation payoff recursively as

$$v_i(m^A, m^A, \delta) = (1 - \delta) v_i' + \delta [\gamma(\delta) v_i(m^\ast, m^\ast, \delta) + (1 - \gamma(\delta)) v_i(m^A, m^A, \delta)] \quad (\text{T.6.4})$$

Substituting the definition of $\gamma(\delta)$ given in (A.4), substituting (T.6.1), and solving for $v_i(m^A, m^A, \delta)$ yields (T.6.2). ■

Lemma T.6.2: Given the beliefs described in Definition T.5.4, for δ sufficiently close to one, no player $\langle i, t \rangle$ has an incentive to deviate from the message strategy described in Definition A.6 at any information set $\mathcal{I}_i^{tE} \in \mathcal{I}_i^{tE}(\ast)$.

^{T.6}See our Point of Notation T.1.1 above.

Proof: From Lemma T.6.1, if player $\langle i, t \rangle$ follows the equilibrium message strategy μ_i^t , then his continuation payoff is as in (T.6.1). If he deviates and sends m^A instead of m^* , we can write his payoff recursively as follows

$$\begin{aligned} v_i(m^A, m^*, \delta) = & \\ & q \left\{ (1 - \delta)u_i(a'_i, a^*_{-i}) + \delta [\gamma(\delta)v_i(m^*, m^*, \delta) + (1 - \gamma(\delta))v_i(m^A, m^A, \delta)] \right\} + \\ & (1 - q) \left\{ (1 - \delta)\hat{v}_i + \delta [\gamma(\delta)v_i(m^*, m^*, \delta) + (1 - \gamma(\delta))v_i(m^A, m^*, \delta)] \right\} \end{aligned} \quad (\text{T.6.5})$$

Substituting (T.6.1) and (T.6.2) and solving for $v_i(m^A, m^*, \delta)$ yields

$$\begin{aligned} v_i(m^A, m^*, \delta) = & \frac{1}{1 - (1 - q)\delta(1 - \frac{(1 - \delta)(1 - \alpha)}{\alpha\delta})} \\ & \left\{ (1 - \delta)[qu_i(a'_i, a^*_{-i}) + (1 - q)\hat{v}_i] + (1 - \delta)(1 - \alpha)\frac{v_i^*}{\alpha} + \right. \\ & \left. q\delta(1 - \frac{(1 - \delta)(1 - \alpha)}{\alpha\delta})(\alpha v'_i + (1 - \alpha)v_i^*) \right\} \end{aligned} \quad (\text{T.6.6})$$

From (T.6.6) we get directly that

$$\lim_{\delta \rightarrow 1} v_i(m^A, m^*, \delta) = \alpha v'_i + (1 - \alpha)v_i^* \quad (\text{T.6.7})$$

and since the right-hand side of (T.6.7) is obviously less than $v_i(m^*, m^*, \delta) = v_i^*$, we can conclude that player $\langle i, t \rangle$ cannot gain by deviating to sending message m^A instead of m^* when δ is large enough.

If player $\langle i, t \rangle$ deviates and sends message m^B instead of m^* we can write his payoff recursively as follows

$$\begin{aligned} v_i(m^B, m^*, \delta) = & \\ & q \left\{ (1 - \delta)u_i(a'_i, a^*_{-i}) + \delta [\gamma(\delta)v_i(m^*, m^*, \delta) + (1 - \gamma(\delta))v_i(m^A, m^A, \delta)] \right\} + \\ & (1 - q) \left\{ (1 - \delta)\hat{v}_i + \delta v_i(m^B, m^*, \delta) \right\} \end{aligned} \quad (\text{T.6.8})$$

Substituting (T.6.1) and (T.6.2) and solving for $v_i(m^B, m^*, \delta)$ yields

$$\begin{aligned} v_i(m^B, m^*, \delta) = & \frac{1}{1 - (1 - q)\delta} \\ & \left\{ (1 - \delta)[qu_i(a'_i, a^*_{-i}) + (1 - q)\hat{v}_i] + (1 - \delta)(1 - \alpha)\frac{qv_i^*}{\alpha} + \right. \\ & \left. q\delta(1 - \frac{(1 - \delta)(1 - \alpha)}{\alpha\delta})(\alpha v'_i + (1 - \alpha)v_i^*) \right\} \end{aligned} \quad (\text{T.6.9})$$

As for the previous case, if δ is sufficiently large, the deviation does not pay since from (T.6.9) we get directly that

$$\lim_{\delta \rightarrow 1} v_i(m^B, m^*, \delta) = \alpha v'_i + (1 - \alpha)v_i^* \quad (\text{T.6.10})$$

Hence the lemma is proved. ■

Lemma T.6.3: *Given the beliefs described in Definition T.5.4, for δ sufficiently close to one, no player $\langle i, t \rangle$ has an incentive to deviate from the message strategy described in Definition A.6 at any information set $\mathcal{I}_i^{tE} \in \mathcal{I}_i^{tE}(A)$.*

Proof: From Lemma T.6.1, if player $\langle i, t \rangle$ follows the equilibrium message strategy μ_i^t , then his continuation payoff is as in (T.6.2). If he deviates and sends m^* instead of m^A , we can write his payoff recursively as follows

$$\begin{aligned} v_i(m^*, m^A, \delta) = & \\ & q \{ (1 - \delta)u_i(a_i^*, a'_{-i}) + \delta [\gamma(\delta)v_i(m^*, m^*, \delta) + (1 - \gamma(\delta))v_i(m^*, m^A, \delta)] \} + \\ & (1 - q) \{ (1 - \delta)\hat{v}_i + \delta [\gamma(\delta)v_i(m^*, m^*, \delta) + (1 - \gamma(\delta))v_i(m^*, m^A, \delta)] \} \end{aligned} \quad (\text{T.6.11})$$

Substituting (T.6.1) and solving for $v_i(m^*, m^A, \delta)$ yields

$$v_i(m^*, m^A, \delta) = \alpha [qu_i(a_i^*, a'_{-i}) + (1 - q)\hat{v}_i] + (1 - \alpha)v_i^* \quad (\text{T.6.12})$$

and since the right-hand side of (T.6.12) is obviously less than $v_i(m^A, m^A, \delta) = \alpha v'_i + (1 - \alpha)v_i^*$, we can conclude that player $\langle i, t \rangle$ cannot gain by deviating to sending message m^* instead of m^A when δ is large enough.

If player $\langle i, t \rangle$ deviates and sends message m^B instead of m^A we can write his payoff recursively as follows

$$v_i(m^B, m^A, \delta) = (1 - \delta)v'_i + \delta [\gamma(\delta)v_i(m^B, m^*, \delta) + (1 - \gamma(\delta))v_i(m^B, m^A, \delta)] \quad (\text{T.6.13})$$

Solving for $v_i(m^B, m^A, \delta)$, using the definition of $\gamma(\delta)$ given in (A.4) yields

$$v_i(m^B, m^A, \delta) = \alpha v'_i + (1 - \alpha)v_i(m^B, m^*, \delta) \quad (\text{T.6.14})$$

Using (T.6.10) this is clearly enough to show that for δ high enough player $\langle i, t \rangle$ cannot gain by deviating in this way. Hence the lemma is proved. ■

Lemma T.6.4: *Given the beliefs described in Definition T.5.4 no player $\langle i, t \rangle$ has an incentive to deviate from the message strategy described in Definition A.6 at any information set $\mathcal{I}_i^{tE} \in \mathcal{I}_i^{tE}(B)$.*

Proof: From Lemma T.6.1, if player $\langle i, t \rangle$ follows the equilibrium message strategy μ_i^t , then his continuation payoff is as in (T.6.3). If he deviates and sends m^* instead of m^B , we can write his payoff recursively as follows

$$v_i(m^*, m^B, \delta) = (1 - \delta) [qu_i(a_i^*, a'_{-i}) + (1 - q)\hat{v}_i] + \delta v_i(m^*, m^B, \delta) \quad (\text{T.6.15})$$

Solving for $v_i(m^*, m^B, \delta)$ yields

$$v_i(m^*, m^B, \delta) = qu_i(a_i^*, a'_{-i}) + (1 - q)\hat{v}_i \quad (\text{T.6.16})$$

Since the right-hand side of (T.6.16) is strictly less than $v_i(m^B, m^B, \delta)$, this is clearly enough to show that for δ high enough player $\langle i, t \rangle$ cannot gain by deviating in this way.

Now suppose that player $\langle i, t \rangle$ deviates and send message m^A instead of m^B . Then we can write his continuation payoff recursively as

$$v_i(m^A, m^B, \delta) = (1 - \delta) [qu_i(a') + (1 - q)\hat{v}_i] + \delta [\gamma(\delta)v_i(m^*, m^B, \delta) + (1 - \gamma(\delta))v_i(m^A, m^B, \delta)] \quad (\text{T.6.17})$$

Solving for $v_i(m^A, m^B, \delta)$ gives us

$$v_i(m^A, m^B, \delta) = \alpha v'_i + (1 - \alpha)v_i(m^*, m^B, \delta) \quad (\text{T.6.18})$$

Since the right-hand side of (T.6.18) is strictly less than $v_i(m^B, m^B, \delta)$, this is clearly enough to show that player $\langle i, i \rangle$ cannot gain by deviating in this way. Hence the lemma is proved. ■

Lemma T.6.5: *Given the beliefs described in Definition T.5.4, for δ sufficiently close to one, no player $\langle i, t \rangle$ has an incentive to deviate from the message strategy described in Definition A.6 at any information set $\mathcal{I}_i^{tE} \in \mathcal{I}_i^{tE}(*B)$.*

Proof: From Lemma T.6.1, and the last case in (T.5.9) if player $\langle i, t \rangle$ follows the equilibrium message strategy μ_i^t , then his continuation payoff is

$$(1 - \beta_i(\delta))v_i(m^B, m^B, \delta) + \beta_i(\delta)v_i(m^B, m^*, \delta) \quad (\text{T.6.19})$$

From Definition T.5.2 it is clear that $\lim_{\delta \rightarrow 1} \beta_i(\delta) = 0$. Hence, using (T.6.3) we can write

$$\lim_{\delta \rightarrow 1} (1 - \beta_i(\delta))v_i(m^B, m^B, \delta) + \beta_i(\delta)v_i(m^B, m^*, \delta) = v'_i \quad (\text{T.6.20})$$

If he deviates and sends message m^* instead of m^B , we can write his continuation payoff as

$$(1 - \beta_i(\delta))v_i(m^*, m^B, \delta) + \beta_i(\delta)v_i(m^*, m^*, \delta) \quad (\text{T.6.21})$$

Using again the fact that $\lim_{\delta \rightarrow 1} \beta_i(\delta) = 0$ and (T.6.15) we can write

$$\lim_{\delta \rightarrow 1} (1 - \beta_i(\delta))v_i(m^*, m^B, \delta) + \beta_i(\delta)v_i(m^*, m^*, \delta) = qu_i(a_i^*, a'_{-i}) + (1 - q)\hat{v}_i \quad (\text{T.6.22})$$

Since the quantity on the right-hand side of (T.6.20) is clearly greater than the quantity on the right-hand side of (T.6.22) we can then conclude that, for δ close enough to one, player $\langle i, t \rangle$ cannot gain by deviating in this way.

Suppose now that player $\langle i, t \rangle$ deviates to sending message m^A instead of m^B . Then we can write his continuation payoff as

$$(1 - \beta_i(\delta))v_i(m^A, m^B, \delta) + \beta_i(\delta)v_i(m^A, m^*, \delta) \quad (\text{T.6.23})$$

Using once more the fact that $\lim_{\delta \rightarrow 1} \beta_i(\delta) = 0$, (T.6.18) and (T.6.16) we can write

$$\lim_{\delta \rightarrow 1} (1 - \beta_i(\delta))v_i(m^A, m^B, \delta) + \beta_i(\delta)v_i(m^A, m^*, \delta) = \alpha v'_i + (1 - \alpha) [qu_i(a_i^*, a'_{-i}) + (1 - q)\hat{v}_i] \quad (\text{T.6.24})$$

Since the quantity on the right-hand side of (T.6.20) is clearly greater than the quantity on the right-hand side of (T.6.24) we can then conclude that, for δ close enough to one, player $\langle i, t \rangle$ cannot gain by deviating in this way. Hence the lemma is proved. ■

Remark T.6.1: From Lemmas T.6.2, T.6.3, T.6.4 and T.6.5 it is clear that there exists a $\underline{\delta} \in (0, 1)$ such that whenever $\delta > \underline{\delta}$ the message strategies of Definition A.6 are sequentially rational given the beliefs of Definition T.5.4.

We now turn to the sequential rationality of the action strategies we have defined in Section A.2.

Definition T.6.3: Recall that at the action stage, player $\langle i, t \rangle$ chooses an action after having received a message m_i^t and having observed a realization x^t of the randomization device \tilde{x}^t .

Let \mathcal{I}_i^{tB} denote period- t action-stage collection of information sets that belong to player $\langle i, t \rangle$, with typical element \mathcal{I}_i^{tB} . Clearly, each element of \mathcal{I}_i^{tB} is identified by a pair (m_i^t, x^t) .

It is convenient to partition \mathcal{I}_i^{tB} into three mutually disjoint exhaustive subsets on the basis of the message m_i^t received by player $\langle i, t \rangle$.^{T.7}

Let $\mathcal{I}_i^{tB}(\ast) \subset \mathcal{I}_i^{tB}$ be the collection of information sets in which player $\langle i, t \rangle$ receives message m^\ast . Notice that using Definition T.5.3 we know that in this case player $\langle i, t \rangle$ believes that m_{-i}^t is equal to (m^\ast, \dots, m^\ast) with probability one. These beliefs will be denoted by $\Phi_i^{tB}(\ast)$.

Let $\mathcal{I}_i^{tB}(A) \subset \mathcal{I}_i^{tB}$ be the collection of information sets in which player $\langle i, t \rangle$ receives message m^A . Notice that using Definition T.5.3 we know that in this case player $\langle i, t \rangle$ believes that m_{-i}^t is equal to (m^A, \dots, m^A) with probability one. These beliefs will be denoted by $\Phi_i^{tB}(A)$.

Finally, let $\mathcal{I}_i^{tB}(B) \subset \mathcal{I}_i^{tB}$ be the collection of information sets in which player $\langle i, t \rangle$ receives message m^B . Notice that using Definition T.5.3 we know that in this case player $\langle i, t \rangle$ believes that m_{-i}^t is equal to (m^\ast, \dots, m^\ast) with probability $\beta_i(\delta)$ and to (m^B, \dots, m^B) with probability $1 - \beta_i(\delta)$. These beliefs will be denoted by $\Phi_i^{tB}(\ast B)$.

Lemma T.6.6: Given the beliefs described in Definition T.5.3, for δ sufficiently close to one, no player $\langle i, t \rangle$ has an incentive to deviate from the action strategy described in Definition A.5 at any information set $\mathcal{I}_i^{tB} \in \mathcal{I}_i^{tB}(\ast)$.

Proof: Given any $x^t \in X$, if player $\langle i, t \rangle$ follows the equilibrium action strategy he achieves a payoff is that bounded below by

$$(1 - \delta)\underline{u}_i + \delta v_i^\ast \tag{T.6.25}$$

Given any $x^t \in X$, (using the notation of Definition T.6.2) following any possible deviation player $\langle i, t \rangle$ achieves a payoff that is bounded above by

$$(1 - \delta)\bar{u}_i + \delta [\gamma(\delta)v_i(m^\ast, m^\ast, \delta) + (1 - \gamma(\delta))v_i(m^A, m^A, \delta)] \tag{T.6.26}$$

Using (T.6.1), (T.6.2) and the definition of $\gamma(\delta)$ given in (A.4), we can re-write (T.6.26) as

$$(1 - \delta)\bar{u}_i + \delta \alpha v_i' \left[1 - \frac{(1 - \delta)(1 - \alpha)}{\delta \alpha} \right] + \delta v_i^\ast \left\{ 1 - \alpha \left[1 - \frac{(1 - \delta)(1 - \alpha)}{\delta \alpha} \right] \right\} \tag{T.6.27}$$

Taking the limit of (T.6.25) as $\delta \rightarrow 1$ gives v_i^\ast . Taking the limit of (T.6.27) as $\delta \rightarrow 1$ gives $\alpha v_i' + (1 - \alpha)v_i^\ast$. Since $v_i^\ast > \alpha v_i' + (1 - \alpha)v_i^\ast$ this is clearly enough to prove the claim. ■

Lemma T.6.7: Given the beliefs described in Definition T.5.3, for δ sufficiently close to one, no player $\langle i, t \rangle$ has an incentive to deviate from the action strategy described in Definition A.5 at any information set $\mathcal{I}_i^{tB} \in \mathcal{I}_i^{tB}(A)$.

^{T.7}In the interest of brevity, we avoid an explicit distinction between the $t = 0$ players and all others. What follows can be interpreted as applying to all players re-defining m_i^0 to be equal to m^\ast for players $\langle i \in I, 0 \rangle$.

Proof: We distinguish between the information set in $\mathcal{I}_i^{tB}(A)$ that has $x^t = x(1)$ and those that have $x^t \neq x(1)$. We begin with the information set in which $x^t = x(1)$.

After having observed the pair $(m^A, x(1))$, if he follows the equilibrium strategy, player $\langle i, t \rangle$ achieves a continuation payoff of

$$(1 - \delta)u_i(a') + \delta [\gamma(\delta)v_i(m^*, m^*, \delta) + (1 - \gamma(\delta))v_i(m^A, m^A, \delta)] \quad (\text{T.6.28})$$

Now consider a deviation to action a_i^* . In this case the continuation payoff is

$$(1 - \delta)u_i(a_i^*, a'_{-i}) + \delta [\gamma(\delta)v_i(m^*, m^*, \delta) + (1 - \gamma(\delta))v_i(m^A, m^A, \delta)] \quad (\text{T.6.29})$$

Clearly this deviation is not profitable since $u_i(a_i^*, a'_{-i}) < u_i(a')$.

A deviation to an action $a_i \notin \{a_i', a_i^*\}$ yields a continuation payoff that is bounded above by

$$(1 - \delta)\bar{u}_i + \delta v_i' \quad (\text{T.6.30})$$

Taking the limit of (T.6.28) as $\delta \rightarrow 1$ yields $\alpha v_i' + (1 - \alpha)v_i^*$. Hence, for δ large enough, the quantity in (T.6.28) is greater than the quantity in (T.6.30). Therefore, for δ close enough to one, this deviation is not profitable either.

Now consider any information set in $\mathcal{I}_i^{tB}(A)$ with $x^t \neq x(1)$. The continuation payoff to player $\langle i, t \rangle$ from following the equilibrium strategy is bounded below by

$$(1 - \delta)\underline{u}_i + \delta [\gamma(\delta)v_i(m^*, m^*, \delta) + (1 - \gamma(\delta))v_i(m^A, m^A, \delta)] \quad (\text{T.6.31})$$

The continuation payoff to player $\langle i, t \rangle$ following any deviation is bounded above by

$$(1 - \delta)\bar{u}_i + \delta v_i' \quad (\text{T.6.32})$$

Taking the limit of (T.6.31) as $\delta \rightarrow 1$ yields $\alpha v_i' + (1 - \alpha)v_i^*$. Hence, for δ large enough, the quantity in (T.6.31) is greater than the quantity in (T.6.32). Therefore, for δ close enough to one, no deviation is profitable at any information set in $\mathcal{I}_i^{tB}(A)$ with $x^t \neq x(1)$. ■

Lemma T.6.8: *Given the beliefs described in Definition T.5.3, for δ sufficiently close to one, no player $\langle i, t \rangle$ has an incentive to deviate from the action strategy described in Definition A.5 at any information set $\mathcal{I}_i^{tB} \in \mathcal{I}_i^{tB}(B)$.*

Proof: We distinguish between the information set in $\mathcal{I}_i^{tB}(B)$ that has $x^t = x(1)$ and those that have $x^t \neq x(1)$. We begin with the information set in which $x^t = x(1)$.

After having observed the pair $(m^B, x(1))$, if he follows the equilibrium strategy, player $\langle i, t \rangle$ achieves a continuation payoff of

$$\beta_i(\delta) [(1 - \delta)u_i(a_i', a'_{-i}) + \delta(\gamma(\delta)v_i(m^*, m^*, \delta) + (1 - \gamma(\delta))v_i(m^A, m^A, \delta))] + (1 - \beta_i(\delta)) [(1 - \delta)u_i(a') + \delta v_i(m^B, m^B, \delta)] \quad (\text{T.6.33})$$

If on the other hand he deviates to action a_i^* his continuation payoff is

$$\beta_i(\delta) [(1 - \delta)u_i(a_i^*) + \delta v_i(m^*, m^*, \delta)] + (1 - \beta_i(\delta)) [(1 - \delta)u_i(a_i^*, a'_{-i}) + \delta v_i(m^B, m^B, \delta)] \quad (\text{T.6.34})$$

Using (T.6.1) (T.6.2) (T.6.3) and Definition T.5.2, we can now write the equilibrium continuation payoff in (T.6.33) minus the deviation continuation payoff in (T.6.34) as

$$(1 - \delta) \{ r_i(1 - \delta) [u_i(a_i', a'_{-i}) - u_i(a_i^*)] + r_i \delta (1 - \gamma(\delta)) \alpha (v_i' - v_i^*) + (1 - (1 - \delta)r_i) [u_i(a') - u_i(a_i^*, a'_{-i})] \} \quad (\text{T.6.35})$$

Dividing (T.6.35) by $1 - \delta$, taking the limit as $\delta \rightarrow 1$ and using (A.4), yields that (up to a factor $1 - \delta$) this difference in payoffs in the limit is equal to

$$r_i \alpha (v'_i - v_i^*) + u_i(a') - u_i(a_i^*, a'_{-i}) \quad (\text{T.6.36})$$

Notice now that using the (upper) bound on r_i given in (T.5.6) we can verify that the quantity in (T.6.36) is positive. Hence we can conclude that deviating to a_i^* is in fact not profitable for player $\langle i, t \rangle$.

Next, consider a deviation to an action $a_i \notin \{a_i^*, a'_{-i}\}$. Following this deviation, the continuation payoff to player $\langle i, t \rangle$ is bounded above by

$$(1 - \delta)\bar{u}_i + \delta v_i(m^B, m^B, \delta) \quad (\text{T.6.37})$$

Using (T.6.1), (T.6.2), (T.6.3) and Definition T.5.2, we can now write the equilibrium continuation payoff in (T.6.33) minus the deviation continuation payoff in (T.6.37) as

$$(1 - \delta) \{ \delta r_i (1 - \alpha (1 - \gamma(\delta))) (v_i^* - v'_i) + (1 - \delta) r_i u_i(a_i^*, a'_{-i}) + (1 - (1 - \delta) r_i) u_i(a') - \bar{u}_i \} \quad (\text{T.6.38})$$

Dividing (T.6.38) by $1 - \delta$, taking the limit as $\delta \rightarrow 1$ and using (A.4) and the fact that $\beta_i(\delta) = (1 - \delta) r_i$, yields that (up to a factor $1 - \delta$) this difference in payoffs in the limit is equal to

$$r_i (1 - \alpha) (v_i^* - v'_i) + u_i(a') - \bar{u}_i \quad (\text{T.6.39})$$

Notice now that using the (lower) bound on r_i given in (T.5.6) we can verify that the quantity in (T.6.39) is positive. Hence we can conclude that deviating to $a_i \notin \{a_i^*, a'_{-i}\}$ is in fact not profitable for player $\langle i, t \rangle$.

Now consider an information set in $\mathcal{I}_i^{tB}(B)$ that has $x^t \neq x(1)$. If he follows his equilibrium strategy, player $\langle i, t \rangle$ achieves a continuation payoff that is bounded below by

$$(1 - \delta)\underline{u}_i + \delta[\beta_i(\delta)v_i(m^B, m^*, \delta) + (1 - \beta_i(\delta))v_i(m^B, m^B, \delta)] \quad (\text{T.6.40})$$

His continuation payoff following any deviation is bounded above by

$$(1 - \delta)\bar{u}_i + \delta v_i(m^B, m^B, \delta) \quad (\text{T.6.41})$$

Using (T.6.3) and Definition T.5.2, we can now write the equilibrium continuation payoff in (T.6.40) minus the deviation continuation payoff in (T.6.41) as

$$(1 - \delta) \{ \underline{u}_i - \bar{u}_i + \delta r_i [v_i(m^B, m^*, \delta) - v'_i] \} \quad (\text{T.6.42})$$

Dividing (T.6.42) by $1 - \delta$, taking the limit as $\delta \rightarrow 1$ and using (T.6.10), yields that (up to a factor $1 - \delta$) this difference in payoffs in the limit is equal to

$$\underline{u}_i - \bar{u}_i + r_i (1 - \alpha) (v_i^* - v'_i) \quad (\text{T.6.43})$$

Notice now that using the (lower) bound on r_i given in (T.5.6) we can verify that the quantity in (T.6.43) is positive. Hence we can conclude that deviating from the equilibrium strategy at any information set in $\mathcal{I}_i^{tB}(B)$ that has $x^t \neq x(1)$ is in fact not profitable for player $\langle i, t \rangle$. Therefore, the proof is now complete. ■

Remark T.6.2: From Lemmas T.6.6, T.6.7 and T.6.8 it is clear that there exists a $\underline{\delta} \in (0, 1)$ such that whenever $\delta > \underline{\delta}$ the action strategies of Definition A.5 are sequentially rational given the beliefs of Definition T.5.3.

T.7. Proof of Theorem 2: Consistency of Beliefs

Definition T.7.1: For every $i \in I$, let

$$\psi_i(\delta) = \frac{(1 - \delta)r_i \sum_{j \in I} (\|A_j\| - 1 - q)}{1 - (1 - \delta)r_i} \quad (\text{T.7.1})$$

where q is as in Definition A.1, and r_i is as in Definition T.5.2.

Remark T.7.1: For δ sufficiently close to 1, it is clear that the $\psi_i(\delta)$ of equation (T.7.1) is in $(0, 1)$.

Remark T.7.2: Let $(g_\varepsilon, \mu_\varepsilon)$ be a completely mixed strategy profile of Definitions A.8 and A.9. It is then straightforward to check that as $\varepsilon \rightarrow 0$ the profile $(g_\varepsilon, \mu_\varepsilon)$ converges pointwise to the equilibrium strategy profile described in Definitions A.5 and A.6, as required.

Lemma T.7.1: Let $(g_\varepsilon, \mu_\varepsilon)$ be the completely mixed strategy profile of Definitions A.8 and A.9. Let any $t \geq 2$ and any quadruple of the type (m_i^t, x^t, a^t, y^t) be given.^{T.8}

Then^{T.9}

$$\lim_{\varepsilon \rightarrow 0} \Pr[m^{t-1} = (m^*, \dots, m^*) \mid m_i^t, x^t, a^t, y^t, g_\varepsilon, \mu_\varepsilon] = 1 \quad (\text{T.7.2})$$

Proof: In order for $m^{t-1} \neq (m^*, \dots, m^*)$ to occur it is necessary that at least one player has deviated from the equilibrium strategy in some period $\tau \leq t-2$, either at the action or at the message stage (or both). Given the completely mixed strategy profile of Definitions A.8 and A.9 and given that from (A.10) we know that trembles become more likely as t increases, we then know that the probability of event $m^{t-1} \neq (m^*, \dots, m^*)$ is an infinitesimal in ε_{t-2} of order no lower than $1/2$.^{T.10} Hence, using (A.10) the probability of $m^{t-1} \neq (m^*, \dots, m^*)$ is an infinitesimal in ε of order no lower than $1/2n^{2(t-2)}$.

The probability of $m^{t-1} \neq (m^*, \dots, m^*)$ needs to be compared with the probability of the quadruple (m_i^t, x^t, a^t, y^t) . Depending on the particular quadruple (m_i^t, x^t, a^t, y^t) , it is possible that many paths of play could have generated this outcome. However, a lower bound on this probability can be computed as follows. Assume no deviations up to and including period $t-2$. From Definition A.9 the probability that message m_i^t is sent by player $\langle i, t \rangle$ is at least (if a deviation is required) an infinitesimal in ε_{t-1} of order 2. From Definition A.8, the probability of any profile a^t (depending on the number of deviations required; clearly no more than n) is at least an infinitesimal in ε_t of order n . Hence, using (A.10), the probability of the quadruple (m_i^t, x^t, a^t, y^t) is no smaller than an infinitesimal in ε of order $2/n^{2(t-1)} + 1/n^{2t-1}$.

Since $n \geq 3$ it is straightforward to check that $1/2n^{2(t-2)} > 2/n^{2(t-1)} + 1/n^{2t-1}$. Hence equation (T.7.2) now follows and the proof is complete.^{T.11} ■

Lemma T.7.2: Let $(g_\varepsilon, \mu_\varepsilon)$ be the completely mixed strategy profile of Definitions A.8 and A.9. Let any $t \geq 2$ and any quadruple of the type (m_i^t, x^t, a^t, y^t) be given.^{T.12} Fix also any array $\hat{m}_{-i} = (\hat{m}_1, \dots, \hat{m}_{i-1}, \hat{m}_{i+1}, \dots, \hat{m}_n)$.

^{T.8}The reason we require that $t \geq 2$ in (T.7.2) below is that of course all players $\langle i \in I, 0 \rangle$ receive message m_i^0 for sure.

^{T.9}See our point of notation T.1.3 above.

^{T.10}See footnote T.2 for an explicit statement of our (standard) terminology concerning infinitesimals.

^{T.11}It is worth pointing out that the bounds on probabilities that we have used in this argument are not “tight.” We have used the ones above simply because they facilitate the exposition. Any tight bounds would necessarily involve a case-by-case treatment according to what the equilibrium strategies prescribe, the particular message vector $m^{t-1} \neq (m^*, \dots, m^*)$ and the particular quadruple (m_i^t, x^t, a^t, y^t) .

^{T.12}The reason we require that $t \geq 2$ in (T.7.2) below is that of course all players $\langle i \in I, 0 \rangle$ receive message m_i^0 for sure.

Then

$$\lim_{\varepsilon \rightarrow 0} \Pr [m_{-i}^t = \hat{m}_{-i} | m_i^t, x^t, a^t, y^t, g_\varepsilon, \mu_\varepsilon] = \lim_{\varepsilon \rightarrow 0} \Pr [m_{-i}^t = \hat{m}_{-i} | m_i^t, x^t, a^t, y^t, m^{t-1} = (m^*, \dots, m^*), g_\varepsilon, \mu_\varepsilon] \quad (\text{T.7.3})$$

Proof: A routine application of Bayes’ rule yields

$$\begin{aligned} & \Pr [m_{-i}^t = \hat{m}_{-i} | m_i^t, x^t, a^t, y^t, g_\varepsilon, \mu_\varepsilon] = \\ & \frac{\Pr [m_{-i}^t = \hat{m}_{-i} | m_i^t, x^t, a^t, y^t, m^{t-1} = (m^*, \dots, m^*), g_\varepsilon, \mu_\varepsilon]}{\Pr [m^{t-1} = (m^*, \dots, m^*) | m_i^t, x^t, a^t, y^t, g_\varepsilon, \mu_\varepsilon]} + \\ & \sum_{m \neq (m^*, \dots, m^*)} \Pr [m_{-i}^t = \hat{m}_{-i} | m_i^t, x^t, a^t, y^t, m^{t-1} = m, g_\varepsilon, \mu_\varepsilon] \Pr (m^{t-1} = m | m_i^t, x^t, a^t, y^t, g_\varepsilon, \mu_\varepsilon) \end{aligned} \quad (\text{T.7.4})$$

Now take the limit as $\varepsilon \rightarrow 0$ on both sides of (T.7.4). Next, observe that by Lemma T.7.1 all terms in the summation sign must converge to zero and the second term on the right-hand-side of (T.7.4) must converge to one. Hence (T.7.3) follows and the proof is now complete. ■

Lemma T.7.3: *The strategy profile (g, μ) described in Definitions A.5 and A.6 and the beginning-of-period beliefs described in Definition T.5.3 are consistent.*

Proof: For the players $\langle i \in I, 0 \rangle$ there is nothing to prove. When $t = 1$, clearly all players $\langle i \in I, 1 \rangle$ believe that all preceding players received their respective m_i^0 . When $t \geq 2$, given Lemma T.7.2 we can reason taking it as given that all players $\langle i \in I, t \rangle$ believe that all players $\langle i \in I, t - 1 \rangle$ received message m^* . Given this observation the claim for period t follows easily by a case-by-case examination, comparing the likelihood of deviations in period $t - 1$ (orders of infinitesimals in ε_{t-1}). We omit the details entirely for the first two cases of (T.5.7) (in which messages m^* and m^A are received).

To see the consistency of the beliefs postulated in the third case of (T.5.7) (when message m^B is received) observe the following. According to Definition A.8 deviations at the action stage of period $t - 1$ after receiving message m^* have probability ε_{t-1} . Moreover, according to Definition A.9, after receiving message m^* player $\langle i, t - 1 \rangle$ sends message m^B with probability $\psi_i(\delta) \varepsilon_{t-1}$ when the equilibrium strategy prescribes to send message m^* . Using Lemma T.7.2, (T.7.1) and Definition T.5.2 it is then immediate to see that Definitions A.8 and A.9 yield

$$\lim_{\varepsilon \rightarrow 0} \Phi_i^{tB} [m_{-i}^t = (m_i^*, \dots, m_i^*) | m_i^t = m^B, g_\varepsilon, \mu_\varepsilon] = \frac{\psi_i(\delta)}{\psi_i(\delta) + \sum_{j \in I} (\|A_j\| - 1 - q)} = (1 - \delta)r_i = \beta_i(\delta) \quad (\text{T.7.5})$$

and

$$\lim_{\varepsilon \rightarrow 0} \Phi_i^{tB} [m_{-i}^t = (m_i^B, \dots, m_i^B) | m_i^t = m^B, g_\varepsilon, \mu_\varepsilon] = \frac{\sum_{j \in I} (\|A_j\| - 1 - q)}{\psi_i(\delta) + \sum_{j \in I} (\|A_j\| - 1 - q)} = 1 - (1 - \delta)r_i = 1 - \beta_i(\delta) \quad (\text{T.7.6})$$

as required. ■

Lemma T.7.4: *The strategy profile (g, μ) described in Definitions A.5 and A.6 and the end-of-period beliefs described in Definition T.5.4 are consistent, as required for an SE.*

Proof: When $t = 0$ all players receive message m_i^0 . When $t = 1$ the beliefs of the players about the messages received by the previous cohort are trivial. When $t \geq 2$, given Lemma T.7.2 we can reason taking it as given that all players $\langle i \in I, t \rangle$ believe that all players $\langle i \in I, t-1 \rangle$ received message m^* . Given this observation the claim for period t follows easily by a case-by-case examination, comparing the likelihood of possible deviations in periods $t-1$ and t (orders of infinitesimals in ε_{t-1} and ε_t). For the sake of brevity we omit most of the details, and we simply draw attention to the following facts.

The beliefs of player $\langle i, t \rangle$ after any realization $x^t = x(\ell)$ with $\ell \geq 2$ can be seen to be consistent in the following way. In period t , either no deviation from the prescription of the action-stage equilibrium strategies is observed or some deviation is observed. When there are no deviations, the revised end-of-period beliefs Φ_i^{tR} of player $\langle i, t \rangle$ are of course the same as the beginning of period beliefs Φ_i^{tB} . Therefore, using Lemma T.7.2 the claim in this case can be verified simply checking that the beliefs described in (T.5.9) correspond to the prescriptions of the strategies described in (A.9) in the appropriate way. In the case in which some deviations occur, observe that the strategies in (A.9) prescribe that all players should send message m^B , regardless of the message they received, which corresponds to the beliefs described in (T.5.9) as required.

Now consider the case in which the realization of the action-stage randomization device is $x^t = x(1)$. Next distinguish further between two cases. First, the action profile a^t is neither equal to a' nor is of the type (a_j^*, a'_{-j}) for some $j \in I$. In this case, from the strategies described in (A.8) it is immediate to check that the message sent by any player $\langle i, t \rangle$ does not depend on the message m_i^t he received. Therefore, the claim in this case follows immediately from the message-stage strategies described in (A.8).

On the other hand, from the strategies described in (A.8) it is immediate to check that when a^t is either equal to a' or is of the type (a_j^*, a'_{-j}) for some $j \in I$ the message sent by player $\langle i, t \rangle$ does depend on the message m_i^t he received.

Using the completely mixed strategies described in Definition A.8 it is easy to check that for all $m \in \{m^*, m^A, m^B\}$ and for all $y^t \in Y$, whenever a^t is either equal to a' or is of the type (a_j^*, a'_{-j}) for some $j \in I$ it must be that

$$\lim_{\varepsilon \rightarrow 0} \Phi_i^{tR}[m_{-i}^t = (m, \dots, m) \mid m_i^t = m, x(1), a^t, y^t, g_\varepsilon, \mu_\varepsilon] = 1 \quad (\text{T.7.7})$$

Given (T.7.7), the claim in this case can now be verified simply checking that the beliefs described in (T.5.8) correspond to the prescriptions of the strategies described in (A.8) in the appropriate way. ■

T.8. Proof of Theorem 2

Given any $v^* \in \text{int}(V)$ and any $\delta \in (0, 1)$, the strategies and randomization devices described in Definitions A.2, A.4, A.5 and A.6 clearly implement the payoff vector v^* .

From Remarks T.6.1 and T.6.2, we know that there exists a $\underline{\delta}$ such that whenever $\delta > \underline{\delta}$ each strategy in the profile described in Definitions A.5 and A.6 is sequentially rational given the beliefs described in Definitions T.5.3 and T.5.4.

From Lemmas T.7.3 and T.7.4 we know that the strategy profile described in Definitions A.5 and A.6 and the beliefs described in Definitions T.5.3 and T.5.4 are consistent, as required by Definition 1 of SE.

Hence, using Lemma T.2.1, the proof of Theorem 2 is now complete. ■

T.9. Proof of Theorem 3: Beliefs

Definition T.9.1. *Beginning-of-Period Beliefs:* Let k be any element of I , and j be any element of I not equal to i .

The beginning-of-period beliefs of all players $\langle i \in I, 0 \rangle$ are trivial. Of course, all players believe that all other players have received the null message $m_i^0 = \emptyset$.

The beginning-of-period beliefs $\Phi_i^{tE}(m_i^t)$ of any other player $\langle i, t \rangle$, depending on the message he receives from player $\langle i, t-1 \rangle$ are as follows^{T.13}

$$\begin{aligned}
& \text{if } m_i^t = m^* \quad \text{then } m_{-i}^t = (m^*, \dots, m^*) \text{ with probability 1} \\
& \text{if } m_i^t = \check{m}^j \quad \text{then } \begin{cases} m_{-i-j}^t = (\check{m}^j, \dots, \check{m}^j) & \text{with pr. 1} \\ m_j^t \in \underline{M}(j, t) & \text{with pr. 1} \\ \Pr(m_j^t = \underline{m}^{j,\tau}) > 0 & \forall \underline{m}^{j,\tau} \in \underline{M}(j, t) \end{cases} \\
& \text{if } m_i^t = \underline{m}^{j,\tau} \quad \text{then } m_{-i}^t = (\underline{m}^{j,\tau}, \dots, \underline{m}^{j,\tau}) \text{ with probability 1} \\
& \text{if } m_i^t = \check{m}^{i,\tau} \quad \text{then } m_{-i}^t = (\check{m}^i, \dots, \check{m}^i) \text{ with probability 1} \\
& \text{if } m_i^t = \bar{m}^k \quad \text{then } m_{-i}^t = (\bar{m}^k, \dots, \bar{m}^k) \text{ with probability 1}
\end{aligned} \tag{T.9.1}$$

Definition T.9.2. End-of-Period Beliefs: Let k be any element of I , and j be any element of I not equal to i .

We begin with period $t = 0$. Recall that $m_i^0 = \emptyset$ for all $i \in I$. As before, let also $g^0(m^0, x^0) = (g_1^0(m_1^0, x^0), \dots, g_n^0(m_n^0, x^0))$, and define $g_{-k}^0(m^0, x^0)$ in the obvious way.

Let $\Phi_i^{0E}(m_i^0, x^0, a^0, y^0)$ be as follows

$$\begin{aligned}
& \text{if } a^0 = g^0(m^0, x^0) \text{ and } y^0 = y(j) && \text{then } m_{-i-j}^1 = (\check{m}^j, \dots, \check{m}^j), m_j^1 = \underline{m}^{j,T} \text{ with pr. 1} \\
& \text{if } a^0 = g^0(m^0, x^0) \text{ and } y^0 = y(i) && \text{then } m_{-i}^1 = (\check{m}^i, \dots, \check{m}^i) \text{ with probability 1} \\
& \text{if } a_{-k}^0 = g_{-k}^0(m^0, x^0) \text{ and } a_k^0 \neq g_k^0(m_k^0, x^0) && \text{then } m_{-i}^1 = (\underline{m}^{k,T}, \dots, \underline{m}^{k,T}) \text{ with prob. 1} \\
& \text{otherwise} && m_{-i}^1 = (m^*, \dots, m^*) \text{ with probability 1}
\end{aligned} \tag{T.9.2}$$

Our next case is $t \geq 1$ and $x^t = x(\kappa)$ with $\kappa > \bar{\kappa}$. Let $x(\ell_{00}, \ell^*)$ denote the realization of x^t . For any player $\langle i, t \rangle$, let $\Phi_i^{tE}(m_i^t, x(\ell_{00}, \ell^*), a^t, y^t)$ be as follows^{T.14}

$$\begin{aligned}
& \text{if } a^t = a(\ell^*) \text{ and } m_i^t = \check{m}^j && \text{then } \begin{cases} m_{-i-j}^{t+1} = (\check{m}^j, \dots, \check{m}^j) & \text{with pr. 1} \\ m_j^{t+1} \in \underline{M}(j, t) & \text{with pr. 1} \\ \Pr(m_j^{t+1} = \underline{m}^{j,\tau}) > 0 & \forall \underline{m}^{j,\tau} \in \underline{M}(j, t) \end{cases} \\
& \text{if } a^t = a(\ell^*) \text{ and } m_i^t = \underline{m}^{j,\tau} && \text{then } m_{-i}^{t+1} = (\underline{m}^{j,\tau}, \dots, \underline{m}^{j,\tau}) \text{ with probability 1} \\
& \text{if } a^t = a(\ell^*) \text{ and } m_i^t = \check{m}^{i,\tau} && \text{then } m_{-i}^{t+1} = (\check{m}^i, \dots, \check{m}^i) \text{ with probability 1} \\
& \text{if } a^t = a(\ell^*) \text{ and } m_i^t = \bar{m}^k && \text{then } m_{-i}^{t+1} = (\bar{m}^k, \dots, \bar{m}^k) \text{ with probability 1} \\
& \text{if } a_{-k}^t = a_{-k}(\ell^*) \text{ and } a_k^t \neq a_k(\ell^*) && \text{then } m_{-i}^{t+1} = (\underline{m}^{k,T}, \dots, \underline{m}^{k,T}) \text{ with probability 1} \\
& \text{otherwise} && m_{-i}^{t+1} = (m^*, \dots, m^*) \text{ with probability 1}
\end{aligned} \tag{T.9.3}$$

We divide the case of $t \geq 1$ and $x^t = x(\kappa)$ with $\kappa \leq \bar{\kappa}$ into several subcases, according to which message player $\langle i, t \rangle$ has received. We begin with $m_i^t = m^*$. Let $x(\cdot, \hat{\ell}, \dots)$ denote the realization of x^t . For any player $\langle i, t \rangle$, with the understanding that $\underline{m}^{j,\tau}$ is a generic element of $\underline{M}(j, t+1)$, let $\Phi_i^{tE}(m^*, x(\cdot, \hat{\ell}, \dots), a^t, y^t)$ be

^{T.13}Notice that the second line of (T.9.1) does not fully specify the probability distribution over the component m_j^t of the beliefs of player $\langle i, t \rangle$. For the rest of the argument, what matters is only that all elements of $\underline{M}(j, t)$ have positive probability, and that no message outside this set has positive probability. The distribution can be computed using Bayes' rule from the equilibrium strategies described in Definitions B.8 and B.9 above. We omit the details for the sake of brevity.

^{T.14}Similarly to (T.9.1), the first line of (T.9.3) does not fully specify the probability distribution over the component m_j^{t+1} of the beliefs of player $\langle i, t \rangle$. For the rest of the argument, what matters is only that all elements of $\underline{M}(j, t)$ have positive probability, and that no message outside this set has positive probability. The distribution can be computed using Bayes' rule from the equilibrium strategies described in Definitions B.8 and B.9 above. We omit the details for the sake of brevity.

as follows

$$\begin{array}{ll}
\text{if } a^t = a(\hat{\ell}) \text{ and } y^t = y(j) & \text{then } \begin{cases} m_{-i-j}^{t+1} = (\check{m}^j, \dots, \check{m}^j) \\ m_j^{t+1} = \underline{m}^{j,\tau} \end{cases} \text{ with pr. } \frac{1}{\|\underline{M}(j, t+1)\|} \\
\text{if } a^t = a(\hat{\ell}) \text{ and } y^t = y(i) & \text{then } m_{-i}^{t+1} = (\check{m}^i, \dots, \check{m}^i) \text{ with probability 1} \\
\text{if } a_{-k}^t = a_{-k}(\hat{\ell}) \text{ and } a_k^t \neq a_k(\hat{\ell}) & \text{then } m_{-i}^{t+1} = (\underline{m}^{k,T}, \dots, \underline{m}^{k,T}) \text{ with probability 1} \\
\text{otherwise} & m_{-i}^{t+1} = (m^*, \dots, m^*) \text{ with probability 1}
\end{array} \tag{T.9.4}$$

The next subcase is that of $m_i^t = \check{m}^j$. Let $x(\dots, j_\ell, \dots)$ denote the realization of x^t . With the understanding that j' is an element of I not equal to i and that $\underline{m}^{j',\tau}$ is a generic element of $\underline{M}(j', t+1)$, let $\Phi_i^{tE}(\check{m}^j, x(\dots, j_\ell, \dots), a^t, y^t)$ be as follows

$$\begin{array}{ll}
\text{if } a^t = \check{a}^j(j_\ell) \text{ and } y^t = y(j') & \text{then } \begin{cases} m_{-i-j}^{t+1} = (\check{m}^{j'}, \dots, \check{m}^{j'}) \\ m_{j'}^{t+1} = \underline{m}^{j',\tau} \end{cases} \text{ with pr. } \frac{1}{\|\underline{M}(j', t+1)\|} \\
\text{if } a^t = \check{a}^j(j_\ell) \text{ and } y^t = y(i) & \text{then } m_{-i}^{t+1} = (\check{m}^i, \dots, \check{m}^i) \text{ with probability 1} \\
\text{if } a_{-k}^t = \check{a}_{-k}^j(j_\ell) \text{ and } a_k^t \neq \check{a}_k^j(j_\ell) & \text{then } m_{-i}^{t+1} = (\underline{m}^{k,T}, \dots, \underline{m}^{k,T}) \text{ with probability 1} \\
\text{otherwise} & m_{-i}^{t+1} = (m^*, \dots, m^*) \text{ with probability 1}
\end{array} \tag{T.9.5}$$

The next subcase is that of $m_i^t = \underline{m}^{i,\tau} \in \underline{M}(i, t)$. Let $x(\dots, i_\ell, \dots)$ denote the realization of x^t . With the understanding that $\underline{m}^{j,\tau}$ is a generic element of $\underline{M}(j, t+1)$, let $\Phi_i^{tE}(\underline{m}^{i,\tau}, x(\dots, i_\ell, \dots), a^t, y^t)$ be as follows

$$\begin{array}{ll}
\text{if } a^t = \check{a}^i(i_\ell) \text{ and } y^t = y(j) & \text{then } \begin{cases} m_{-i-j}^{t+1} = (\check{m}^j, \dots, \check{m}^j) \\ m_j^{t+1} = \underline{m}^{j,\tau} \end{cases} \text{ with pr. } \frac{1}{\|\underline{M}(j, t+1)\|} \\
\text{if } a^t = \check{a}^i(i_\ell) \text{ and } y^t = y(i) & \text{then } m_{-i}^{t+1} = (\check{m}^i, \dots, \check{m}^i) \text{ with probability 1} \\
\text{if } a_{-k}^t = \check{a}_{-k}^i(i_\ell) \text{ and } a_k^t \neq \check{a}_k^i(i_\ell) & \text{then } m_{-i}^{t+1} = (\underline{m}^{k,T}, \dots, \underline{m}^{k,T}) \text{ with probability 1} \\
\text{if } a_{-k}^t = a_{-k}(i_\ell) \text{ and } a_k^t \neq a_k(i_\ell) & \text{then } m_{-i}^{t+1} = (\underline{m}^{k,T}, \dots, \underline{m}^{k,T}) \text{ with probability 1} \\
\text{if } a^t = a(i_\ell) & \text{then } m_{-i}^{t+1} = (\underline{m}^{i,\tau-1}, \dots, \underline{m}^{i,\tau-1}) \text{ with probability 1} \\
\text{otherwise} & m_{-i}^{t+1} = (m^*, \dots, m^*) \text{ with probability 1}
\end{array} \tag{T.9.6}$$

where we set $\underline{m}^{i,0} = \bar{m}^i$.

The next subcase of $t \geq 1$ and $x^t = x(\kappa)$ with $\kappa \leq \bar{\kappa}$ that we consider is that of $m_i^t = \underline{m}^{j,\tau} \in \underline{M}(j, t)$. Let $x(\dots, j_\ell, \dots)$ denote the realization of x^t . Let $\Phi_i^{tE}(\underline{m}^{j,\tau}, x(\dots, j_\ell, \dots), a^t, y^t)$ be as follows

$$\begin{array}{ll}
\text{if } a^t = a(j_\ell) & \text{then } m_{-i}^{t+1} = (\underline{m}^{j,\tau-1}, \dots, \underline{m}^{j,\tau-1}) \text{ with probability 1} \\
\text{if } a_{-k}^t = a_{-k}(j_\ell) \text{ and } a_k^t \neq a_k(j_\ell) & \text{then } m_{-i}^{t+1} = (\underline{m}^{k,T}, \dots, \underline{m}^{k,T}) \text{ with probability 1} \\
\text{otherwise} & m_{-i}^{t+1} = (m^*, \dots, m^*) \text{ with probability 1}
\end{array} \tag{T.9.7}$$

where we set $\underline{m}^{j,0} = \bar{m}^j$.

The final subcase to consider is that of $m_i^t = \bar{m}^{k'}$ for some $k' \in I$. Let $x(\dots, \ell_{k'}, \dots)$ denote the realization of x^t . Let $\Phi_i^{tE}(\bar{m}^{k'}, x(\dots, \ell_{k'}, \dots), a^t, y^t)$ be as follows

$$\begin{array}{ll}
\text{if } a^t = a(\ell_{k'}) & \text{then } m_{-i}^{t+1} = (\bar{m}^{k'}, \dots, \bar{m}^{k'}) \text{ with probability 1} \\
\text{if } a_{-k}^t = a_{-k}(\ell_{k'}) \text{ and } a_k^t \neq a_k(\ell_{k'}) & \text{then } m_{-i}^{t+1} = (\underline{m}^{k,T}, \dots, \underline{m}^{k,T}) \text{ with probability 1} \\
\text{otherwise} & m_{-i}^{t+1} = (m^*, \dots, m^*) \text{ with probability 1}
\end{array} \tag{T.9.8}$$

T.10. Proof of Theorem 3: Sequential Rationality

Definition T.10.1: Let \mathcal{I}_i^{tE} denote the end-of-period- t collection of information sets that belong to player $\langle i, t \rangle$, with typical element \mathcal{I}_i^{tE} .

It is convenient to partition \mathcal{I}_i^{tE} into mutually disjoint exhaustive subsets on the basis of the associated beliefs of player $\langle i, t \rangle$. The fact that they exhaust \mathcal{I}_i^{tE} can be checked directly from Definition T.9.2 above.

Let $\mathcal{I}_i^{tE}(\ast) \subset \mathcal{I}_i^{tE}$ be the collection of information sets in which player $\langle i, t \rangle$ believes that m_{-i}^{t+1} is equal to (m^\ast, \dots, m^\ast) with probability one. These beliefs will be denoted by $\Phi_i^{tE}(\ast)$.

Let $\mathcal{I}_i^{tE}(\surd i) \subset \mathcal{I}_i^{tE}$ be the collection of information sets in which player $\langle i, t \rangle$ believes that m_{-i}^{t+1} is equal to $(\check{m}^i, \dots, \check{m}^i)$ with probability one. These beliefs will be denoted by $\Phi_i^{tE}(\surd i)$.

For every $j \in I$ not equal to i , let $\mathcal{I}_i^{tE}(\surd j, t) \subset \mathcal{I}_i^{tE}$ be the collection of information sets in which player $\langle i, t \rangle$ believes that m_{-i-j}^{t+1} is equal to $(\check{m}^j, \dots, \check{m}^j)$ with probability one, that $\Pr(m_j^{t+1} = \underline{m}^{j,\tau}) > 0 \forall \underline{m}^{j,\tau} \in \underline{M}(j, t)$, and that $\Pr(m_j^{t+1} \in \underline{M}(j, t)) = 1$.^{T.15} These beliefs will be denoted by $\Phi_i^{tE}(\surd j, t)$.

For every $j \in I$ not equal to i , let $\mathcal{I}_i^{tE}(\surd j, t+1) \subset \mathcal{I}_i^{tE}$ be the collection of information sets in which player $\langle i, t \rangle$ believes that m_{-i-j}^{t+1} is equal to $(\check{m}^j, \dots, \check{m}^j)$ with probability one, that $\Pr(m_j^{t+1} = \underline{m}^{j,\tau}) = \|\underline{M}(j, t+1)\|^{-1} \forall \underline{m}^{j,\tau} \in \underline{M}(j, t+1)$. These beliefs will be denoted by $\Phi_i^{tE}(\surd j, t+1)$.

For every $k \in I$, let $\mathcal{I}_i^{tE}(\bar{k}) \subset \mathcal{I}_i^{tE}$ be the collection of information sets in which player $\langle i, t \rangle$ believes that m_{-i}^{t+1} is equal to $(\bar{m}^k, \dots, \bar{m}^k)$ with probability one. These beliefs will be denoted by $\Phi_i^{tE}(\bar{k})$.

For every $k \in I$, and every $\tau = \max\{T-t, 1\}, \dots, T$ let $\mathcal{I}_i^{tE}(\underline{k}, \tau) \subset \mathcal{I}_i^{tE}$ be the collection of information sets in which player $\langle i, t \rangle$ believes that m_{-i}^{t+1} is equal to $(\underline{m}^{k,\tau}, \dots, \underline{m}^{k,\tau})$ with probability one. These beliefs will be denoted by $\Phi_i^{tE}(\underline{k}, \tau)$.

Definition T.10.2: Let the strategy profile (g, μ) described in Definitions B.8 and B.9 be given. Fix a period t and an n -tuple of messages $m^{t+1} = (m_1^{t+1}, \dots, m_n^{t+1})$, with $m_k^{t+1} \in M_k^{t+1}$ for every $k \in I$.

Clearly, the profile (g, μ) together with m^{t+1} uniquely determine a probability distribution over action profiles over all future periods, beginning with $t+1$.

Therefore, we can define the expected discounted (from the beginning of period $t+1$) payoff to player $\langle i, t \rangle$, given (g, μ) and m^{t+1} in the obvious way. This will be denoted by $\ddot{v}_i^t(m^{t+1})$. Moreover, since they play a special role in some of the computations that follow, we reserve two pieces of notation for two particular instances of m^{t+1} . The expression $\ddot{v}_i^t(\ast)$ stands for $\ddot{v}_i^t(m^{t+1})$ when $m^{t+1} = (m^\ast, \dots, m^\ast)$. Moreover, for any $k \in I$, the expression $\ddot{v}_i^t(k, \tau)$ stands for $\ddot{v}_i^t(m^{t+1})$ when $m_{-k}^{t+1} = (\check{m}^k, \dots, \check{m}^k)$ and $m_k^{t+1} = \underline{m}^{k,\tau} \in \underline{M}(k, t+1)$.

Lemma T.10.1: For any $i \in I$, any $k \in I$, any t , and any $\tau = \max\{T-t, 1\}, \dots, T$, we have that

$$\ddot{v}_i^t(\ast) = \frac{(1-\delta) \left[q \hat{v}_i + (1-q) z_i \right] + \delta q v_i^\ast}{1 - \delta(1-q)} \quad (\text{T.10.1})$$

and

$$\ddot{v}_i^t(k, \tau) = \frac{(1-\delta) \left[q \check{u}_i^k + (1-q) z_i \right] + \delta q v_i^\ast}{1 - \delta(1-q)} \quad (\text{T.10.2})$$

where $\ddot{v}_i^t(\ast)$ and $\ddot{v}_i^t(k, \tau)$ are as in Definition T.10.2, \hat{v}_i is as in (B.6), z_i is as in Remark B.4, v_i^\ast is as in the statement of the Theorem, and \check{u}_i^k is as in (B.3).

^{T.15}See footnote T.14 above.

Proof: Assume first that $t \geq T$. Using Definitions B.8 and B.9 we can write $\check{v}_i^t(*)$ and $\check{v}_i^t(k, \tau)$ recursively as

$$\check{v}_i^t(*) = q \left\{ (1 - \delta)\hat{v}_i + \delta \left[(1 - \eta)\check{v}_i^{t+1}(*) + \frac{\eta}{n} \sum_{k' \in I} \sum_{\tau=1}^T \frac{\check{v}_i^{t+1}(k', \tau)}{T} \right] \right\} + \frac{1}{(1 - q)} [(1 - \delta)z_i + \delta\check{v}_i^{t+1}(*)] \quad (\text{T.10.3})$$

and

$$\check{v}_i^t(k, \tau) = q \left\{ (1 - \delta)\check{v}_i^k + \delta \left[(1 - \eta)\check{v}_i^{t+1}(*) + \frac{\eta}{n} \sum_{k' \in I} \sum_{\tau=1}^T \frac{\check{v}_i^{t+1}(k', \tau)}{T} \right] \right\} + \frac{1}{(1 - q)} [(1 - \delta)z_i + \delta\check{v}_i^{t+1}(k, \tau)] \quad (\text{T.10.4})$$

Since the strategy profile (g, μ) described in Definitions B.8 and B.9 is stationary for $t \geq T$, we immediately have that $\check{v}_i^t(*) = \check{v}_i^{t+1}(*)$ and, for any $k \in I$ and any $\tau = 1, \dots, T$, $\check{v}_i^t(k, \tau) = \check{v}_i^{t+1}(k, \tau)$. Hence we can solve (T.10.3) and (T.10.4) simultaneously for the $NT + 1$ variables $\check{v}_i^t(*)$ and $\check{v}_i^t(k, \tau)$ ($k \in I$ and $\tau = 1, \dots, T$). Using (B.7) this immediately gives (T.10.1) and (T.10.2), as required.

Proceeding by induction backwards from $t = T$, it is also immediate to verify that the statement holds for any $t < T$. The details are omitted for the sake of brevity. ■

Lemma T.10.2: *Let the strategy profile (g, μ) and system of beliefs Φ described in Definitions B.8, B.9, T.9.1 and T.9.2 be given. Then the end-of-period continuation payoffs for any player $\langle i, t \rangle$ (discounted as of the beginning of period $t + 1$) at any information set $\mathcal{I}_i^t \in \mathcal{I}_i^{tE}$ (as categorized in Definition T.10.1) are as follows.^{T.16}*

$$v_i^t(g, \mu | \Phi_i^{tE}(*)) = \frac{(1 - \delta) [q\hat{v}_i + (1 - q)z_i] + \delta q v_i^*}{1 - \delta(1 - q)} \quad (\text{T.10.5})$$

$$v_i^t(g, \mu | \Phi_i^{tE}(\surd i)) = \frac{(1 - \delta) [q\check{v}_i^i + (1 - q)z_i] + \delta q v_i^*}{1 - \delta(1 - q)} \quad (\text{T.10.6})$$

$$v_i^t(g, \mu | \Phi_i^{tE}(\surd j, t)) = v_i^t(g, \mu | \Phi_i^{tE}(\surd j, t + 1)) = \frac{(1 - \delta) [q\check{v}_i^j + (1 - q)z_i] + \delta q v_i^*}{1 - \delta(1 - q)} \quad \forall j \neq i \quad (\text{T.10.7})$$

$$v_i^t(g, \mu | \Phi_i^{tE}(\bar{k})) = q\bar{v}_i^k + (1 - q)z_i \quad \forall k \in I \quad (\text{T.10.8})$$

$$v_i^t(g, \mu | \Phi_i^{tE}(\underline{k}, \tau)) = \left[1 - \left(\frac{\delta q}{1 - \delta(1 - q)} \right)^\tau \right] [q\omega_i^k + (1 - q)z_i] + \left(\frac{\delta q}{1 - \delta(1 - q)} \right)^\tau [q\bar{v}_i^k + (1 - q)z_i] \quad \forall k \in I \quad \forall \tau = \max\{T - t, 1\}, \dots, T \quad (\text{T.10.9})$$

where \hat{v}_i is as in (B.6), z_i is as in Remark B.4, v_i^* is as in the statement of the Theorem, \check{v}_i^k is as in (B.3), and ω_i^k is as in (B.2).

^{T.16}See our Point of Notation T.1.1 above.

Proof: Equations (T.10.5), (T.10.6) and (T.10.7) are a direct consequence of Definition T.10.1 and Lemma T.10.1.

Equation (T.10.8) follows directly from Definition T.10.1 and the description of the profile (g, μ) in Definitions B.8 and B.9.

Using the notation established in Definition T.10.2, consider the quantity $\ddot{v}_i^t(\underline{m}^{k,\tau}, \dots, \underline{m}^{k,\tau})$. Given the strategies described in Definitions B.8 and B.9 it is evident that this quantity does not depend on t . Therefore, for any $k \in I$ and $\tau = \max\{T - t, 1\}, \dots, T$, we can let $\ddot{v}_i(\underline{k}, \tau) = \ddot{v}_i^t(\underline{m}^{k,\tau}, \dots, \underline{m}^{k,\tau})$, for all t . Clearly, using Definition T.10.1, we have that for all k, τ and t , $v_i^t(g, \mu | \Phi_i^{tE}(\underline{k}, \tau)) = \ddot{v}_i(\underline{k}, \tau)$.

From the description of (g, μ) in Definitions B.8 and B.9, for any $k \in I$ and for any $\tau = 2, \dots, T$, the quantity $\ddot{v}_i(\underline{k}, \tau)$ obeys a difference equation as follows.

$$\ddot{v}_i(\underline{k}, \tau) = q [(1 - \delta)\underline{\omega}_i^k + \delta\ddot{v}_i(\underline{k}, \tau - 1)] + (1 - q) [(1 - \delta)z_i + \delta\ddot{v}_i(\underline{k}, \tau)] \quad (\text{T.10.10})$$

Using again Definitions B.8 and B.9, the terminal condition for (T.10.10) is

$$\ddot{v}_i(\underline{k}, 1) = q [(1 - \delta)\underline{\omega}_i^k + \delta[q\underline{v}_i^k + (1 - q)z_i]] + (1 - q) [(1 - \delta)z_i + \delta\ddot{v}_i(\underline{k}, 1)] \quad (\text{T.10.11})$$

Solving (T.10.10) and imposing the terminal condition (T.10.11) now yields (T.10.9), as required. ■

Purely for expositional convenience, before completing the proof of sequential rationality at the message stage, we now proceed with the argument that establishes sequential rationality at the action stage.

Definition T.10.3: Recall that at the action stage, player $\langle i, t \rangle$ chooses an action after having received a message m_i^t and having observed a realization x^t of the randomization device \tilde{x}^t .

Let \mathcal{I}_i^{tB} denote period- t action-stage collection of information sets that belong to player $\langle i, t \rangle$, with typical element \mathcal{I}_i^{tB} . Clearly, each element of \mathcal{I}_i^{tB} is identified by a pair (m_i^t, x^t) .

It is convenient to partition \mathcal{I}_i^{tB} into mutually disjoint exhaustive subsets. The fact that they exhaust \mathcal{I}_i^{tB} can be checked directly from Definition T.9.1 above.

Let $\mathcal{I}_i^{tB}(\ast) \subset \mathcal{I}_i^{tB}$ be the collection of information sets in which player $\langle i, t \rangle$ believes that m_{-i}^t is equal to (m^\ast, \dots, m^\ast) with probability one.^{T.17} These beliefs will be denoted by $\Phi_i^{tB}(\ast)$.

Let $\mathcal{I}_i^{tB}(\surd i) \subset \mathcal{I}_i^{tB}$ be the collection of information sets in which player $\langle i, t \rangle$ believes that m_{-i}^t is equal to $(\check{m}^i, \dots, \check{m}^i)$ with probability one. These beliefs will be denoted by $\Phi_i^{tB}(\surd i)$.

For every $j \in I$ not equal to i , let $\mathcal{I}_i^{tB}(\surd j) \subset \mathcal{I}_i^{tB}$ be the collection of information sets in which player $\langle i, t \rangle$ believes that m_{-i-j}^t is equal to $(\check{m}^j, \dots, \check{m}^j)$ with probability one, that $\Pr(m_j^t = \underline{m}^{j,\tau}) > 0 \forall \underline{m}^{j,\tau} \in \underline{M}(j, t)$, and that $\Pr(m_j^t \in \underline{M}(j, t)) = 1$.^{T.18} These beliefs will be denoted by $\Phi_i^{tB}(\surd j)$.

For every $j \in I$ not equal to i , and every $\tau = \max\{T - t + 1, 1\}, \dots, T$ let $\mathcal{I}_i^{tB}(j, \tau) \subset \mathcal{I}_i^{tB}$ be the collection of information sets in which player $\langle i, t \rangle$ believes that m_{-i}^t is equal to $(\underline{m}^{j,\tau}, \dots, \underline{m}^{j,\tau})$ with probability one. These beliefs will be denoted by $\Phi_i^{tB}(j, \tau)$.

For every $k \in I$, let $\mathcal{I}_i^{tB}(\bar{k}) \subset \mathcal{I}_i^{tB}$ be the collection of information sets in which player $\langle i, t \rangle$ believes that m_{-i}^t is equal to $(\bar{m}^k, \dots, \bar{m}^k)$ with probability one. These beliefs will be denoted by $\Phi_i^{tE}(\bar{k})$.

Lemma T.10.3: There exists a $\underline{\delta} \in (0, 1)$ such that whenever $\delta > \underline{\delta}$ the action-stage strategies described in Definition B.8 are sequentially rational given the beliefs described in Definition T.9.1 for every player $\langle i, t \rangle$.^{T.19}

^{T.17}In the interest of brevity, we avoid an explicit distinction between the $t = 0$ players and all others. What follows can be interpreted as applying to all players re-defining m_i^0 to be equal to m^\ast for players $\langle i \in I, 0 \rangle$.

^{T.18}See footnote T.13.

^{T.19}It should be understood that we are, for now, taking it as given that each player $\langle i, t \rangle$ follows the prescriptions of the

Proof: Consider any information set $\mathcal{I}_i^{tB} \in \{\mathcal{I}_i^{tB}(\ast) \cup \mathcal{I}_i^{tB}(\surd i) \cup \mathcal{I}_i^{tB}(\surd j)\}$.^{T.20}

Using Definition B.8, Lemma T.10.2 and Definition T.10.3, it is immediate to check that, as $\delta \rightarrow 1$, the limit expected continuation payoff to player $\langle i, t \rangle$ from following the action-stage strategies described in Definition B.8 at any of these information sets is

$$v_i^* = q\hat{v}_i + (1 - q)z_i \quad (\text{T.10.12})$$

In the same way, it can be checked that, as $\delta \rightarrow 1$, the limit expected continuation payoff to player $\langle i, t \rangle$ from deviating at any of these information sets is

$$q\bar{v}_i^i + (1 - q)z_i \quad (\text{T.10.13})$$

Since by assumption $\hat{v}_i > \bar{v}_i^i$ this is of course sufficient to prove our claim for any information set $\mathcal{I}_i^{tB} \in \{\mathcal{I}_i^{tB}(\ast) \cup \mathcal{I}_i^{tB}(\surd i) \cup \mathcal{I}_i^{tB}(\surd j)\}$.

Now consider any information set \mathcal{I}_i^{tB} either in $\mathcal{I}_i^{tB}(\underline{j}, \tau)$ or in $\mathcal{I}_i^{tB}(\bar{j})$ (with $j \neq i$).

Using Definition B.8, Lemma T.10.2 and Definition T.10.3, it is immediate to check that, as $\delta \rightarrow 1$, the limit expected continuation payoff to player $\langle i, t \rangle$ from following the action-stage strategies described in Definition B.8 at any of these information sets is

$$q\bar{v}_i^j + (1 - q)z_i \quad (\text{T.10.14})$$

In the same way, it can be checked that, as $\delta \rightarrow 1$, the limit expected continuation payoff to player $\langle i, t \rangle$ from deviating at any of these information sets is exactly as in (T.10.13).

Since by assumption for any $j \neq i$ we have that $\bar{v}_i^j > \bar{v}_i^i$ this is of course sufficient to prove our claim for any of these information sets.

To conclude the proof of the lemma, we now consider any information set $\mathcal{I}_i^{tB} \in \mathcal{I}_i^{tB}(\bar{i})$. Using Definition B.8, Lemma T.10.2 and Definition T.10.3, it can be checked that the expected continuation payoff to player $\langle i, t \rangle$ from following the action-stage strategies described in Definition B.8 at any of these information sets is bounded below by

$$(1 - \delta)\underline{u}_i + \delta [q\bar{v}_i^i + (1 - q)z_i] \quad (\text{T.10.15})$$

In the same way it can be readily seen that the expected continuation payoff to player $\langle i, t \rangle$ from deviating at any of these information sets is bounded above by

$$(1 - \delta)\bar{u}_i + \delta \left\{ \left[1 - \left(\frac{\delta q}{1 - \delta(1 - q)} \right)^T \right] [q\omega_i^i + (1 - q)z_i] + \left(\frac{\delta q}{1 - \delta(1 - q)} \right)^T [q\bar{v}_i^i + (1 - q)z_i] \right\} \quad (\text{T.10.16})$$

message-stage strategies described in Definition B.9. Of course, we have not demonstrated yet that this is in fact sequentially rational given the beliefs described in Definition T.9.2. We will come back to this immediately after the current lemma is proved.
^{T.20}See Definition T.10.3.

The difference given by (T.10.15) minus (T.10.16) can be written as

$$(1 - \delta) \left\{ \frac{\delta q \left[1 - \left(\frac{\delta q}{1 - \delta(1 - q)} \right)^T \right] (\bar{v}_i^i - \underline{\omega}_i^i)}{(1 - \delta)} - (\bar{u}_i - \underline{u}_i) \right\} \quad (\text{T.10.17})$$

Consider now the term inside the curly brackets in (T.10.17). We have that

$$\lim_{\delta \rightarrow 1} \frac{\delta q \left[1 - \left(\frac{\delta q}{1 - \delta(1 - q)} \right)^T \right] (\bar{v}_i^i - \underline{\omega}_i^i)}{(1 - \delta)} - (\bar{u}_i - \underline{u}_i) = T(\bar{v}_i^i - \underline{\omega}_i^i) - (\bar{u}_i - \underline{u}_i) \quad (\text{T.10.18})$$

Using (B.11), we know that the quantity on the right-hand side of (T.10.18) is strictly positive. Hence we can conclude our claim is valid at any information set $\mathcal{I}_i^{tB} \in \mathcal{I}_i^{tB}(\bar{i})$. ■

Lemma T.10.4: *Consider the notation we established in Definition T.10.2. For any given t and $\tau = \max\{T - t, 1\}, \dots, T$ let $\ddot{v}_i^t(m, \underline{m}^{i,\tau})$ denote $\ddot{v}_i^t(m^{t+1})$ when the vector m^{t+1} has the i -th component equal to a generic $m \in M_i^{t+1}$ and $m_{-i}^{t+1} = (\underline{m}^{i,\tau}, \dots, \underline{m}^{i,\tau})$. As in the proof of Lemma T.10.2, let $\ddot{v}_i(\underline{i}, \tau) = \ddot{v}_i^t(\underline{m}^{i,\tau}, \dots, \underline{m}^{i,\tau})$.*

Then there exists a $\underline{\delta} \in (0, 1)$ such that whenever $\delta > \underline{\delta}$ for every player $\langle i, t \rangle$, for every $m \in M_i^{t+1}$, and for every $\tau = \max\{T - t, 1\}, \dots, T$

$$\ddot{v}_i(\underline{i}, \tau) \geq \ddot{v}_i^t(m, \underline{m}^{i,\tau}) \quad (\text{T.10.19})$$

Proof: We prove the claim for the case $t \geq T$. The treatment of $t < T$ has some completely non-essential complications due to the fact that the players’ message spaces increase in size for the first T periods. The details are omitted for the sake of brevity.

We now introduce a new random random variable \tilde{w} , independent of \tilde{x} and \tilde{y} (see Definitions B.4 and B.5), and uniformly distributed over the finite set $\{1, \dots, T\}$. This will be used in the rest of the proof of the lemma to keep track of the ‘‘private’’ randomization across messages that members of dynasty i may be required to perform (see Definition B.9). Just as we did for the action-stage and the message-stage randomization devices, we consider countably many independent ‘‘copies’’ of \tilde{w} , one for each time period, denoted by \tilde{w}^t , with typical realization w^t .

To keep track of all ‘‘future randomness’’ looking ahead for $t' = 1, 2, \dots$ periods from t , it will also be convenient to define the random vectors $\tilde{s}^{t,t'}$

$$\tilde{s}^{t,t'} = [(\tilde{x}^{t+1}, \tilde{y}^{t+1}, \tilde{w}^{t+1}), \dots, (\tilde{x}^{t+t'}, \tilde{y}^{t+t'}, \tilde{w}^{t+t'})] \quad (\text{T.10.20})$$

A typical realization of $\tilde{s}^{t,t'}$ will be denoted by $s^{t,t'} = [(x^{t+1}, y^{t+1}, w^{t+1}), \dots, (x^{t+t'}, y^{t+t'}, w^{t+t'})]$. The set of all possible realizations of $\tilde{s}^{t,t'}$ (which obviously does not depend on t) is denoted by $S^{t'}$.

Recall that the profile (g, μ) described in Definitions B.8 and B.9 is taken as given throughout. Now suppose that in period t , player $\langle i, t \rangle$ sends a generic message $m \in M_i^{t+1}$ and that $m_{-i}^{t+1} = (\underline{m}^{i,\tau}, \dots, \underline{m}^{i,\tau})$. Then, given any realization $s^{t,t'}$ we can compute the actual action profile played by all players $\langle k \in I, t + t' \rangle$. This will be denoted by $\mathbf{a}^{t+t'}(m, \underline{m}^{i,\tau}, s^{t,t'})$. Similarly, we can compute the profile of messages $m_{-i}^{t+t'}$ received by all players $\langle j \neq i, t + t' \rangle$. This $n - 1$ -tuple will be denoted by $\mathbf{m}^{t+t'}(m, \underline{m}^{i,\tau}, s^{t,t'})$.

Recall that the messages received by all time- $t + t'$ players are the result of choices and random draws that take place on or before period $t + t' - 1$. Therefore it is clear that if we are given two realizations $s^{t,t'}$

$= [s^{t,t'-1}, (\hat{x}^{t+t'}, \hat{y}^{t+t'}, \hat{w}^{t+t'})]$ and $\hat{s}^{t,t'} = [s^{t,t'-1}, (\hat{x}^{t+t'}, \hat{y}^{t+t'}, \hat{w}^{t+t'})]$, then it must be that

$$\mathbf{m}^{t+t'}(m, \underline{m}^{i,\tau}, s^{t,t'}) = \mathbf{m}^{t+t'}(m, \underline{m}^{i,\tau}, \hat{s}^{t,t'}) \quad (\text{T.10.21})$$

Notice next that from the description of the profile (g, μ) in Definitions B.8 and B.9 it is also immediate to check that for any t' , any $m \in M_i^{t'+1}$ and any realization $s^{t,t'}$ the message profile $\mathbf{m}^{t+t'}(m, \underline{m}^{i,\tau}, s^{t,t'})$ can only take one out of two possible forms. Either we have $\mathbf{m}^{t+t'}(m, \underline{m}^{i,\tau}, s^{t,t'}) = (\bar{m}^i, \dots, \bar{m}^i)$ or it must be that $\mathbf{m}^{t+t'}(m, \underline{m}^{i,\tau}, s^{t,t'}) = (\underline{m}^{i,\tau'}, \dots, \underline{m}^{i,\tau'})$ for some $\tau' = 1, \dots, T$.

Lastly, notice that, given an arbitrary message $m \in M_i^{t'+1}$ we can write

$$\ddot{v}_i^t(m, \underline{m}^{i,\tau}) = (1 - \delta) \sum_{t'=1}^{\infty} \delta^{t'-1} \sum_{s^{t,t'} \in S^{t'}} \Pr(\tilde{s}^{t,t'} = s^{t,t'}) u_i[\mathbf{a}^{t+t'}(m, \underline{m}^{i,\tau}, s^{t,t'})] \quad (\text{T.10.22})$$

Since the strategies described in Definitions B.8 and B.9 are stationary for $t \geq T$, and the distribution of $\tilde{s}^{t,t'}$ is independent of t , it is evident from (T.10.22) that $\ddot{v}_i^t(m, \underline{m}^{i,\tau})$ does not depend on t . From now on we drop the superscript and write $\ddot{v}_i(m, \underline{m}^{i,\tau})$.

We now proceed with the proof of inequality (T.10.19) of the statement of the lemma. In order to do so, from now on we fix a particular $t = \hat{t}$, $m = \hat{m}$ and $\tau = \hat{\tau}$, and we prove (T.10.19) for these fixed values of t , m and τ . Since the lower bound on δ that we will find will clearly not depend on t , and since there are finitely many values that m and τ can take, this will be sufficient to prove the claim.

Inequality (T.10.19) in the statement of the lemma is trivially satisfied (as an equality) if $m = \underline{m}^{i,\tau}$. From now on assume that \hat{m} and $\hat{\tau}$ are such that $\hat{m} \neq \underline{m}^{i,\hat{\tau}}$.

Given any $t' = 1, 2, \dots$, we now partition the set of realizations $S^{t'}$ into five disjoint exhaustive subsets; $S_1^{t'}$, $S_2^{t'}$, $S_3^{t'}$, $S_4^{t'}$ and $S_5^{t'}$. This will allow us to decompose the right-hand side of (T.10.22) in a way that will make possible the comparison with (a similar decomposition of) the left-hand side of (T.10.19) as required to prove the lemma.

Let

$$S_1^{t'} = \{s^{\hat{t},t'} \mid \mathbf{m}^{\hat{t}+t'}(\underline{m}^{i,\hat{\tau}}, \underline{m}^{i,\hat{\tau}}, s^{\hat{t},t'}) = (\underline{m}^{i,\tau'}, \dots, \underline{m}^{i,\tau'}) \text{ for some } \tau' = 1, \dots, \hat{\tau}\} \quad (\text{T.10.23})$$

and notice that if $t' \leq \hat{\tau}$ then $S_1^{t'} = S^{t'}$.

Assume now that $t' > \hat{\tau}$ and let

$$S_2^{t'} = \{s^{\hat{t},t'} \mid \mathbf{m}^{\hat{t}+t'}(\hat{m}, \underline{m}^{i,\hat{\tau}}, s^{\hat{t},t'}) = (\bar{m}^i, \dots, \bar{m}^i) \text{ and } u_i(\mathbf{a}^{\hat{t}+t'}(\hat{m}, \underline{m}^{i,\hat{\tau}}, s^{\hat{t},t'})) \leq u_i(\mathbf{a}^{\hat{t}+t'}(\underline{m}^{i,\hat{\tau}}, \underline{m}^{i,\hat{\tau}}, s^{\hat{t},t'}))\} \quad (\text{T.10.24})$$

and

$$S_3^{t'} = \{s^{\hat{t},t'} \mid \mathbf{m}^{\hat{t}+t'}(\hat{m}, \underline{m}^{i,\hat{\tau}}, s^{\hat{t},t'}) = (\bar{m}^i, \dots, \bar{m}^i) \text{ and } u_i(\mathbf{a}^{\hat{t}+t'}(\hat{m}, \underline{m}^{i,\hat{\tau}}, s^{\hat{t},t'})) > u_i(\mathbf{a}^{\hat{t}+t'}(\underline{m}^{i,\hat{\tau}}, \underline{m}^{i,\hat{\tau}}, s^{\hat{t},t'}))\} \quad (\text{T.10.25})$$

Notice that if the first condition in (T.10.24) holds, then $\mathbf{m}^{\hat{t}+t'}(\underline{m}^{i,\hat{\tau}}, \underline{m}^{i,\hat{\tau}}, s^{\hat{t},t'}) = (\bar{m}^i, \dots, \bar{m}^i)$. Therefore, $S_1^{t'}$ and $S_2^{t'}$ and $S_3^{t'}$ are disjoint.

Next, let any $s^{\hat{t},t''} \in S_3^{t''}$ with $t'' < t'$ be given and define

$$S_4^{t'}(s^{\hat{t},t''}) = \{s^{\hat{t},t'} \mid s^{\hat{t},t'} = (s^{\hat{t},t''}, s^{t'',t'}) \text{ for some } s^{t'',t'} \text{ and } \|\{t \in (t'' + 1, \dots, t' - 1) \mid x^t = x(\kappa) \text{ with } \kappa \leq \bar{\kappa}\}\| \leq T - 1\} \quad (\text{T.10.26})$$

Now let

$$S_4^{t'} = \bigcup_{\substack{t'' < t' \\ s^{\hat{i}, t''} \in S_3^{t''}}} S_4^{t''} \quad (\text{T.10.27})$$

From the strategies described in Definitions B.8 and B.9 it can be checked that if $s^{\hat{i}, t'} \in S_4^{t'}$ then $\mathbf{m}^{\hat{i}+t'}(\hat{m}, \underline{m}^{\hat{i}, \hat{\tau}}, s^{\hat{i}, t'}) = (\underline{m}^{\hat{i}, \tau'}, \dots, \underline{m}^{\hat{i}, \tau'})$ for some τ' and $\mathbf{m}^{\hat{i}+t'}(\underline{m}^{\hat{i}, \hat{\tau}}, \underline{m}^{\hat{i}, \hat{\tau}}, s^{\hat{i}, t'}) = (\overline{m}^{\hat{i}}, \dots, \overline{m}^{\hat{i}})$. Therefore, it is clear that $S_4^{t'}$ is disjoint from $S_1^{t'}$, $S_2^{t'}$ and $S_3^{t'}$.

The last set in the partition of $S^{t'}$ is defined as the residual of the previous four.

$$S_5^{t'} = S^{t'} / \{S_1^{t'} \cup S_2^{t'} \cup S_3^{t'} \cup S_4^{t'}\} \quad (\text{T.10.28})$$

Using (T.10.22), we can now proceed to compare the two sides of inequality (T.10.19) of the statement of the lemma for the five distinct (conditional) cases $s^{\hat{i}, t'} \in S_1^{t'}$ through $s^{\hat{i}, t'} \in S_5^{t'}$. Notice first of all that when $s^{\hat{i}, t'} \in S_2^{t'}$, we know immediately from (T.10.24) that there is nothing to prove.

We begin with $s^{\hat{i}, t'} \in S_1^{t'}$. Notice first of all that if we fix any $\bar{s}^{\hat{i}, t'} \in S_1^{t'}$, then it follows from (T.10.21) and (T.10.23) that any $s^{\hat{i}, t'}$ of the form $[\bar{s}^{\hat{i}, t'-1}, s^{t'-1, t'}]$ (where $\bar{s}^{\hat{i}, t'-1}$ are the first $t' - 1$ triples of $\bar{s}^{\hat{i}, t'}$) is in fact in $S_1^{t'}$.

Using, (T.10.23) and Definitions 3, B.8 and B.9 we get

$$\begin{aligned} \sum_{s^{t'-1, t'} \in S^1} \Pr(\bar{s}^{t'-1, t'} = s^{t'-1, t'}) u_i(\mathbf{a}^{\hat{i}+t'}(\underline{m}^{\hat{i}, \hat{\tau}}, \underline{m}^{\hat{i}, \hat{\tau}}, [\bar{s}^{\hat{i}, t'-1}, s^{t'-1, t'}])) &= q\omega_i^i + (1-q)z_i \geq \\ \sum_{s^{t'-1, t'} \in S^1} \Pr(\bar{s}^{t'-1, t'} = s^{t'-1, t'}) u_i(\mathbf{a}^{\hat{i}+t'}(\hat{m}, \underline{m}^{\hat{i}, \hat{\tau}}, [\bar{s}^{\hat{i}, t'-1}, s^{t'-1, t'}])) & \end{aligned} \quad (\text{T.10.29})$$

Therefore, since the $\bar{s}^{\hat{i}, t'}$ that we fixed is an arbitrary element of $S_1^{t'}$, we can now conclude that

$$\sum_{s^{\hat{i}, t'} \in S_1^{t'}} \Pr(\bar{s}^{\hat{i}, t'} = s^{\hat{i}, t'}) u_i(\mathbf{a}^{\hat{i}+t'}(\underline{m}^{\hat{i}, \hat{\tau}}, \underline{m}^{\hat{i}, \hat{\tau}}, s^{\hat{i}, t'})) \geq \sum_{s^{\hat{i}, t'} \in S_1^{t'}} \Pr(\bar{s}^{\hat{i}, t'} = s^{\hat{i}, t'}) u_i(\mathbf{a}^{\hat{i}+t'}(\hat{m}, \underline{m}^{\hat{i}, \hat{\tau}}, s^{\hat{i}, t'})) \quad (\text{T.10.30})$$

Now fix any $\bar{s}^{\hat{i}, t'} \in S_3^{t'}$. Using, (T.10.25), (T.10.26) and (T.10.27), and Definitions B.8 and B.9 we get that the difference given by

$$\begin{aligned} \Pr(\bar{s}^{\hat{i}, t'} = \bar{s}^{\hat{i}, t'}) u_i(\mathbf{a}^{\hat{i}+t'}(\underline{m}^{\hat{i}, \hat{\tau}}, \underline{m}^{\hat{i}, \hat{\tau}}, \bar{s}^{\hat{i}, t'})) + \\ \sum_{t''=t'+1}^{\infty} \delta^{(t''-t')} \sum_{s^{\hat{i}, t''} \in S_4^{t''}(\bar{s}^{\hat{i}, t'})} \Pr(\bar{s}^{\hat{i}, t''} = \bar{s}^{\hat{i}, t''}) u_i(\mathbf{a}^{\hat{i}+t''}(\underline{m}^{\hat{i}, \hat{\tau}}, \underline{m}^{\hat{i}, \hat{\tau}}, \bar{s}^{\hat{i}, t''})) \end{aligned} \quad (\text{T.10.31})$$

minus

$$\begin{aligned} \Pr(\bar{s}^{\hat{i}, t'} = \bar{s}^{\hat{i}, t'}) u_i(\mathbf{a}^{\hat{i}+t'}(\hat{m}, \underline{m}^{\hat{i}, \hat{\tau}}, \bar{s}^{\hat{i}, t'})) + \\ \sum_{t''=t'+1}^{\infty} \delta^{(t''-t')} \sum_{s^{\hat{i}, t''} \in S_4^{t''}(\bar{s}^{\hat{i}, t'})} \Pr(\bar{s}^{\hat{i}, t''} = \bar{s}^{\hat{i}, t''}) u_i(\mathbf{a}^{\hat{i}+t''}(\hat{m}, \underline{m}^{\hat{i}, \hat{\tau}}, \bar{s}^{\hat{i}, t''})) \end{aligned} \quad (\text{T.10.32})$$

is greater or equal to

$$\Pr(\bar{s}^{\hat{t},t'} = \bar{s}^{\hat{t},t'}) \left\{ \frac{\delta q \left[1 - \left(\frac{\delta q}{1 - \delta(1 - q)} \right)^T \right] (\bar{v}_i^i - \underline{\omega}_i^i)}{(1 - \delta)} - (\bar{u}_i - \underline{u}_i) \right\} \quad (\text{T.10.33})$$

Notice now that we know that the quantity in (T.10.33) is in fact positive for δ sufficiently close to 1. This is simply because the term in curly brackets in (T.10.33) is the same as the right-hand side of (T.10.18). Therefore, we have dealt with any $\bar{s}^{\hat{t},t'} \in S_3^{t'}$ and with all its relevant ‘‘successors’’ of the form $S_4^{t''}(\bar{s}^{\hat{t},t'})$. Since t' is arbitrary, by (T.10.27), this exhausts $S_3^{t'}$ and $S_4^{t'}$ for all possible values of t' .

Finally, we deal with $s^{\hat{t},t'} \in S_5^{t'}$. Notice first of all that if we fix any $\bar{s}^{\hat{t},t'} \in S_5^{t'}$, then it follows from (T.10.21) and (T.10.28) that any $s^{\hat{t},t'}$ of the form $[\bar{s}^{\hat{t},t'-1}, s^{t'-1,t'}]$ (where $\bar{s}^{\hat{t},t'-1}$ are the first $t' - 1$ triples of $\bar{s}^{\hat{t},t'}$) is in fact in $S_5^{t'}$.

Using, (T.10.28) and Definitions B.8 and B.9 we get

$$\begin{aligned} \sum_{s^{t'-1,t'} \in S^1} \Pr(\bar{s}^{t'-1,t'} = s^{t'-1,t'}) u_i(\mathbf{a}^{\hat{t}+t'}(\underline{m}^{i,\hat{\tau}}, \underline{m}^{i,\hat{\tau}}, [\bar{s}^{\hat{t},t'-1}, s^{t'-1,t'}])) &= q\bar{v}_i^i + (1 - q)z_i > \\ q\omega_i^i + (1 - q)z_i &\geq \sum_{s^{t'-1,t'} \in S^1} \Pr(\bar{s}^{t'-1,t'} = s^{t'-1,t'}) u_i(\mathbf{a}^{\hat{t}+t'}(\hat{m}, \underline{m}^{i,\hat{\tau}}, [\bar{s}^{\hat{t},t'-1}, s^{t'-1,t'}])) \end{aligned} \quad (\text{T.10.34})$$

Therefore, since the $\bar{s}^{\hat{t},t'}$ that we fixed is an arbitrary element of $S_5^{t'}$, we can now conclude that

$$\sum_{s^{\hat{t},t'} \in S_5^{t'}} \Pr(\bar{s}^{\hat{t},t'} = s^{\hat{t},t'}) u_i(\mathbf{a}^{\hat{t},t'}(\underline{m}^{i,\hat{\tau}}, \underline{m}^{i,\hat{\tau}}, s^{\hat{t},t'})) \geq \sum_{s^{\hat{t},t'} \in S_5^{t'}} \Pr(\bar{s}^{\hat{t},t'} = s^{\hat{t},t'}) u_i(\mathbf{a}^{\hat{t},t'}(\hat{m}, \underline{m}^{i,\hat{\tau}}, s^{\hat{t},t'})) \quad (\text{T.10.35})$$

Hence, the proof of the lemma is now complete. ■

Remark T.10.1: Let the strategy profile (g, μ) described in Definitions B.8 and B.9 be given. Consider a player $\langle i, t \rangle$, and a realization of future uncertainty $s^{t,t'}$ as defined in the proof of Lemma T.10.4.

Let any message $m \in M_i^{t+1}$ be given, and fix any information set \mathcal{I}_i^{tE} and associated beliefs $\Phi_i^{tE}(\cdot)$.

It is then clear from Definitions B.8 and B.9 and T.10.1, that for any t' the action that player $\langle i, t \rangle$ expects player $\langle i, t + t' \rangle$ to take is uniquely determined by m , $s^{t,t'}$ and \mathcal{I}_i^{tE} .

For the rest of the argument we will denote this by $\mathbf{a}_i^{t+t'}(m, s^{t,t'}, \mathcal{I}_i^{tE})$.

Lemma T.10.5: There exists a $\underline{\delta} \in (0, 1)$ such that whenever $\delta > \underline{\delta}$ the message-stage strategies described in Definition B.9 are sequentially rational given the beliefs described in Definition T.9.2 for every player $\langle i, t \rangle$.

Proof: Consider any information set $\mathcal{I}_i^{tE} \in \mathcal{I}_i^{tE}(\underline{z}, \tau)$, where $\mathcal{I}_i^{tE}(\underline{z}, \tau)$ is as in Definition T.10.1. It is then evident from Lemma T.10.4 and from the beliefs $\Phi_i^{tE}(\underline{z}, \tau)$ described in Definition T.10.1 that for δ sufficiently close to 1, the message strategies described in Definition B.9 are sequentially rational at any such information set.

From now on, consider any information set $\mathcal{I}_i^{tE} \notin \mathcal{I}_i^{tE}(\underline{z}, \tau)$. Let $m \in M_i^{t+1}$ be the message that player $\langle i, t \rangle$ should send according to the strategy μ_i^t , and let \hat{m} be any other message in M_i^{t+1} . Consider a particular realization $\bar{s}^{t,t'}$, and for any $t'' \in \{1, \dots, t' - 1\}$, let $\bar{s}^{t,t''}$ denote the first t'' triples of $\bar{s}^{t,t'}$.

Next, assume that $\mathbf{a}_i^{t+t'}(m, \bar{s}^{t,t'}, \mathcal{I}_i^{tE}) \neq \mathbf{a}_i^{t+t'}(\hat{m}, \bar{s}^{t,t'}, \mathcal{I}_i^{tE})$, and that either $t' = 1$, or alternatively that $\mathbf{a}_i^{t+t''}(m, \bar{s}^{t,t''}, \mathcal{I}_i^{tE}) = \mathbf{a}_i^{t+t''}(\hat{m}, \bar{s}^{t,t''}, \mathcal{I}_i^{tE})$ for every $t'' \in \{1, \dots, t' - 1\}$.

Clearly, in periods $\{t+1, \dots, t'-1\}$, conditional on $\bar{s}^{t,t'}$, the payoff to player $\langle i, t \rangle$ is unaffected by the deviation to \hat{m} . Now consider the payoff to player $\langle i, t \rangle$, conditional on $\bar{s}^{t,t'}$, from the beginning of period t' on, for simplicity discounted from the beginning of period t' . If player $\langle i, t \rangle$ sends message m as prescribed by μ_i^t , and δ is close enough to 1, the payoff in question is bounded below by

$$(1-\delta)\underline{u}_i + \delta(q\bar{v}_i^i + (1-q)z_i) \quad (\text{T.10.36})$$

Now consider the payoff to player $\langle i, t \rangle$ if he sends message \hat{m} , conditional on $\bar{s}^{t,t'}$, from the beginning of period t' on, for simplicity discounted from the beginning of period t' . In period t' the action played cannot yield him more than \bar{u}_i . From Lemma T.10.4, we know that, for δ close enough to 1, from the beginning of period $t'+1$ the payoff is bounded above by $\check{v}_i(i, T)$. Hence, for δ close enough to 1, using (T.10.9) the payoff in question is bounded above by

$$\delta\bar{u}_i + (1-\delta) \left\{ \left[1 - \left(\frac{\delta q}{1-\delta(1-q)} \right)^T \right] [q\omega_i^i + (1-q)z_i] + \left(\frac{\delta q}{1-\delta(1-q)} \right)^T [q\bar{v}_i^i + (1-q)z_i] \right\} \quad (\text{T.10.37})$$

Notice now that the quantity in (T.10.36) is the same as the quantity in (T.10.15), and the quantity in (T.10.37) is in fact the same as the quantity in (T.10.16). Hence, exactly as in the proof of Lemma T.10.3, we know that, for δ sufficiently close to 1, the quantity in (T.10.36) is greater than the quantity in (T.10.37). This is clearly enough to conclude the proof. ■

T.11. Proof of Theorem 3: Consistency of Beliefs

Remark T.11.1: Let $(g_\varepsilon, \mu_\varepsilon)$ be the completely mixed strategy profile of Definitions B.11 and B.12. It is then straightforward to check that as $\varepsilon \rightarrow 0$ the profile $(g_\varepsilon, \mu_\varepsilon)$ converges pointwise (in fact uniformly) to the equilibrium strategy profile described in Definitions B.8 and B.9, as required.

Lemma T.11.1: The strategy profile (g, μ) described in Definitions B.8 and B.9 and the beginning-of-period beliefs described in Definition T.9.1 are consistent.

Proof: When $t = 0$, there is nothing to prove. Assume $t \geq 1$. We consider two cases. First assume that player $\langle i, t \rangle$ receives message $m \in \{m^*\} \cup \check{M}_{-i} \cup \underline{M}(i, t)$. Clearly, this is on the equilibrium path generated by the profile of strategies (g, μ) described in Definitions B.8 and B.9. Therefore, consistency in this case simply requires checking that the beginning-of-period beliefs described in Definition T.9.1 are obtained via Bayes' rule from the profile (g, μ) . This is a routine exercise, and we omit the details.

Now assume that player $\langle i, t \rangle$ receives message $m \notin \{m^*\} \cup \check{M}_{-i} \cup \underline{M}(i, t)$. From Definition T.9.1 it is immediate to check that in this case player $\langle i, t \rangle$ assigns probability one to the event that $m_{-i}^t = (m, \dots, m)$. Given (g, μ) , this event may of course have been generated by several possible histories. Notice however, that the profile (g, μ) is such that a *single* deviation by one player at the action stage is sufficient to generate the message profile $m^t = (m, \dots, m)$. Therefore, upon observing $m \notin \{m^*\} \cup \check{M}_{-i} \cup \underline{M}(i, t)$ the probability that $m_{-i}^t = (m, \dots, m)$ is an infinitesimal in ε of order no higher than 2^{T-21} . This needs to be compared with the probability that $m_{-i}^t \neq (m, \dots, m)$ and $m_i^t = m$. The latter event is impossible given the profile (g, μ) unless a deviation at the message stage has occurred at some point. Therefore its probability is an infinitesimal in ε of order no lower than $2n+1$. This is obviously enough to prove the claim. ■

Lemma T.11.2: The strategy profile (g, μ) described in Definitions B.8 and B.9 and the end-of-period beliefs described in Definition T.9.2 are consistent.

^{T.21}See footnote T.2 above for a specification of our (standard) use of terminology concerning the orders of infinitesimals.

Proof: The case $t = 0$ is trivial. Assume $t \geq 1$, and consider any player $\langle i, t \rangle$ after having observed (m_i^t, x^t, a^t, y^t) .

We deal first with the case in which $x^t = x(\kappa)$ with $\kappa > \bar{\kappa}$. Let $x(\ell_{00}, \ell^*)$ denote the realization x^t . In this case, the action-stage strategies described in Definition B.8 prescribe that every player $\langle k \in I, t \rangle$ should play $a_k^t(\ell^*)$. Therefore, if the observed action profile a^t is equal to $a(\ell^*)$, player $\langle i, t \rangle$ does not revise his beginning-of-period beliefs during period t . Hence consistency in this case follows immediately from the profile μ and from the consistency of beginning-of-period beliefs, which of course was proved in Lemma T.11.1. Notice now that if $a^t \neq a(\ell^*)$, then the message strategies described in Definition B.9 prescribe that each player $\langle k \in I, t \rangle$ should send a message that does not depend on the message m_k^t he received. Hence, in this case consistency is immediate from Definition T.9.2 and the profile μ .

We now turn to the case in which $x^t = x(\kappa)$ with $\kappa \leq \bar{\kappa}$. Here, it is necessary to consider several subcases, depending on the message m received by player $\langle i, t \rangle$. Assume first that $m \notin \check{M}_{-i} \cup \underline{M}(i, t)$. Then for any possible triple (x^t, a^t, y^t) we have that

$$\lim_{\varepsilon \rightarrow 0} \Pr(m_{-i}^t = (m, \dots, m) \mid m_i^t = m, x^t, a^t, g_\varepsilon, \mu_\varepsilon) = 1 \quad (\text{T.11.1})$$

To see this consider two sets of possibilities. First, $m = m^*$, $x^t = x(\cdot, \hat{\ell}, \dots)$, and $a^t = (a_1(\hat{\ell}), \dots, a_n(\hat{\ell}))$. Then play is as prescribed by the equilibrium path generated by the profile (g, μ) , and from Definitions B.8 and B.9 there is nothing more to prove. For all other possibilities, notice that the event $m^t = (m, \dots, m)$ is consistent with any a^t together with n deviations at the action stage of the second type described in Definition B.11. Therefore, for any a^t , the probability of $m^t = (m, \dots, m)$ and a^t is an infinitesimal in ε of order no higher than $2n$. On the other hand, from Definition B.12 it is immediate that the probability that $m_{-i}^t \neq (m, \dots, m)$ (since it requires at least one deviation at the message stage) is an infinitesimal in ε of order no lower than $2n + 1$. Hence (T.11.1) follows. From (T.11.1) it is a matter of routine to check the consistency of end-of-period beliefs from using the profile (g, μ) . We omit the details.

Still assuming that $x^t = x(\kappa)$ with $\kappa \leq \bar{\kappa}$, now consider the case $m = \check{m}_j \in \check{M}_{-i}$. In this case we can show that

$$\lim_{\varepsilon \rightarrow 0} \Pr(m_{-i-j}^t = (\check{m}_j, \dots, \check{m}_j) \text{ and } m_j^t \in \underline{M}(j, t) \mid m_i^t = \check{m}_j, x^t, a^t, g_\varepsilon, \mu_\varepsilon) = 1 \quad (\text{T.11.2})$$

using an argument completely analogous to the one we used for (T.11.1). The details are omitted. As in the previous case, from (T.11.2) it is a matter of routine to check the consistency of end-of-period beliefs from using the profile (g, μ) .

The last case remaining is $x^t = x(\kappa)$ with $\kappa \leq \bar{\kappa}$ and $m = \underline{m}^{i, \tau}$. In this case we have that

$$\begin{aligned} \lim_{\varepsilon \rightarrow 0} \Pr(m_{-i}^t = (\check{m}_i, \dots, \check{m}_i) \mid m_i^t = \underline{m}^{i, \tau}, x^t, a^t, g_\varepsilon, \mu_\varepsilon) + \\ \lim_{\varepsilon \rightarrow 0} \Pr(m_{-i}^t = (\underline{m}^{i, \tau}, \dots, \underline{m}^{i, \tau}) \mid m_i^t = \underline{m}^{i, \tau}, x^t, a^t, g_\varepsilon, \mu_\varepsilon) = 1 \end{aligned} \quad (\text{T.11.3})$$

Again, the argument is completely analogous to the one used for (T.11.1) and (T.11.2), and the details are omitted. Now take (T.11.3) as given and let $x^t = x(\dots, i_\ell, \dots)$.

Suppose next that $a_{-i}^t = \check{a}_{-i}^i(i_\ell)$. Then player $\langle i, t \rangle$ does not revise his beginning-of-period beliefs, and hence, using the profile μ and Lemma T.11.1 it is immediate to check that his end-of-period beliefs are consistent in this case.

Now suppose that for some $j \neq i$ we have that $a_j^t \neq \check{a}_j^i(i_\ell)$ and $a_{-i-j}^t = \check{a}_{-i-j}^i(i_\ell)$. Consistency of beliefs in this case requires showing that the first element in the sum in (T.11.3) is equal to 1. Of course given (T.11.3) it suffices to compare the probabilities of the two events $m_{-i}^t = (\check{m}_i, \dots, \check{m}_i)$ and $m_{-i}^t = (\underline{m}^{i, \tau}, \dots, \underline{m}^{i, \tau})$. The first is compatible with a single deviation at the action stage on the part of player $\langle j, t \rangle$. Therefore its probability is an infinitesimal in ε of order no higher than 2. The latter requires an action-stage deviation

in some period $t' < t$ (order 2 in ε), and $n - 2$ action-stage deviations in period t (order 1 each). Hence, player $\langle i, t \rangle$ has consistent beliefs if he assigns probability 1 to $m_{-i}^t = (\check{m}_i, \dots, \check{m}_i)$. The consistency of his end-of-period beliefs can then be checked from the profile μ .

Now suppose that for some $j \neq i$ we have that $a_j^t \neq a_j^i(i_\ell)$ and $a_{-i-j}^t = a_{-i-j}^i(i_\ell)$. Consistency of beliefs in this case requires showing that the second element in the sum in (T.11.3) is equal to 1. Of course given (T.11.3) it suffices to compare the probabilities of the two events $m_{-i}^t = (\check{m}_i, \dots, \check{m}_i)$ and $m_{-i}^t = (\underline{m}^{i,\tau}, \dots, \underline{m}^{i,\tau})$. The first requires $(n - 2)$ deviations at the action-stage of period t , each of order 2 in ε . Since $n \geq 4$, this is therefore an infinitesimal in ε of order no lower than 4. The second is consistent with a deviation of order 2 in ε at the action-stage of some period $t' < t$, together with a deviation of order 1 in ε at the action stage of period t . Therefore its probability is an infinitesimal in ε of order no higher than 3. Hence, player $\langle i, t \rangle$ has consistent beliefs if he assigns probability 1 to $m_{-i}^t = (\underline{m}^{i,\tau}, \dots, \underline{m}^{i,\tau})$. The consistency of his end-of-period beliefs can then be checked from the profile μ . The same argument applies to show the consistency of his end-of-period beliefs when $a_{-i}^t = a_{-i}^i(i_\ell)$. We omit the details.

In all other possible cases for a^t , the messages sent by all players $\langle j \neq i, t \rangle$ do not in fact depend on a^t , provided that m_j^t is either \check{m}_i or $\underline{m}^{i,\tau}$. Given (T.11.3), the consistency of the end-of-period beliefs of player $\langle i, t \rangle$ can then be checked directly from the profile μ . ■

T.12. Proof of Theorem 3

Given any $v^* \in \text{int}(V)$ and any $\delta \in (0, 1)$, using (B.8), (B.7) and the strategies and randomization devices described in Definitions B.4, B.5, B.8 and B.9 clearly implement the payoff vector v^* .

From Lemmas T.10.3 and T.10.5 we know that there exists a $\underline{\delta} \in (0, 1)$ such that whenever $\delta > \underline{\delta}$ each strategy in the profile described in Definitions B.8 and B.9 is sequentially rational given the beliefs described in Definitions T.9.1 and T.9.2.

From Lemmas T.11.1 and T.11.2 we know that the strategy profile described in Definitions B.8 and B.9 and the beliefs described in Definitions T.9.1 and T.9.2 are consistent.

Hence, using Lemma T.2.1, the proof of Theorem 3 is now complete. ■