ISSN 2279-9362

Research
Education
Outreach

CCA

Measuring Macroeconomic Tail Risk

Roberto Marfè and Julien Penasse

No. 715
February 2024

Carlo Alberto Notebooks

www.carloalberto.org/research/working-papers

Measuring Macroeconomic Tail Risk*

Roberto Marfè† Julien Pénasse‡

This draft: January 17, 2024

Abstract

This paper estimates consumption and GDP tail risk dynamics over the long run (1900–2020). Our predictive approach circumvents the scarcity of large macroeconomic crises by exploiting a rich information set covering 42 countries. This flexible approach does not require asset price information and can thus serve as a benchmark to evaluate the empirical validity of rare disasters models. Our estimates covary with asset prices and forecast future stock returns, in line with theory. A calibration disciplined by our estimates supports the prediction that macroeconomic tail risk drives the equity premium.

JEL: E44, G12, G17
Keywords: rare disasters, equity premium, return predictability

*We thank Daniel Andrei, Patrick Augustin, Bo Becker, Jules van Binsbergen, Pierre Collin-Dufresne, Stefano Colonnello, George Constantinides, Max Croce, Magnus Dahlquist, Darell Duffie, Leland Farmer, Xavier Gabaix, Anisha Ghosh, Eric Ghysels, Anisha Ghosh, François Gourio, Daniel Greenwald, Robin Greenwood, Benjamin Holcblat, Hendrik Hülsbusch, Christian Julliard, Leonid Kogan, Peter Kondor, Christos Koulovatianos, Hening Liu, André Lucas, Sydney Ludvigson, Rajnish Mehra, Christoph Meinerding, Alan Moreira, Tyler Muir, Stijn Van Nieuwerburgh, Marcus Opp, Loriana Pelizzon, Riccardo Sabbatucci, Marti Subrahmanyam, Urszula Szczerbowicz, Adrien Verdelhan, Emil Verner, Tan Wang, Michael Weber, Rüdiger Weber, and Irina Zviadadze for helpful comments. We also benefitted from useful comments from seminar participants at the Columbia University Macro Lunch, the University of Luxembourg, the University of Chile, the University of Venice, Fundação Getulio Vargas (Rio), McGill University, and the Stockholm School of Economics and conference participants at the CEPR ESSFM Gerzensee 2017 meeting, the SAFE Asset Pricing Workshop 2017, NFA 2017, the Paris December 2017 Finance Meeting, the MFA 2018, and the EFA 2020. A previous version of this paper was titled "The Time-Varying Risk of Macroeconomic Disasters."

†Collegio Carlo Alberto and ESOMAS, University of Turin, Piazza Arbarello, 8, 10122 Torino, Italy; +39 (011) 15630834; roberto.marfe@carloalberto.org, http://robertomarfe.altervista.org/

‡University of Luxembourg, 6, rue Richard Coudenhove-Kalergi L-1359 Luxembourg, Luxembourg; +352 466644 5824; julien.penasse@uni.lu, https://sites.google.com/view/jpenasse

Many puzzles in macro-finance arise from the inability of theories to reconcile asset prices and macroeconomic risk quantitatively. A classic example is the equity premium puzzle: averaged stock returns are far too high to be explained by the observed risk in consumption (Mehra and Prescott, 1985). A potential resolution of these puzzles is that we do not observe the risk that agents genuinely care about. As shown in Rietz (1988) and Barro (2006), a small probability of a macroeconomic disaster can greatly increase the equity premium. Models in which this probability is time varying can generate a volatile equity premium and thus can rationalize macro-finance puzzles, including excess volatility and predictability puzzles (Gabaix, 2012; Gourio, 2012; Wachter, 2013). The success of rare disaster models to rationalize asset pricing puzzles unfortunately relies on assumptions that are difficult to test. With standard preferences, asset prices are very sensitive to tail risk, meaning key parameters are difficult to pin down without relying on asset price data (Martin, 2013b). The assumption that tail risk is time-varying only makes matters worse. For these reasons, disaster risk has been described as "dark matter for economists" (e.g., Campbell, 2017; Cochrane, 2017; Chen et al., 2022).

This paper proposes a simple approach to measuring time-varying macroeconomic tail risk. Our approach circumvents the scarcity of large macroeconomic crises by making use of a broad cross-section of countries and by exploiting the informational content of variables that forecast the lower quantiles of consumption growth. Our approach delivers the time-varying probability of a large macroeconomic crisis by combining variables according to their ability to forecast the conditional distribution of consumption growth. We apply this approach to a large international panel, exploiting a rich information set that includes macroeconomic, political, and financial variables. We write our probability estimate as macro risk, or $\hat{\pi}$ for short. An attractive feature of our approach is that it does not impose structural assumptions or require the use of asset price information. This feature allows us to construct macro risk estimates that do not use asset price data, denoted by $\hat{\pi}^{-}$, which we use as a benchmark with which to evaluate the predictions of rare disasters models. In particular, backing up a time-varying estimate allows us to directly test whether asset prices covary with macroeconomic tail risk over time.

We estimate $\hat{\pi}$ semiparametrically using a two-step approach, similar to Adrian et al. (2019). First, we forecast the lower quantiles of consumption growth using quantile regressions. The dependent variable is consumption growth in a given country i between years t and $t+3$. The right-hand-side variables consist of 18 time-t macroeconomic, war and political, natural disaster, financial conditions, and asset price predictors. We use this predictor set to forecast the lower time-t quantiles of consumption growth. To avoid overfitting, the regressions are pooled across countries and include a penalization that shrinks the coefficient estimates toward zero. Second,

we interpolate the conditional quantiles to infer the probability of a large macroeconomic crisis, which we define as a two-standard-deviation drop in consumption below its long-term growth path. This approach delivers $\hat{\pi}$, the probability of a macroeconomic crisis in country i in year $t+1$. (The variant $\hat{\pi}^{-}$ excludes asset price predictors). Our data set covers 1900 to 2020, comprising 42 countries.

Figure 1 shows $\hat{\pi}$ for the United States. Macro risk mostly varies between 1% and 5%, with higher levels during the Knickerbocker Crisis, the Great Depression, the two World Wars, the Korean War, and the Great Recession of 2009. This time variation is the direct result of the ability of our model to forecast the lower quantiles of consumption growth. The most reliable predictors include realized recessions at home and in neighboring countries, world growth, wars and crises abroad, and the U.S. (or world) dividend-price ratio. We find that the forecasting model performs remarkably well out of sample and could, therefore, be used to predict crises in real time. Equally importantly, we document redundancy among predictor variables. For instance, $\hat{\pi}$ and $\hat{\pi}^{-}$ are highly correlated, and we cannot reject that they have equal forecasting power. Overall, this redundancy is quite robust across predictor categories, which suggests that the predictor set approximately spans the information set of investors.

Under this spanning condition, $\hat{\pi}$ can be interpreted as the rational expectation probability of a macroeconomic disaster. It thus provides a benchmark from which to evaluate theories in which macroeconomic tail risk plays a role. In the second part of the paper, we ask whether macro risk is related to the equity premium. We find that asset prices tend to be low (high dividend-price ratio) when $\hat{\pi}$ is high (corr. = 0.37). By itself, the U.S. dividend-price ratio—a standard equity premium proxy in the literature (e.g., van Binsbergen and Koijen, 2010)—captures a lot of information about future crises. A one-standard-deviation increase in the D/P ratio yields a 2.3% increase in macro risk. One would expect this result if investors were shunning stocks when macro risk is high, as predicted by rare disasters models. We find similar results using the equity premium proxy proposed by Martin (2017). Furthermore, macro risk directly forecasts international stock returns, confirming the link between tail risk and the equity premium. In contrast, macro risk is only weakly related to future consumption growth, meaning that our estimates reproduce the well-known disconnect between the equity premium and consumption. Importantly, these results hold when we estimate macro risk excluding asset price predictors, which means they are not the mechanical consequence of asset prices forecasting macroeconomic risk. We also show that predictability by $\hat{\pi}$ is distinct from predictability by realized recession and realized financial crises, which is useful to distinguish rare disaster risk from competing theories of asset price fluctuations.

Next, we ask whether rare disaster models help rationalize the equity premium puzzle and the other patterns we observe in the data. Following the Mehra and Prescott (1985) approach, we calibrate consumption dynamics in a consumption-based asset pricing model, derive asset pricing moments, and compare those moments to their empirical counterparts. A benefit of this approach is that we can work with estimated consumption dynamics that exclude asset price information, meaning that consumption parameters are not reverse-engineered to fit the asset pricing data.

We consider a rare disaster model similar to that of Wachter (2013), in which agents have recursive preferences. Consumption growth follows an exogenous stream, which is subject to rare disasters that occur with a time-varying probability. The model generates a high and volatile equity premium and a low risk-free rate under conservative preferences. It can also reproduce the option volatility skew, the predictability of stock returns, and the lack of predictability of consumption growth by asset prices that we observe in the data. An alternative dividend specification documents that time-varying disaster probability is still responsible for a sizable component of the equity premium and its dynamics, when its imperfect correlation with the D/P ratio matches that in the actual data. Importantly, our π estimates reveal a moderately persistent yet volatile disaster probability π. Chen et al. (2022) demonstrates that models like Wachter (2013), characterized by high persistence and low volatility in disaster probability π, tend to over-rely on "dark matter." This over-reliance often leads to reduced internal refutability and diminished out-of-sample performance. In contrast, our model's lower persistence and higher volatility in estimating disaster probability π do not necessitate significant "dark matter," resulting in enhanced internal refutability and stronger out-of-sample robustness.

Related literature

This paper draws inspiration from the seminal study of Barro (2006) and the subsequent literature that uses international macroeconomic data to measure macroeconomic tail risk. Nakamura et al. (2013) show that a rare disaster model can rationalize the equity premium even if crises unfold over several years and are followed by economic recoveries. As noted above, these claims remain controversial because of the high sensitivity of the equity premium to disaster risk. Indeed, several papers use asset prices to inform disaster risk estimates and conclude that rare disasters cannot rationalize equity or options prices (e.g., Julliard and Ghosh, 2012; Backus et al., 2011).

This early literature focused on the unconditional properties of macroeconomic tail risk, which is sufficient to evaluate models in which disaster risk and the equity premium are constant. In contrast, we are interested in studying the dynamics of macroeconomic tail risk. This is

important because variation in macroeconomic tail risk can rationalize puzzles related to the time-series behavior of equity prices. Furthermore, measuring tail risk dynamics is critical to comparing rare disasters to alternative mechanisms, such as time-varying risk aversion.

The literature has so far used two approaches to measure time-varying tail risk. One approach consists in using proxy variables that are likely to capture variation in tail risk. Berkman et al. (2011) follows that route using international political crises to proxy for disaster risk. They document that increases in the number of international political crises correlate with increases in the mean and volatility of world stock market returns, consistent with rare disaster models. A unique feature of their data is that it contains crises that can be precisely dated, which is ideal for event studies, i.e., to study relative *changes* in disaster risk. Manela and Moreira (2017) propose a news implied volatility index encoding information from the title and abstract of front-page articles published in the *Wall Street Journal*. They find that a large fraction of the variation in their index reflects words associated with war and policy, which can be interpreted as disaster concerns. A drawback of the proxy approach is that it is not best suited to study disaster risk in levels, because individual proxies do not span the information set of economic agents. For instances, many macroeconomic crises begin with other crises that are not political in nature. Another limitation is that proxies are not directly related to consumption, which means that additional assumptions are needed to evaluate the rare disaster model quantitatively.

An alternative approach consists in extracting tail risk from option prices. Options provide a hedge against tail events and are thus informative of the disaster concerns of option market participants. Siriwardane (2015) extracts a jump risk factor from short-term option prices for a large panel of U.S. firms. He shows that this factor can proxy for variation in the risk-neutral probability of a macroeconomic disaster. Under stronger assumptions, Barro and Liao (2021) back out the physical probability of a macroeconomic crisis using implied volatility surfaces for seven equity market indices. An important limitation of using option prices is that it requires structural assumptions about investors' preferences and the joint behavior of the stock markets and macroeconomic outcomes. In practice, these assumptions imply that option-based estimates are highly correlated with equity volatility, and thus likely capture market movements unrelated to macroeconomic tail risk. We indeed show that while volatility forecasts consumption tail risk, many other variables exhibit incremental forecasting power. Besides, option-based estimates that impose disaster model assumptions are subject to the critique that they cannot be used to falsify the very assumptions they are imposing.

Our paper generalizes the use of proxies by combining and weighing candidate proxies according to their ability to forecast the lower quantiles of consumption growth. One benefit of

our approach is that we do not take a stand on whether option prices or other risk proxies reflect disaster risk. The fact that $\hat{\pi}$ varies over time is the direct consequence of these candidate proxies forecasting macroeconomic tail risk. As our set of predictive variables includes a volatility index, our framework can accommodate volatility information without imposing structural assumptions, and on a much longer time frame. We find that volatility strongly forecasts macroeconomic risk, consistent with Manela and Moreira (2017) and Barro and Liao (2021). However, $\hat{\pi}$ delivers superior forecasts, reflecting the incremental forecasting power of the additional predictors in our data set. Consequently, $\hat{\pi}$ and volatility indices are positively but only imperfectly correlated, confirming that stock volatility varies for reasons others than macroeconomic tail risk. Likewise, the correlation between $\hat{\pi}$ and the number of political crises proposed by Berkman et al. (2011) is only 0.15, reflecting the fact that many variables beyond political crises forecast tail risk.[1]

Our paper is also related to the literature concerned with the estimation of rare disaster models. As noted in Campbell (2017) and Chen et al. (2022), rare disaster models have powerful effects on asset pricing moments and consequently standard efficient methods that rely excessively on asset pricing restrictions are often fragile. Several recent papers (see, e.g., Hansen (2008), Julliard and Ghosh (2012), Chen et al. (2022), Cheng et al. (2022)) develop robust estimation and testing methods to address these difficulties. We use a method of moments approach that exclusively targets consumption dynamics moments and then calibrate parameters that are relevant to asset prices. Our estimation approach can be seen as an extreme case of robust approaches in that it does not rely on asset pricing moments. Chen et al. (2022) show that such an approach is especially valuable when asset pricing moments provide most of the identification strength (i.e., when the dark matter measure is large). Our main contribution to this literature is that we measure π, which allows us to target moments directly related to π's dynamics. In other words, our estimations and quantitative analysis jointly help to uncover the "dark matter" often embedded in time-varying disaster risk models. We seek the parameters that match the following moments: mean consumption growth rate; standard deviation of consumption growth rate; expected number of crises (2-SD events); autocovariance of the unobserved disaster probability; variance of the unobserved disaster probability; expected disaster size. Our approach uses one moment condition for each parameter. We fit the model using rescaled international consumption data and the fitted disaster probability. The inference method does not assume that the disaster probability is observed and explicitly accounts for estimation error. The method also accounts for time aggregation, since we estimate the model

[1]The Online Appendix offers a more detailed comparison with the disaster risk index proposed by Berkman et al. (2011).

using three-year consumption data.

Measuring and understanding macroeconomic tail risk is critical for policymakers. A recent literature explores the possible nonlinear relationship between financial and economic conditions. Giglio et al. (2016) and Adrian et al. (2019) use quantile regressions to estimate the predictive U.S. GDP growth distribution conditional on measures of financial market distress. Similar to Adrian et al. (2019), we use quantile regressions to study the properties of international consumption growth, but our focus is on evaluating rare disaster models of asset prices.

The rest of this paper is organized as follows. Section I presents our data and methodology. Section II documents the forecastability of the lower quantiles of consumption growth. Section III presents macro risk estimates and studies the relationship between macro risk and variables, such as stock returns and macroeconomic growth. Section IV calibrates a rare disaster model using macro risk dynamics. Section V concludes. An Online Appendix contains additional results and proofs.

I. Methodology and Data

We begin by presenting our econometric framework. Models of time-varying disaster risk generally assume that the left tail of consumption growth is driven by a single state variable. This variable, which we denote macro risk $\pi_{i,t}$, is the probability that a given country i experiences a macroeconomic crisis at time t. As macroeconomic disasters generally unfold slowly over time, we will primarily work with consumption growth aggregated over multiple years. We write H-year consumption growth as

$$\Delta c_{i,t}^{(H)} \equiv \ln\left(\frac{C_{i,t}}{C_{i,t-H}}\right), \tag{1}$$

where C_t is real consumption per capita. A macroeconomic crisis is defined as a 2 standard deviation (SD) decline in consumption growth from its long-term growth path:

$$\text{Crisis}_{i,t}^{(H)} \equiv \begin{cases} 1, & \text{if } \Delta c_{i,t}^{(H)} < \text{mean}(\Delta c_{i,t}^{(H)}) - k \times \text{SD}(\Delta c_{i,t}^{(H)}). \\ 0, & \text{otherwise.} \end{cases} \tag{2}$$

We define macro risk as the probability of a macroeconomic crisis between times $t+1$ and $t+H$,

$$\pi_{i,t}^{(H)} \equiv \Pr(\text{Crisis}_{i,t+H}^{(H)} = 1 | I_{i,t}), \tag{3}$$

where the probability $\Pr(.|I_{i,t})$ is taken with respect to information $I_{i,t}$ available to economic agents at time t.[2] In what follows, we mostly use a three-year window (i.e., $H = 3$). To ease notations, we drop the horizon superscript whenever there is no ambiguity. We later study the sensitivity of our empirical results to this baseline definition.

This definition conforms with the approach in the rare disasters literature to define tail events as declines in consumption (or GDP) that exceed a fixed threshold value. Using a fixed threshold amounts to treating the long-term growth rate and volatility in normal times as constant parameters. Variation in $\pi_{i,t}$ thus summarizes movements in recession risk, which depends on the level, volatility, and higher moments of the growth process. Our approach slightly departs from the prior literature, which compares consumption peak to its trough to define macroeconomic disasters. For example, macroeconomic disasters in Barro and Ursúa (2008) have variable durations, with an average of three and a half years. We define macroeconomic crises over a fixed window instead since it is better suited to our forecasting framework.

A. Empirical Implementation

In the spirit of Barro (2006) and subsequent literature, we use international data to supplement the relatively small US sample that is commonly used in empirical work. A difficulty is that average growth rates and standard deviations tend to differ greatly across countries, even among industrialized ones. To ensure that international series are approximately comparable, we standardize the data such that consumption growth in each country has mean zero and unit standard deviation.[3] This assumption is similar to Nakamura et al. (2013), where countries are allowed to have different long-term growth rates and volatilities. We find that the rescaled consumption series exhibit a distribution that is comparable to the US series. Formally, we cannot reject that US and non-US series are from the same continuous distribution, using the two-sample Kolmogorov-Smirnov test (p-value = 0.11).[4]

A common concern in the rare disaster literature is that macroeconomic crises are so rare that even pinning down the unconditional probability of a crisis poses an empirical challenge. In addition to using international data, we gain traction by modeling the entire left tail of

[2] Remark that Eq. (3) allows for an arbitrary correlation between consumption growth conditioning on $I_{i,t}$.

[3] To ensure we do not use information that is not available in real time, we standardize consumption growth using a rolling window approach. Using a rolling window ensures that realized crises have similar magnitudes over the long run. To maximize the sample size, we initially require only 25 years of observations and then rescale using a 50-year rolling window.

[4] To further confirm that US and non-US macroeconomic crises capture similar events, we use the multidimensional two-sample Kolmogorov-Smirnov test, as described by Fasano and Franceschini (1987). We examine whether consumption growth, GDP growth, and the investment rate have statistically indistinguishable joint distributions, conditional on realized 2-SD crises. Although the test may not be effective due to the sample imbalance between the US and the rest of the world, we cannot reject that the conditional distributions are identical (p-value = 0.64).

consumption growth. Specifically, we work with a set of variables that have predictive power for the left tail, allowing us to model variation in the conditional distribution of consumption growth, summarized by π.

Denote a predictor set available to the econometrician as $X_{i,t}$. While $\pi_{i,t}$ is unobservable, we assume that $\pi_{i,t}$, is a function of the information available at time t: $\pi_{i,t} = f(X_{i,t})$ for some function f. We compute f semiparametrically using quantile regressions, similar to Adrian et al. (2019). In a first stage, we estimate conditional quantiles covering the left tail of consumption growth.[5] The quantile regression model of Koenker and Bassett (1978) assumes the conditional quantile of the dependent variable has an affine form. Specifically, we estimate the model

$$Q_{X_{i,t}}(\tau|X_{i,t}) = X_{i,t}B, \tag{4}$$

where $Q_{X_{i,t}}(\tau|X_{i,t})$ is a vector of conditional quantiles of three-year consumption growth given $X_{i,t}$, and τ is a set of probabilities covering the interval $[0.01, 0.5]$. Each column of the matrix B contains coefficients with respect to each quantile, which are estimated by minimizing the sum of absolute errors across quantiles. In a second stage, we fit the log-linear function $\hat{Q}_{X_{i,t}}(\tau \mid X_{i,t}) \approx g_{i,t}(\tau)$ for each country and year. This gives us an estimate of the probability of a 2-SD crisis as $\hat{\pi}_{i,t} = g_{i,t}^{-1}(-2)$. (Throughout the paper, we use "hats" to denote the empirical counterparts of the true unobserved values.)

This approach allows efficient use of consumption data since it exploits variation in the whole left tail and not only very rare events. We also explore an alternative approach where f takes the form of a standard normal distribution (i.e., $f(X_{i,t}) = \Phi(X'_{i,t} \cdot b)$ for some vector b. In that case, we obtain estimates of π by regressing realized crises (2) on the predictor variables using probit regressions.[6]

Our goal is that the vector $X_{i,t}$ approximates as closely as possible the information set $I_{i,t}$ of economic agents. We include as many variables as possible while maximizing the sample coverage over our international sample. While it is impossible to verify that we cast a wide enough net in terms of predictive variables, we expect $X_{i,t}$ to encode much more information than past realizations of consumption. We also expect to observe "decreasing returns to scale", such that the forecasting performance of our model is not too sensitive to small changes in the predictor set.

Overfitting is an immediate difficulty when working with a large number of predictors. To

[5] The Online Appendix provides further details on the estimation.

[6] We also use a probit model to construct alternative estimates that distinguish between global crises, like world wars, and more local crises, thus allowing different crises to be forecasted by different predictors. We present the results in the Online Appendix.

circumvent that difficulty, we exploit the panel structure of our data by restricting regression coefficients to be equal across countries. This implies that each quantile regression features less than one free parameter per country. While such homogeneity restrictions are unlikely to be literally true, they lead to more efficient parameter estimates, at the cost of a hopefully small bias. Pooled regression models are common in the forecasting literature and typically lead to improved out-of-sample forecasting performance (see, e.g., Garcia-Ferrer et al. (1987), for an early application to macro data).[7]

In addition, we also allow for the quantile minimization problem to include a penalty term to shrink possibly uninformative predictors. We consider penalties based on the L_1 and L_2 norms, as well as a combination of the two norms. A penalty based on the L_1 norm yields the quantile equivalent of the LASSO estimator ('least absolute shrinkage selection operator') of Tibshirani (1996). As is well known, the LASSO estimator imposes a sparsity assumption in that it leads some coefficient estimates to be exactly zero for large enough values of the LASSO parameter. The L_2 norm penalty corresponds to the ridge estimator, in which estimates will almost never be zero exactly. The ridge estimator tends to dominate the LASSO in settings with many correlated regressors. The combination of the two penalties yields the elastic net estimator of Zou and Hastie (2005). Whenever we estimate penalized quantile regressions, we select the penalization parameters using cross-validation. Namely, we slit the sample into ten parts, fit the model using nine-tenth of the data, and evaluate the model out-of-sample on the remaining tenth (i.e., ten-fold cross-validation). We select the penalization parameter(s) that minimize the out-of-sample quantile loss function averaged over the ten subsamples.

B. Consumption and GDP Data

Our consumption and GDP data were initially constructed by Barro and Ursúa (2010). The data consist of real per capita consumption expenditures and GDP. The sample covers 42 countries, and for many of these countries the data are available since the early nineteenth century. We extend the data to the years 2010–2020 using the World Bank's World Development Indicators. (The Online Appendix provides details on data sources and the construction of all variables used in the analysis.) The data span 1870 to 2020 and comprise 25 OECD countries; 14 countries from Latin America and Asia; and Egypt, Russia, and South Africa.

Figure 2 gives an overview of the tail events occurring over our sample. The figure shows a heatmap of negative realizations of consumption growth below −0.5 SD. Small black dots

[7]The Online Appendix presents a variant in which we relax the homogeneity restrictions for the quantile loadings on global variables. We show that the corresponding macro estimates remain strongly correlated to the baseline $\hat{\pi}$ (correlation = 0.86). However, this model sharply underperforms out-of-sample, even when overfitting is accounted for using penalized regressions.

represent data unavailability, and grey circles represent 2-SD macroeconomic crisis years (see the Online Appendix for a list of the crises). Data coverage begins at various times across countries but is generally continuous once it starts, except for Austria, Singapore, and Malaysia. We do not interpolate missing observations, although our results are unaffected when we do so. We observe an apparent clustering of tail events, particularly around the two world wars, the GFC, and the Covid-19 crisis.

Our baseline crisis definition (2) identifies a large macroeconomic crisis as a decline in the log consumption two SD below a country's long-term growth rate. This definition yields 156 crisis years, corresponding to a crisis frequency of 4.1%, which is similar to the 3.6% disaster frequency reported in Barro and Ursúa (2008). Realized crises are more common in the prewar period. The United States is a prime example, having experienced several macroeconomic crises, predominantly in the prewar years (1918–1919; 1930–1931; 2007–2008). Over the sample, the United States was statistically just as likely as any other country to face a macroeconomic crisis.

Note that the consumption series might be less reliable at the beginning of the sample, despite the rigorous precautions employed by Barro and Ursúa (2008) and others in assembling the data. A possible concern therefore is that measurement errors may lead us to exaggerate tail risk predictability. To see if measurement errors artificially inflate crisis persistence in the data, we test if realized crises are less autocorrelated after 1945. Reassuringly, Online Appendix Table A.VII shows that the autocorrelation of crises has not changed after 1945. This indicates that measurement errors are unlikely to exaggerate crisis predictability.

C. Predictive Variables

We assemble a relatively large annual data set of predictor variables spanning 1900 to 2020. We group the variables into five categories: macroeconomic; war and political; natural disaster; financial conditions; and asset prices. Table I gives an overview of our sample. We select variables based on their availability of comparable data for all countries over the most extended sample available. Many of these variables are binary indicators encoding recessions and other non-macro crises such as wars and banking crises. These variables naturally capture the possibility that macroeconomic risk increases during periods of stress. For instance, it is generally accepted that volatility is higher during recessions. Beyond volatility, including a recession indicator allows for capturing the possibility that macroeconomic risk increases during a recession. A limitation of the data is the potential for look-ahead bias in events such as wars, civil wars, political crises, and natural disasters. This issue stems from the difficulty agents may face in recognizing the onset of major upheavals as they unfold in real time. To address this, we delay

the start of such crises by one year. Additionally, crises in our sample typically span multiple years, meaning a crisis event in our data usually corresponds to an event that has been ongoing for at least one year. In the Online Appendix, we investigate an alternative approach where we substitute crisis indicators with predicted probabilities, calculated in a first step. These predicted probabilities prove to be very persistent, resulting in a high correlation with the crisis indicators. Consequently, this alternative approach yields results similar to our original method. The Online Appendix provides further details.

Macroeconomic predictors

Our sample of macro predictors is aimed at capturing the serial and spatial correlation in macroeconomic risk. "Recession" is a variable equal to one when a country experienced a negative growth rate in a given year. Recession events are quite frequent; Table I indicates that 28.20% of country-years are in a recession on average. "Recession abroad" captures the possible transmission of the realization of a foreign recession. It is constructed as a distance-weighted average of the recessions that occurred abroad:

$$\text{Recession abroad}_{i,t} = \sum_{j \neq i} d_{j,t}^{-1} \text{Recession}_{j,t}, \tag{5}$$

where $d_{j,t}$ measures the geographic distance between countries, normalized such that $\sum_{j \neq i} d_{j,t}^{-1} = 1$, for all j. We also include predictors based on past consumption growth rates and a "world" growth rate constructed by GDP-weighting consumption growth for the countries in the sample with available data.

War and Political Predictors

We create three variables encoding wars, civil wars, and political crises, primarily obtained from Sarkees and Wayman (2010) and the Center for Systemic Peace (CSP). Our wars and civil wars data consist of armed conflicts that resulted in deaths in battle. We define political crises as events of power disruption (e.g., a military coup), as well as periods of political instability (e.g., a country under occupation by a foreign power). We also construct a variable capturing foreign wars and political crises as in Eq. (5), where Recession is replaced by a variable equal to one whenever a foreign country enters into one of the three above-mentioned crises. With the exception of coups and similar events, these data describe spells that typically last for several years. Though defined based on systematically recorded crisis events, they might inadvertently contain a hindsight bias, giving the econometrician information that is only available ex-post. We thus delay the beginning of these crises by one year.

Natural Disaster Predictors

We create two variables capturing natural disasters. A "natural disaster" indicator equals one when a major earthquake, tsunami, or volcano eruption occurs in a given country. We focus on disasters that reached the highest grade in event classifications from the National Centers for Environmental Information. We also set our "natural disaster" indicator to one for all countries during the years corresponding to the Great Influenza Epidemic (1919-1920). (We forecast consumption growth up to 2020 using predictors up to 2017; this approach thus excludes the Covid-19 pandemic from the predictor data.) We also construct an indicator variable capturing the most devastating famines that occurred over the sample. To mitigate hindsight bias in the data, we delay natural disaster indicators by one year; for example, we start the Great Influenza indicator in 1919, although the pandemic began in 1918.

Financial Condition Predictors

Financial condition predictors include crisis dummies capturing banking, currency, and sovereign default crises, as well as extensive inflation periods. Our primary data source is Reinhart and Rogoff (2009), but we also checked the financial crisis indicators obtained from the Jordà-Schularick-Taylor Macrohistory Database (Jordà et al., 2017). We create a variable capturing the global financial cycle, proxied for by worldwide credit growth. We compute the three-year changes in bank loans-to-GDP ratio over three years using data from the Jordà-Schularick-Taylor (JST) database. Following Mian et al. (2017), we use one more lag that would be necessary in a standard forecasting regression, i.e., we use credit growth from four years ago to last year instead of credit growth from three years ago to the current year. The JST database covers 17 advanced economies, and we construct a GDP-weighed global factor, which we assign to the 42 countries in our sample.[8]

Asset Price Predictors

Our primary source for asset price data is Global Financial Data. To maximize the sample size, we work with the U.S. series to proxy for world variables. Estimates based on individual country predictors yield similar results, except for shorter periods. Asset pricing theory and empirical evidence suggest the dividend-price ratio is a natural proxy for the equity premium (e.g., van Binsbergen and Koijen, 2010). If the equity premium varies in proportion to disaster concerns, then the equity premium, and thus the dividend-price ratio, should forecast future

[8] We obtain very similar estimates using local credit growth for the 17 countries where the individual data are available, see Online Appendix F.

crises. Our proxy for the world dividend-price ratio is the S&P 500 dividend-price ratio.[9] We also construct a volatility index based on the CBOE VXO index from 1986. We extend the series back to 1871 using realized volatility as in Bloom (2009); to do so, we use daily data from 1885 through 1927, and monthly data before 1885, following Schwert (1989). Finally, we use a yield-curve slope predictor (e.g., Harvey, 1988; Estrella and Mishkin, 1998), based on the difference between the 10-year Treasury bond and the 1-year short rate.

II. Forecasting Macroeconomic Tail Risk

A. Quantile Regression Estimates

Figure 3 shows conditional quantile regression estimates for three-year consumption growth for the United States. The estimates are standardized and constructed at the .01, .02, .03, .04, .05, .075, .1, .15, .2, .3, .4, and .5 quantiles. The solid black line shows $\hat{\pi}$, which is interpolated as the probability of a 2-SD decline in consumption. We observe substantial variation in the quantiles over time; however, the extreme quantiles appear much more volatile than central quantiles. This is confirmed in Figure 4, which plots the associated quantile regression estimates for all predictor variables. The figure shows quantile coefficient curves for univariate and multivariate regressions. The curves are rarely flat along quantiles; instead, the slope coefficients tend to be larger in absolute value at the lowest percentiles.

We use a saturated regression to construct $\hat{\pi}$, as a large set of variables is desirable to approximate the information set of economic agents. In a time-series context, a saturated regression would likely suffer from in-sample overfitting as well as from near-collinearity between the predictive variables. However, the number of predictor variables (18) remains small relative to the number of countries (42) in our sample. We also show in the Online Appendix that the average pairwise correlation among variables is low.[10] We nevertheless implement a variable selection procedure using penalized regressions. The cross-validation procedure tends to select low values for the penalization parameter, resulting in modest shrinkage. We report ridge estimates in Figure 4 (dotted black line). These estimates are close to unpenalized estimates (dotted blue lines); we obtain comparable results with LASSO and elastic net penalizations.

We estimate $\hat{\pi}$ using multivariate regressions, but it is interesting to examine univariate regressions as well as they are easier to interpret. We report these coefficients in red. The multivariate slopes are generally smaller in magnitudes, which is expected, but tend to keep

[9] GFD offers a world dividend-price ratio, which is available from 1925. In our sample, the correlation between the two series is 0.91.

[10] The collinearity diagnostics of Belsley et al. (2004) further confirm that there are no statistically significant near linear dependencies in the predictor matrix.

the same sign as univariate slopes. The slope signs conform with the intuition that adverse shocks tend to have persistent negative effects. For instance, a country that experiences a recession in the forecasting year is exposed to more tail risk in the future. The same is true for non-macroeconomic shocks such as war or political crises, at home or abroad. Low equity prices (i.e., high dividend-price ratio) are also associated with more risk in the tail. This pattern is consistent with rare disaster models, which posit a positive relationship between the equity premium (proxied by the dividend-price ratio) and the likelihood of a macroeconomic disaster. Also, in line with rare disaster models (Barro and Liao, 2021; Farhi and Gabaix, 2016), stock volatility and currency crises forecast macroeconomic risk. Interestingly, credit growth is associated with lower tail risk, which contrasts with the body of evidence that credit growth forecasts worse macroeconomic outcomes (e.g., Mian et al., 2017). Most of the evidence in the literature is based on the postwar period and we find the evidence to be much weaker in our longer sample (see Online Appendix F for an extensive discussion).

To summarize the importance of individual predictors across quantiles, we calculate "marginal effects." We define these marginal effects as the increase in $\hat{\pi}$ for a unit increase in the variable of interest.[11] Figure 5 reports the marginal effects for each predictor in univariate regressions. To facilitate comparisons, we standardize the variables, so that each bar represents the effect of a unit standard deviation increase in the predictor. The figure also shows bootstrapped 90% confidence bands for the null of no predictability.[12] Most variables exhibit substantial forecasting power for the conditional distribution of consumption growth, which translates into statistically significant marginal effects. For instance, the standardized effect of a recession at the time of forecasting is 0.014. Given that the standard deviation of the crisis variable is 0.450 (see Table I), being in a recession raises the future probability of a crisis by $0.014/0.450 \approx 3.1\%$. Among the remaining group of variables, War, War/political crises abroad, credit growth, and the dividend-price ratio have the larger forecasting power. Being at war raises the likelihood of a crisis by about twice as much as being in a recession (6.8%). The predictive power of the U.S. dividend-price ratio is fairly large. A one-standard-deviation increase in the dividend-price ratio raises the likelihood of a disaster by 2.3%. We return to this link when we discuss the link between macro risk and the equity premium in Section III.B.

[11] We shock the predictor variable up and down and report the average of the two marginal effects.

[12] We bootstrap one-year consumption growth to account for the fact that the dependent variable is aggregated over three years.

B. *Macroeconomic Risk Estimates*

Figure 6 displays $\hat{\pi}$, the probability of a consumption crisis three years ahead (solid lines), together with the realized crisis years (shaded areas) for the 42 countries in our sample. (Figure 1 in the introduction zooms in on U.S. risk.) We construct $\hat{\pi}$ using the full set of predictors and the ridge model, but the unpenalized model and alternative penalized regressions yield very similar results. Later, we will work with estimates for a specification excluding asset price predictors, denoted by $\hat{\pi}^{-}$.

Macroeconomic tail risk is volatile, countercyclical, and persistent. Crisis probabilities usually range between 1% and 5%, with occasional spikes, in particular when crises arise, reflecting the serial correlation in tail risk in the data. Regressing $\hat{\pi}$ on time fixed effects yields a 65% R^2, meaning about three-quarters of the variance in country risk is due to common forces. This high correlation between country risk estimates is a result and not the product of a high pairwise correlation between predictor variables. Excluding global variables delivers a 62% R^2. (In the Online Appendix, we show that most variables have a low average pairwise correlation, with the exception of the Recessions and War and Political Crises Abroad variables.) Empirically, variables that are highly correlated tend to have more forecasting power, which results in a high correlation between macro risk estimates. We obtain a similar R^2 when we regress the U.S. estimates on the average cross-country probability (see the Online Appendix for a graphical comparison). This strong factor structure is consistent with Lewis and Liu (2016), which shows that a high degree of common macroeconomic risk is necessary to reproduce the fact that international asset return correlations are higher than consumption growth correlations.

We assess the forecasting performance of $\hat{\pi}$, using the Receiver Operating Characteristic (ROC) curve, a standard tool to assess the accuracy of binary indicators. Because crises are binary events, one is interested in converting probabilities into binary forecasts, which requires choosing a threshold over which to assign a value of one. The ROC curve shows the true positive rate (i.e., of all the crises that did happen, what fraction did the model predict?) against the false positive rate (i.e., how often the model wrongly predicts a macroeconomic crisis in $t+1$?). The curve obtains by varying all possible threshold values. As the threshold increases, the number of crisis signals drops, so fewer crises are correctly identified and incorrectly signaled. In contrast, for a lower threshold, more crises are correctly identified; the cost is that the frequency of false signals also increases. A model with no forecasting power results in a 45-degree line, whereas a model with a perfect fit would have an elbow-shaped ROC running from (0,0) to (0,1) to (1,1). The goodness of fit is measured by the area under the ROC curve (AUROC), with 0.5 corresponding to no explanatory power and 1 to a perfect fit. We plot the ROC curve for $\hat{\pi}$

in Figure 7. Figure 7 indicates that our predictive model has substantial forecasting power, with an AUROC statistic of 0.67. Of course, this remarkable forecasting power could reflect overfitting, as the curve uses in-sample estimates. We show in Section II.D that probabilities fitted out-of-sample are similar to in-sample estimates and that the model continues to perform well in that context.

C. Decomposition of Macroeconomic Risk

To better understand what drives variation in macro risk, we construct risk estimates for each predictor category. Figure 8 plots the breakdown of US macro risk across four predictor categories: Macro, War and Political, Financial Conditions, and Asset Prices. We leave out Natural Disasters predictors, as there was no major famine in the US in our sample (other natural disasters do not have a statistically significant predictive ability). Shaded areas highlight important events, namely recessions, wars, and banking crises.

The Macro component illustrates how even a small number of variables can capture meaningful patterns in a country's economic history. The component indicates elevated tail risk during the two World Wars, the Great Depression and the Great Influenza Pandemic, and finally during the Great Recession. The US economy experienced a disproportionally large number of recessions (highlighted by shaded areas) in the late nineteenth century and the beginning of the twentieth century. This does not however translate into an excessively volatile Macro component in the early part of the sample, which reflects the impact of other macroeconomic variables, such as the realization of crises abroad.

The Macro component only imperfectly captures the economic consequences of warfare. The War and Political component fills this gap. It shows marked spikes during the two World Wars, as well as the Korean and Vietnam conflicts after 1945. The bottom left panel of Figure 8 shows the US crisis probability given by the Financial Conditions component. Periods of low credit growth globally, such as the two World Wars and the recession of 1937-1938, are associated with increased risk. Banking crises identified in the Reinhart-Rogoff dataset (highlighted) also have significant importance.

The bulk of the variation in the Asset Prices component in Figure 8 is driven by movements in the dividend-price ratio. We again observe some redundancy with respect to the remaining components, with elevated macroeconomic risk during important events in US history, such as the entry of the United States into WW1 in 1917, the Depression of 1920-21, and the Great Depression. Around WW2, the Asset Prices component reaches a peak in 1937, then another one at a critical turning point, in 1941, when the US declared war on Japan. The probability

rises again during the Korean War, the 1973 oil crisis, the 1981-1982 recession, the 1987 crash, in 2002 after 9-11, and the Great Financial Crisis.

Figure 8 suggests that the predictor categories capture risks that often overlap. To get a better sense of the relative importance of each category, we decompose the variance of macro risk into each component. To do so, we regress $\hat{\pi}_{i,t}$ on the estimates associated to each category $\hat{\pi}^c_{i,t}$ (for $c = 1, \ldots, 5$):

$$\hat{\pi}_{i,t} = b^0 + \sum_{c=1}^{5} b^c \hat{\pi}^c_{i,t} + \hat{\pi}^{\varepsilon}_{i,t}. \tag{6}$$

We calculate the contribution of each component and the residual $\hat{\pi}^{\varepsilon}_{i,t}$ on the variance of $\hat{\pi}_{i,t}$ using

$$1 = \frac{\text{Cov}(\hat{\pi}_{i,t}, \hat{\pi}_{i,t})}{\text{Var}(\hat{\pi}_{i,t})} = \frac{\sum_c \text{Cov}(b^c \hat{\pi}^c_{i,t}, \hat{\pi}_{i,t}) + \text{Cov}(\hat{\pi}^{\varepsilon}_{i,t}, \hat{\pi}_{i,t})}{\text{Var}(\hat{\pi}_{i,t})}. \tag{7}$$

We present the results of this decomposition in Table II. The first column reports the results for the entire sample. Throughout this table, we report the decomposition for variations of the baseline specification. The table also reports the AUROC statistics, the volatility of the estimated macro risk, and the correlation of the various risk estimates with the baseline estimates. In Panel A, we exclude each predictor category in turn. In Panel B, we estimate the model in subsamples including the periods before and after 1945 and separating OECD and non-OECD countries. In Panel C, we vary the forecasting horizon and present decompositions of consumption risk cumulated from one to five years.

Table II indicates that all predictor categories except Natural Disasters contribute to a large fraction of the variance of $\hat{\pi}$. War and Political predictors contribute to about a third of the variance, followed by Macro and Asset Prices predictors that contribute to about a quarter of the variance each.

Excluding categories, however, does not result in large forecasting losses. An important subset, in particular, is the set that excludes Asset Price predictors, which we use to construct $\hat{\pi}^-$. We see in column (5) of Panel A that the AUROC statistic for $\hat{\pi}^-$ is 0.67, which is identical as that for $\hat{\pi}$. We show in the Online Appendix that the two estimates are indeed very similar. We believe this is important in assessing the role of macroeconomic tail risk in asset pricing, as discussed in Section IV. This redundancy across predictors is not limited to just asset price predictors. As illustrated in Figure Table II, when any one group of predictors is excluded from the set of predictors, the forecasting power remains largely the same and the probabilities remain highly correlated. This is an important result, as it suggests our predictor set approximately spans the information set of investors.

In Panel B, we see that the model performs similarly well in the different subsamples, with AUROCs well above 0.5. There is unsurprisingly more variation between the pre- and postwar periods. War and Political predictors explain more than half of the variance of $\hat{\pi}$ when estimated in the prewar sample, which is twice more than the postwar period. Correspondingly, excluding the prewar period results in macro risk estimates that are much less volatile. In contrast, the results are remarkably similar for the sample consisting of OECD and non-OECD countries. This suggests that our assumption that coefficients are homogeneous across countries results in a small bias. Turning to decompositions for alternative forecasting horizons in Panel C, we see that the model's forecasting performance remains very similar for crises measured over longer periods. We also see that the probability estimates are strongly correlated across horizons. This reflects the slowly unfolding nature of large macroeconomic crises. Interestingly, we do not find a similar forecasting power for mean consumption growth over longer horizons, as we discuss in Section III.A.

Overall, we have shown that the contributions of each predictor category and AUROC statistics are fairly stable. We show in the Online Appendix that the corresponding crisis probabilities are comparable across subsamples and specifications. In addition, series based on linear probit and penalized regressions are also very similar. In the next section, we evaluate the forecasting performance of the model out of sample.

D. Out-of-Sample Predictability

In this section, we evaluate the accuracy of our model in providing early-warning signals of upcoming crises. We back-test our model by replicating the analysis an econometrician would have done, using the proposed methodology to forecast crises out-of-sample. To do so, we employ a backtesting exercise and a sample split exercise.

First, we backtest the model by constructing $\hat{\pi}$ in real-time over an expanding window. We use a 50-year period to train the model (1900-1924). We move to the following year using a 50-year rolling window, so that each year t, we estimate $\hat{\pi}$ using data available from $t-50$ to t and use coefficient estimates to forecast realized crises in $t+H$. We run this test over horizons H from one to five years. The forecasts are compared over the 1925–2020 period.

Second, we run a sample split forecasting exercise. A limitation of the two out-of-sample tests outlined above is that they use the early part of the sample to train the forecasting model and test the forecast on the later portion. Any differences in predictability between the early and late parts of the sample could be missed by these exercises. In addition, we would like to provide further evidence that our homogeneity assumption results in a small bias. For this

exercise, we pick a random sample of 21 countries to train the model and test the model on the remaining 21 countries. We repeat this operation 100 times and report the averaged statistics.

We report the results of the two exercises in Table III, which shows out-of-sample AUROC statistics, as well as in-sample statistics evaluated over the same sample. The results are shown for realized crises defined over horizons of one to five years, for macro risk estimates based on the quantile regression models as well as the probit model. The forecasting performance is mechanically higher in sample, but we observe that the model remains quite successful out of sample. For instance, for the first forecasting test and baseline horizon $H = 3$, the QR model achieves an AUROC of 0.69 in sample, which is essentially equal to the out-of-sample AUROC (0.68). We find similar magnitudes in the remaining tests. The forecasting performance tends to be stronger at shorter horizons, yet the AUROCs remain significantly above 0.5 for all horizons. Penalized regressions slightly outperform unpenalized regression forecasts in some specifications. The probit model often underperforms, which suggests the quantile regression approach captures important information that is lost when fitting $\hat{\pi}$ using only realized crises.

E. Comparison with Option-Based Estimates

We interpret $\hat{\pi}$ as an estimate of the rational expectation probability of a macroeconomic crisis. This estimate is accurate as long as our predictor set approximately spans the information set of economic agents. Another natural approach to recover the rational expectation crisis probability is to use asset prices. The prices of derivatives, in particular deep out-of-the-money options, reveal the probability of a stock market crash, taken from the viewpoint of a risk-neutral agent, that is, under the risk-neutral measure. Under assumptions about investors' preferences, it is possible to recover the physical probability of a stock market crash (see, e.g., Bollerslev and Todorov, 2011). Imposing further assumptions about the link between stock returns and consumption, one can compute the physical probability of a consumption crisis. Backus et al. (2011) tackle the challenge, assuming a constant crisis probability. Siriwardane (2015) does not impose this assumption, but focuses on the risk-neutral crisis probability, while Barro and Liao (2021) back out the time-varying probability of a macroeconomic disaster, under the physical measure.

Barro and Liao produce risk estimates for six countries over the 1994–2018 period, which we report in Figure 9 (black lines) next to ours (blue lines). The unconditional option-implied probability equals 6.2%, which is much higher than $\hat{\pi}$ over that period. They are also more volatile, with several large spikes exceeding 20%. The presence of such spikes occurs in part because of option-based estimates are sampled monthly. In comparison, our yearly estimates

appear more stable.

The benefit of option-based estimates is that they do not require assumptions about investors' information set. The inconvenience is that these estimates rely on structural assumptions that may not be satisfied. For instance, the option-based estimates indicate that macroeconomic risk was about equally high in the United Kingdom, Germany, or Sweden during the Long-Term Capital Management crisis of 1998 as during the Great Financial Crisis. Likewise, using a longer U.S. sample that extends to 1983, Barro and Liao report that the probability of a macroeconomic disaster reached 89% at the month of the October 1987 stock market crash. Overall, option-based estimates tend to strongly covary with implied volatility, which indicates that their model attributes most of the variation in implied volatility to disaster concerns.

If the rare disaster model is literally true, the option-based estimates will contain all relevant information about future disasters. This is a testable assumption: option-based estimates will forecast crises, and no additional variable will have additional forecasting power. However, such a test is not practicable given the short sample over which options are available. To extend the sample back in time, we can proxy option-implied probabilities with realized volatility.[13] We already presented evidence of the ability of the volatility index to forecast future crises. However, on a stand-alone basis, the volatility index has weaker forecasting performances than $\hat{\pi}$. We find that the corresponding AUROC statistics 0.55 is considerably smaller than the corresponding statistics for $\hat{\pi}$ (0.67). A Wald test rejects the null that the volatility index has forecasting performance equal to $\hat{\pi}$, meaning that option prices do not fully summarize macroeconomic tail risk. This is to some extent unsurprising. The rare disaster model assumes that asset prices are driven by a single state variable. The simple fact that the U.S. dividend-price ratio and the VIX are not perfectly correlated is by itself a rejection of the model.

In the Online Appendix, we also compare $\hat{\pi}$ to the news-based volatility index (NVIX) proposed by Manela and Moreira (2017). While the NVIX aggregates information from *Wall Street Journal* headlines, it is designed to forecast the VIX and should thus reflect the risk perception and risk attitude of investors. Interestingly, Manela and Moreira (2017) show that the NVIX has forecasting power for U.S. consumption tail risk. This is consistent with our findings that volatility forecasts crises. Over the 1889–2016 period, the NVIX weakly forecasts international disasters, with an AUROC statistic, which is similar to the volatility index (0.55). Again, $\hat{\pi}$ exhibit substantial incremental forecasting power, consistent with the interpretation that equity volatility varies for reasons unrelated to macroeconomic tail risk.

[13] We show in the Online Appendix that option-implied probabilities and realized volatility indeed strongly covary, with a correlation coefficient of 0.92. Another caveat is that we use U.S. volatility to proxy for international volatility, although international option-implied estimates also strongly covary (Barro and Liao, 2021).

III. Macroeconomic Tail Risk over the Long Run

Having established the forecasting performance of our model, we now document new stylized facts about macroeconomic tail risk over our historical sample.

A. Consumption and Output Growth

In this section, we study the link between macro risk and macroeconomic growth. We run predictive regressions of H-year-ahead log of consumption and GDP growth on past country crisis probabilities $\hat{\pi}_{i,t}$:

$$\Delta c_{i,t+H}^{(H)} = a_i + b\hat{\pi}_{i,t} + u_{i,t+H}. \tag{8}$$

Table IV presents estimation results at 1-, 3-, and 5-year-ahead cumulative horizons. For comparison purposes, we also report similar results for cumulative one-year crises. We report slope estimates, standard errors, and R^2 statistics for the full sample (1900–2020), as well as individually for the prewar (1900–1945) and postwar (1945–2020) periods.

We use a bootstrap procedure to calculate standard errors and test whether slopes are different from zero. In particular, the procedure accounts for the fact that $\hat{\pi}$ is a generated regressor. We begin with the main dependent variable (say, excess returns) sampled at yearly frequency and standardized consumption growth. We cumulate the dependent variable to the appropriate horizon. We then construct $\hat{\pi}$ by regressing standardized consumption growth on the set of predictors and interpolating $\hat{\pi}$ from the fitted quantiles. We obtain the estimates of interest by regressing the dependent variable on $\hat{\pi}$. Next, we perform 1,000 bootstrap iterations. At each iteration, we simulate random draws for the dependent variable and standardized consumption growth. To simulate the dependent variable, we regress the yearly sampled dependent variable on year fixed effects and keep the fixed effect coefficients and the residuals. We simulate the yearly sampled dependent variable by drawing without replacement from the fixed effects and the residuals. As this variable is still sampled yearly, we cumulate the dependent variable to the appropriate horizon. To simulate standardized consumption growth from the quantiles, we draw a random uniform vector. The realizations above 0.5 are transformed to follow a standard normal distribution, while the realizations below 0.5 are interpolated from the conditional quantile estimates. We use these simulated samples of the dependent variable and standardized consumption growth to estimate the model and calculate standard errors.

We see that an increase in macro risk predicts a decline in consumption growth, but the effect dies out after three years, and magnitudes are relatively small. For instance, the full-sample coefficient for consumption growth is -0.32, meaning that a one-standard-deviation increase

in $\hat{\pi}$ (2.1%) lowers the next-year growth rate by 0.6%. Macro risk holds its forecasting power over longer horizons for GDP growth, but with similar magnitudes. In contrast, macro risk has robust forecasting power for future crises for all horizons.

We can relate these results to the well-known disconnect between asset prices and macroeconomic outcomes. In particular, Beeler and Campbell (2012) emphasizes the limited ability of the U.S. dividend-price ratio to forecast future dividends and consumption, meaning that most of the variation in the dividend-price ratio stems from changes in risk premia. We show in Section IV that time-varying disaster models can reproduce this feature of the data. Before turning to the model, we study the relationship between macro risk and the equity premium.

B. Macroeconomic Risk and Equity Premium Proxies

Rare disaster models predict that the equity premium varies over time in tandem with macroeconomic tail risk. The reason is that investors shun stocks when the probability of a disaster increases, effectively raising the equity premium. In what follows, we ask whether this link is present in the data. Since the equity premium is unobserved, we use three candidate proxies: the dividend-price ratio and two volatility-based series. The first series, SVIX^2, is the risk-neutral market return variance scaled by the gross risk-free rate. Martin (2017) shows that this index constitutes a lower bound on the equity premium that can be computed from option prices (the index is available from Martin's webpage for the 1996–2012 period). The second volatility series approximates SVIX^2 using realized volatility instead of option prices, which gives us a long time series. We compute this second proxy as our volatility index squared, which we then multiply by the real U.S. risk-free rate and scale to have the same mean and standard deviation as the SVIX^2. This longer series provides a reasonable approximation for the SVIX^2 (corr. = 0.78).

Figure 10 plots the probability of a macroeconomic crisis in the United States, together with the S&P 500 dividend-price ratio in Panel A, and with the two volatility proxies (Panel B). Because we use both the D/P ratio and the volatility index to construct our baseline crisis probability, $\hat{\pi}$, we present $\hat{\pi}^{-}$ in Figure 10, the variant excluding asset price predictors. We see that macro risk is positively related to both equity premium proxies. For instance, both macro risk and the D/P ratio spiked during the Great Depression and the two world wars. The 1990s stock market boom (low D/P) coincided with a marked decline in macroeconomic risk during the Great Moderation, as previously noted by Lettau et al. (2008). The correlation coefficient between the two series is about 0.45. The correlation between macro risk and the SVIX^2 appears more muted. For instance, squared volatility did not increase during the two world wars, in contrast with macro risk and the D/P ratio. This is in line with our earlier

observation, in Section II.E, of a positive but imperfect association between our volatility index and macro risk. The SVIX index is also noticeably less persistent, with a half-life of about six months, as compared with a half-life of 3 years for the D/P ratio. Consequently, our baseline macro risk measure, which captures the information of the D/P ratio and volatility, exhibits persistence that lies somewhere in between (a half-life of 1.7 years).

C. Macroeconomic Risk and Stock Returns

We further examine the link between macroeconomic tail risk and the equity premium using stock returns. We first test whether stock markets fall around macroeconomic crises and then ask whether macro risk forecasts future stock returns. We obtain stock market and one-year interest rate data from Global Financial Data, which we use to construct excess stock returns series. These are generally broad-market and market capitalization indices, such as the S&P 500 and the FTSE All-Share index (for the United Kingdom). However, these data are less comprehensive than macro series and often contain gaps, in particular during wars. We obtain total return and short rate data for 40 (of 42) countries, corresponding to a smaller panel of 2,030 year-country observations and 32 realized crises.

We first investigate how stock returns behave around macroeconomic crises. Rare disaster models unambiguously predict stock markets to crash during crises. The primary reason for such crashes is the fall in output and consumption, which coincides with lower demand for equities. However, stock markets may also reflect movement in macro risk. In particular, to the extent that crises can be anticipated, stock markets may crash before consumption and output fall. Furthermore, as stock markets attach a risk premium to macroeconomic tail risk, we may observe above-average returns after crises, when output and consumption return to normal and macro risk dissipates. Consistent with this reasoning, Barro and Ursúa (2017) find that large macroeconomic crises tend to coincide with stock market crashes and that crises are more likely, conditional on a stock market crash. We complement their results with an event-study design, to characterize the behavior of stock returns and consumption growth around macroeconomic crises. We adopt a flexible specification with five leads and lags around the year of a large macroeconomic crisis:

$$x_{i,t} = a_i + \sum_{j=-5}^{5} b_{-j}\text{Crisis}_{i,t+j} + u_{i,t}, \tag{9}$$

where $x_{i,t}$ is either the yearly standardized consumption growth or the local currency excess return in country i in year t. The fitted b coefficients thus give the relative returns during a

crisis, which occurs over the periods t, $t+1$, and $t+2$.

We report the fitted coefficients in Figure 11. Stock returns are depicted with error bars and consumption growth is shown with a dotted line. Consumption growth falls on average for the three crisis years, with the worst year occurring in the third year. These crises tend to be anticipated by equity markets. Stock prices tend to fall a year before the beginning of a crisis and remain negative for the two subsequent years. The cumulative decline equals 40.6%. Stock returns somehow rebound the next two years as the economy is exiting the crisis. The presence of a stock market decline followed by a rebound is consistent with rare disaster models' prediction that the equity premium increases when the economy enters a crisis.

We explore this prediction further in Table V, where we forecast stock returns in excess of the risk-free rate. We expect macro risk to positively forecast stock returns, reflecting the risk premium attached to periods of above-average risk. The table has a similar layout as Table IV. We estimate predictive regressions of the cumulative log excess returns on lagged macro risk over horizons of 1, 3, and 5 years. The top panel uses our baseline $\hat{\pi}$ probability; the middle panel uses the measure excluding asset price predictors, $\hat{\pi}^{-}$; the bottom panel shows results for regressions including time fixed effects.

We see that macro risk strongly predicts future returns. A 1% increase in $\hat{\pi}$ forecasts a 0.91-percentage-point larger return the next year, with the forecasts adding up to 2.3% over a five-year period. The regression coefficient suggests a volatile equity premium. For instance, given the volatility of macroeconomic risk (2.1%), the 0.91 coefficient implies a volatility for the one-year equity premium of about $2.1\% \times 0.91 \approx 1.9\%$. This is a large number, albeit smaller than estimates based on valuation ratios. For example, using the U.S. dividend-price ratio, Cochrane (2011) estimates that the U.S. equity premium varies by about 5.4% per year. Remark that $\hat{\pi}$ could forecast returns only incidentally because of its correlation with asset prices, which themselves contain information about future stock returns. This concern is only valid to the extent that our set of predictor variables is much smaller than the information set of market participants, in that asset prices carry too large a weight in $\hat{\pi}$ than in the true crisis probability. We address this concern by forecasting returns with $\hat{\pi}^{-}$, which is constructed by excluding these predictor variables. As the point estimates are almost identical, our conclusions remain unchanged.

Turning to subperiods, Table V indicates that predictability is weaker in the prewar period. One difference between the two subperiods is that much fewer disasters occurred during the latter. As we will see in the next section, rare disaster models can qualitatively reproduce this

feature of the data, with higher predictive coefficients in the absence of disasters.[14]

In the bottom panel of Table V, we examine the same predictability regressions including year fixed effects. Before looking at the results, we think it is useful to note that in complete markets with endogenous trading, all the variation in macro risk would be due to common forces. In fact, such models predict that consumption growth is perfectly correlated across countries. This is because, in such economies, agents endogenously share risk and can thus trade away country-specific risk. Furthermore, as only systematic risk should be priced, adding time fixed effects should result in insignificant coefficients. Of course, as international financial markets are clearly not fully integrated, we expect positive and perhaps statistically significant coefficients. Table V indicates that coefficients remain positive with year fixed effects. In the full sample, the coefficients are all statistically significant. In subperiods, the estimates are all positive but less precisely estimated, resulting in fewer statistically significant estimates. Predictability is relatively weaker after WW2 for short horizons. Interestingly, this weaker predictability is consistent with more financial market integration after WW2.

Next, we ask whether return predictability reflects actual macroeconomic tail risk or the realization of tail events. Table VI estimates the same return predictability regressions as the previous table, replacing the right-hand side variable to include realized recessions and financial crises. Rare disaster models predict that the equity premium varies over time because macroeconomic tail risk changes over time. Another plausible theory is that agents sell risky assets during recessions, causing asset prices to fall. This theory predicts that agents require a higher risk premium during crises, not because of the risk that their economic situation worsens, but because their willingness to take on risk has decreased. Models of habit formation, such as that of Campbell and Cochrane (1999), predict that, during recessions, agents' consumption approaches their habit level, effectively raising their risk aversion. In this class of models, movements in risk aversion rather than movements in risk cause the equity premium to move over time. This prediction is consistent with the pattern in Figure 11 that stock prices fall and rebound around crises. Table VI (first row) provides inconclusive support for this prediction, with insignificant or negative coefficients. In the second row, we add $\hat{\pi}$ to the right-hand side of the forecasting equation, with unchanged results.

Many models based on financial intermediation frictions propose that the price of risk increases during financial crises. The third and fourth rows repeat this forecasting exercise forecasting returns with realized financial crises. We see that the coefficients associated with financial crises are either insignificant or negative. This is consistent with the evidence in Baron

[14]Interestingly, Manela and Moreira (2017) also find that NVIX forecasts returns in the postwar sample, but not in the prewar sample. They also interpret their findings in terms of disaster concerns.

et al. (2021) that financial crises are not associated with higher future returns.[15]

IV. Asset Pricing Implications

In this section, we quantitatively evaluate the asset pricing implications of a workhorse equilibrium asset pricing model with time-varying macroeconomic tail risk. Our model is a discrete-time version of the one offered by Wachter (2013) and belongs to the discrete-time affine class proposed in Drechsler and Yaron (2011). We also consider an alternative dividend specification. In the spirit of Mehra and Prescott (1985), we ask whether a calibrated economy can reproduce key asset pricing moments. To do so, we take advantage of our methodology, which can produce macro risk estimates that do not exploit asset price information. Moreover, we investigate model-implied option prices.

A. Model

We consider a pure exchange closed economy, à la Lucas (1978), in which the log of consumption growth, Δc_{t+1}, evolves as follows:

$$\Delta c_{t+1} = \mu + \sigma\varepsilon_{t+1} + v_{t+1}, \tag{10}$$

where ε_{t+1} and v_{t+1} are two mutually independent shocks. The first shock is a standard normal random variable, and the second shock captures rare consumption disasters. We model v_{t+1} as a compound Poisson shock: $v_{t+1} = J_{t+1}\mathbf{1}_{\Delta n_{t+1}>0}$, where n_{t+1} is a Poisson counting process such that $\Delta n_{t+1} > 0$ describes a disaster event occurring at time $t+1$. The probability that the economy encounters a disaster in $t+1$,[16] π_t, follows a discretized square-root process:

$$\pi_{t+1} - \bar{\pi} = \rho(\pi_t - \bar{\pi}) + \nu\sqrt{\pi_t}u_{t+1}, \tag{11}$$

where $\bar{\pi} > 0$, $0 < \rho < 1$, and $\nu > 0$ are constants and where u_t is a standard normal random variable uncorrelated with ε_t and v_t. Following Cheng et al. (2022), disaster size J_{t+1} follows

[15] Another interesting predictor of equity returns is credit growth. Empirically, credit growth is often associated to worse macroeconomic outcomes and lower stock returns (Baron and Xiong, 2017). In the Onine Appendix, we replicate the negative association between credit growth and future returns documented in Baron and Xiong (2017) and show that both credit growth and $\hat{\pi}$ forecast returns in bivariate regressions.

[16] We slightly abuse the notation since π_t is only an approximation of the conditional probability of a disaster on the unit time interval (i.e., yearly). The Poisson counting process n_{t+1} has an intensity of π_t. The exact conditional probability of a single disaster occurring over the horizon h is, therefore, $\pi_t h \exp(-\pi_t h)$, while the probability that a disaster does not occur is $\exp(-\pi_t h)$. Hence, the residual probability that more than a single disaster occurs is $1 - \exp(-\pi_t h)(1 + \pi_t h)$. According to our estimates, the latter value is about 0.05% when π_t is at its steady state and less than 1% when π_t is at its 99th percentile. So to facilitate this terminology and ease the notation, we will use π_t for the conditional disaster probability and $\sum_t \Delta n_{t+1}$ (instead of $\sum_t \mathbf{1}_{\Delta n_{t+1}>0}$) for the number of disasters.

a parsimonious shifted negative exponential distribution with a moment-generating function given by

$$\varphi(u) = e^{-u\theta}(1+u\beta)^{-1}, \tag{12}$$

where disasters have support from $(-\infty, -\theta)$. The mean and variance are equal to $-(\theta+\beta)$ and β^2, respectively.

These dynamics are standard and can *potentially* capture several asset pricing regularities. This baseline model is most closely related to the continuous-time model of Wachter (2013). Gabaix (2012) calibrates a richer model that allows for movements in the disaster probability and the expected disaster size. The literature contains several alternative specifications. For example, Barro and Jin (2011) considers different laws for disaster size, such as single and double power laws. Gabaix (2011) and Gourio (2012) introduce time-varying disaster risk in production economies. Gourio (2008), Nakamura et al. (2013), Branger et al. (2016), and Hasler and Marfè (2016) consider more complex disaster dynamics, including unfolding consumption declines and subsequent recovery, as well as disasters leading to economic regime changes.[17] We intentionally keep the model simple, tractable, and parsimonious to focus on the time-series relationship between disaster probability and equilibrium asset prices.

We assume the economy is populated by a representative investor with recursive preferences (Epstein and Zin, 1989):

$$V_t = \left[(1-\delta)C_t^{1-1/\psi} + \delta(\mathbb{E}_t[V_{t+1}^{1-\gamma}])^{\frac{1-1/\psi}{1-\gamma}}\right]^{1/(1-1/\psi)}. \tag{13}$$

In this expression, δ, $\gamma \neq 1$, and ψ, respectively, capture time discounting, relative risk aversion, and the elasticity of intertemporal substitution. To ensure tractability, we focus on the case of ψ being equal to one. Similar to Collin-Dufresne et al. (2016), we solve for asset prices by expressing the stochastic discount factor in terms of the investor's log value function $vc_t = v_0 + v_\pi \pi_t$, which is affine in disaster probability π_t. The stochastic discount factor is given by[18]

$$M_{t+1} = \underbrace{\delta e^{-\gamma \Delta c_{t+1}}}_{\substack{\text{discounting of} \\ \text{expected utility}}} \times \underbrace{\frac{e^{(1-\gamma)vc_{t+1}}}{\mathbb{E}_t[e^{(1-\gamma)(\Delta c_{t+1}+vc_{t+1})}]}}_{\substack{\text{discounting of} \\ \text{continuation utility}}}. \tag{14}$$

[17] Another strand of the literature maintains a constant disaster probability but assumes that the representative agent learns about the model parameters or states so that the subjective disaster probability is time-varying (e.g., Weitzman, 2007; Koulovatianos and Wieland, 2011; Orlik and Veldkamp, 2014; Johannes et al., 2016; Ghaderi et al., 2022). Martin (2013a) shows that constant disaster risk spreads across assets in a multiple-tree economy and leads to endogenous risk premia dynamics.

[18] The Online Appendix provides details about the solution.

Note that the stochastic discount factor variance (i.e., the priced risk in the economy) is state-dependent and increases with the disaster probability. The log of the risk-free rate, $r_{f,t} = -\log \mathbb{E}_t[M_{t+1}]$, is affine in the disaster probability:

$$r_{f,t} = -\log \delta + \mu - \gamma\sigma^2 + \pi_t(\varphi(1-\gamma) - \varphi(-\gamma)). \tag{15}$$

This risk-free rate is stationary and decreases with the disaster probability π_t. Consumption risk increases with the likelihood of a disaster, and so does the demand for safe assets, causing the risk-free rate to fall. This effect strengthens with relative risk aversion.

We consider two dividend specifications. In the first specification, which we name the levered consumption model, the log aggregate dividend is proportional to consumption:

$$\Delta d_{t+1} = \phi \Delta c_{t+1}, \tag{16}$$

as in Abel (1999), Campbell (2003), Wachter (2013), and others. When $\phi > 1$, this ensures, in a parsimonious way, that dividends fall more than consumption when a disaster hits, as observed in the U.S. data (Longstaff and Piazzesi, 2004). In the second specification, which we name the co-integration model, the dividend dynamics follow:

$$\Delta d_t = \Delta c_t + \Delta s_t, \tag{17}$$

$$s_t = (1-\phi)\bar{s} + \phi s_{t-1} + \eta z_t + \kappa v_t, \tag{18}$$

where $s_t = \log D_t/C_t$ is the log dividend share of consumption and z_t is a normal random variable—independent of the other shocks for the sake of simplicity. The term κv_t implies that dividends load more on disaster than consumption, for $\kappa > 0$. Although simplistic, this specification is empirically motivated. First, s_t allows to break the unit correlation between consumption and dividend growth rates. Second, s_t induces co-integration between consumption and dividends. Third, recent evidence suggests that dividend growth risk decreases with the length of the horizon (Belo et al., 2015; Marfè, 2017). The assumption of Eq. (16) implies that dividend growth risk is higher than consumption growth risk at any horizon. Instead, co-integration preserves the larger dividend risk at short horizons but imposes that consumption and dividend growth risks are equal in the long run. This produces a downward-sloping dividend growth risk. Fourth, as documented by Gourio (2008) and Nakamura et al. (2013), disasters are often followed by recoveries. Co-integration generates dividend recoveries, as the dividend-consumption ratio increases after disasters, in line with the evidence in Longstaff and Piazzesi

(2004) and Hasler and Marfè (2016).

In addition, the second dividend specification lets us benchmark our model against an additional moment. In the previous section, we highlighted the relatively high correlation between macroeconomic tail risk and the dividend-price ratio. For the United States, the correlation with $\hat{\pi}^-$ is about 45.2%. Dynamics in Eq. (16) lead to a single state variable, π, which mechanically produces a unit correlation between π and the dividend-price ratio. This assumption could make us overestimate the correlation between prices and discount rates and, hence, the effect of disaster risk on the equity premium. Introducing co-integration gives rise to an additional state variable, s_t, which breaks the perfect correlation between π and the dividend-price ratio.

To solve for the equity claim price, we log-linearize returns around the unconditional mean of the log D/P ratio: $d_t - p_t \equiv \log D_t/P_t$. Campbell et al. (1997) and Bansal et al. (2012) document the high accuracy of such a log-linearization, which we assume to be exact. We use the Euler equation $1 = \mathbb{E}_t[M_{t+1}e^{r_{d,t+1}}]$ to recover that the log D/P ratio is affine in the disaster probability:

$$d_t - p_t = A_0 + A_\pi \pi_t + A_s s_t, \tag{19}$$

with $A_s = 0$ for the levered consumption model (see Eq. (16)). The D/P ratio is stationary and increases with the disaster probability π_t. These dynamics reflect a preference for the early resolution of uncertainty about time variation in π_t. An increase in disaster probability makes it more likely that crises will affect future consumption. An investor who prefers early resolution of uncertainty ($\gamma > 1$) is worried about current disaster risk and uncertainty in future disaster risk. Hence, equity prices are low relative to dividends when π_t is high (and vice versa). The substitution effect and the income effect offset each other because the elasticity of intertemporal substitution is equal to one.

The log of the equity premium is given by

$$\begin{aligned}\log \mathbb{E}_t[e^{r_{d,t+1}}] - r_{f,t} = \underbrace{\Omega_1 \gamma \sigma^2}_{\text{nondisaster risk}} + \underbrace{(\gamma - 1) k_1 A_\pi v_\pi \nu^2 \pi_t}_{\text{disaster probability risk}} \\ + \underbrace{(\varphi(\Omega_2) + \varphi(-\gamma) - \varphi(\Omega_2 - \gamma) - 1)\pi_t}_{\text{disaster size risk}},\end{aligned} \tag{20}$$

where $\Omega_1 = \Omega_2 = \phi$ for the dividend specification of Eq. (16) and $\Omega_1 = 1$, $\Omega_2 = 1 + \kappa(1 - k_1 A_s)$ for the dividend specification of Eq. (17), with k_1 being the endogenous log-linearization constant. The equity premium is given by three terms. The first term concerns nondisaster risk and gives rise to the usual consumption-CAPM compensation (Lucas, 1978). The second term reflects the compensation for the variation in disaster probability and increases with its

persistence and volatility for $\gamma > 1$. The third term is associated with disaster size risk and increases with disaster size and with its variance. Both the second and third terms increase with current disaster probability π_t.

B. Calibration

In this section, we assess the asset pricing plausibility of the rare disaster model using our macro risk estimates. In the spirit of Mehra and Prescott (1985), we estimate the parameters of the time-series processes of aggregate consumption and compare model-implied asset pricing moments to their empirical counterparts. This approach is similar to the recursive estimation/calibration approach highlighted by Hansen (2008) and Chen et al. (2022). Specifically, we use a simulated method of moments approach to estimate the parameters of the model (10)–(12) to match baseline moments of consumption growth data. We next calibrate the nuisance parameters to target selected asset pricing moments, while keeping the consumption parameter values equal to their estimates from the baseline moments. Chen et al. (2022) shows that the recursive estimation approach delivers better out-of-sample performance for models that exhibit large 'dark matter' such as rare disaster models.

We focus on the U.S. economy, consistently with our closed-economy setting, but we estimate the consumption parameters using international data. To do so, we rescale the standardized consumption growth series to match the first and second moments of US consumption growth. We assume the econometrician observes rescaled consumption growth, $\Delta c_{i,t}^{(3)}$, together with a candidate predictor of the left tail of consumption growth $z_{i,t}$. The econometrician also observes realized crises defined according to (2) (using a 2-SD threshold). Since we use rescaled consumption growth, this amounts to imposing country-specific thresholds to define crises based on the original consumption data for each country. The estimation approach does not rely on asset pricing data: we set the predictor $z_{i,t} = \hat{\pi}_{i,t}^{-}$ (i.e., the econometrician does not observe asset prices). The decision interval of the agent in the model is assumed to be annual, but we estimate the model using three-year data to match the baseline definition of macroeconomic crises. In particular, disasters in the model occur instantaneously and we account for the fact that disasters are measured over three-year periods in the data. Our modeling approach is thus immune to the measurement issues highlighted by Constantinides (2008), Donaldson and Mehra (2008), and Julliard and Ghosh (2012).

The model for consumption features six parameters: μ, σ, $\bar{\pi}$, ρ, ν, and β. We use six moment conditions that target each parameter. To ensure that estimation error is accounted for in the estimation, we calculate macro risk moments using the econometrician's estimate of $\pi_{i,t}$. In

a first step, the econometrician uses the predictor $z_{i,t}$ to construct this estimate, which we denote $\hat{\pi}^z_{i,t}$. For simplicity, we use ordinary least squares, such that the left-hand side variable is the realized crisis indicator (2) and the right-hand side variable is $z_{i,t}$. We view this as a simple way to tackle estimation error (i.e., $\hat{\pi}^z_{i,t}$ is affine in $\hat{\pi}^-_{i,t}$). In a second step, we calculate the six moment conditions, two of which depend on $\hat{\pi}^z_{i,t}$: $\mathrm{E}(\Delta c_{i,t+3})$, $\mathrm{E}(\Delta c^2_{i,t+3})$, $\mathrm{E}(\mathrm{Crisis}_{i,t})$, $\mathrm{E}(\hat{\pi}^z_{i,t}\hat{\pi}^z_{i,t-1})$, $\mathrm{E}(\hat{\pi}^z_{i,t}{}^2)$, and $\mathrm{E}(\Delta c_{i,t+3}\mathrm{Crisis}_{i,t})$. We calculate the model's population moment using a long sample of 10^7 observations, substituting $\hat{\pi}_{i,t}$ for $\pi_{i,t}$ in the simulation. As the model is exactly identified, we seek the parameters such that the observed moments equal the model-implied moments (up to simulation precision).

We find that this approach delivers an excellent fit of the distribution of three-year consumption growth, as can be seen in Figure 12. Consumption growth is left-skewed with a skewness of -0.57 and leptokurtic with a kurtosis of 4.91. Correspondingly, realized 2-SD crises are more frequent (4.1%) than the corresponding percentile for the normal distribution (2.3%). In contrast, the frequency of 2-SD booms is slightly lower than the normal percentile (2.1%). The model reproduces the skewness and, by design, the proportion of left-tail events observed in the data.

Table VII presents the point estimates, along with the associated standard errors Newey and West (1987) corrected using one lag in parentheses. Consumption growth parameters in "normal times" belong to the usual range of values: the long-term consumption growth rate is 2.3%, and consumption volatility is 2.8%. The subsequent rows show that the key disaster risk parameters that we estimate from the data are close to the calibrated parameters reported by Wachter (2013) and subsequent literature. Disaster risk is both persistent ($\rho = 0.74$) and volatile, with a scale parameter ν of about 0.22. These two parameters contribute to the unconditional volatility of disaster risk, which is $\sqrt{\nu^2\bar{\pi}/(1-\rho^2)} = 6\%$.[19]

The mean disaster size is $\hat{\beta} + \hat{\theta} = 12.2\%$ with a 6.6% standard deviation (the minimum disaster size is $\theta = 2\sigma = 5.6\%$). This is considerably smaller than the estimates found in previous calibrations in the literature. For instance, Barro and Jin (2011) estimates a mean size that is more than double (21.5%). There are two reasons for this discrepancy. First, our estimated disaster size corresponds to one-year contractions, while the size estimates in the literature are typically based on cumulative multiyear contractions (see Constantinides, 2008, Donaldson and Mehra, 2008, and Julliard and Ghosh, 2012). The corresponding three-year size targeted in our estimation is indeed larger (15.9%). Second, consumption tends to be more

[19] Wachter (2013) choose a more persistent $\rho = 0.92$ (using our notation), which is set to match the autocorrelation of the price-dividend ratio and is not estimated from macroeconomic data. The scale parameter is lower than in our setup ($\nu = 0.067$). Altogether, these parameters imply a smaller volatility of disaster risk (3.2%).

volatile in international data. We rescale international consumption data to match the mean and standard deviation of the US moments, thereby conservatively eliminating this effect.

We complete the calibration by setting the dividend parameters and the preference parameters. Evidence from the Depression era suggests that macroeconomic crises have a much larger effect on dividends than on consumption (Longstaff and Piazzesi, 2004). Then, concerning the levered consumption model, we follow Wachter (2013), who argues that a value of $\phi = 2.6$ is conservative and captures well the increased risk of dividends relative to consumption in both normal times and disaster periods. Concerning the co-integration model, we set $\bar{s} = \log(0.10), \phi = 0.90, \eta = 0.30$, and $\kappa = 4$ to capture the high persistence of the log D/P ratio, the correlation between disaster probability and the log D/P ratio, and the high level of short-term shareholders remuneration.[20] Lastly, we calibrate the parameters δ and γ to fit the main asset pricing moments. We find that reasonable values of relative risk aversion are sufficient to generate a high equity premium that conforms with the data. Hereafter, we consider the pair ($\delta = 99\%$, $\gamma = 5$) as our baseline calibration.

C. Asset Pricing Results

Tables VIII and IX show that this calibration leads to a relatively good fit of the main unconditional moments of asset prices. Panel A reports historical asset pricing moments and the distribution of their model counterparts. In particular, we report several percentiles from 1,000 simulated paths of the economy of length equal to 1,000 years. (We choose 1,000 years to approximate the statistical power of the panel regressions in Section II.) To avoid generating a negative probability of disaster, we simulate monthly series and then convert them to annual frequency.[21] In our calibration, the median risk-free rate is about 1.8% with 1.5% median volatility. In the levered consumption model, the median equity premium is 6.2%, and the median return volatility is 23.1%. These moments are slightly lower in the co-integration model (5.5% and 21.6% respectively). In the levered consumption model, the median D/P ratio is 3.4%, with 1.6% median volatility. These figures are slightly higher in the co-integration model (5.7% and 3.4% respectively). Overall, these moments line up well with their empirical counterparts. However, in the levered consumption model, the D/P ratio has a median first-order autocorrelation of the D/P ratio of 0.65, which is lower than in the data (about 0.78). This is the direct consequence of the persistence parameter ρ being smaller than in standard calibrations

[20] Belo et al. (2015) document that the volatility of aggregate dividends plus net repurchases is about 27% at the one-year horizon and then decreases substantially, a pattern that is captured by the dynamics of Eq. (17).

[21] Simulations of π_t at a monthly frequency (i.e., $\pi_{t+\triangle} = (1-\rho^{\triangle})\bar{\pi} + \rho^{\triangle}\pi_t + \nu\sqrt{\pi_t \triangle}u_t$ for $\triangle = 1/12$) have a negligible likelihood (less than 0.1%) of realizations lower than zero, which we replace by a small positive threshold.

directly targeting the D/P ratio persistence. Instead, the co-integration model is calibrated to capture well the persistence of the D/P ratio and the correlation between the log D/P ratio and the disaster probability π (about 50.3% vs 45.2%). Indeed, the co-integration model relaxes the correlation between the log D/P ratio and π because the former is an affine function of both π and the log dividend share of consumption s.[22]

Variation in macroeconomic tail risk contributes to generating a high equity premium. Agents who prefer an early resolution of uncertainty dislike not only bad news but also the possibility of more bad news in the future. They thus require compensation for any uncertainty associated with variation in the probability of a crisis. In particular, the equity premium and return variance can be decomposed into three terms associated with normal times, variation in the probability of a crisis, and severity of a crisis. Table X reports these values. Variation in macroeconomic tail risk generates the bulk of return variance in both the levered consumption model and the co-integration model (82% and 68% respectively). In the former case, it also commands a large fraction of the equity premium (about 74%), whereas it decreases substantially in the latter (about 45%). This is due to the fact that *(i)* we assume that normal risk of s is not priced for the sake of exposition and *(ii)* the co-integration model is calibrated in such a way to match the empirical imperfect correlation between π and the D/P ratio, which is instead mechanically one in the levered consumption model.

To further understand the role of risk in the equity premium, we report in Panel B of Tables VIII and IX the same moments, in a restricted simulated sample without disasters. There were indeed few realized disasters in the U.S. postwar economy and a common argument is that the high U.S. equity premium reflects a peso problem. For our model, we see that the equity premium is indeed slightly higher in this restricted sample (the median is about 1% higher than in the unrestricted case for both model specifications). However, return volatility and other moments are essentially unchanged. The only exception concerns the correlation between π and the D/P ratio in the co-integration model, which is somewhat higher in the restricted case. This happens because s is not hit by disasters anymore and, in turn, induces less variation to the D/P ratio. As a consequence, in absence of disasters, s weakens to a lesser extent the correlation between π and the D/P ratio.

Next, we investigate the extent of predictability implied by the model. Tables IV and V showed that π_t is a good predictor of excess returns, but not of consumption growth. We run similar regressions for either cumulative excess returns or cumulative consumption growth rates,

[22]This correlation depends on the log D/P ratio sensitivities with respect to each state variable. Of the two sensitivities, only the one with respect to π depends on γ, because π is a priced state variable. As a result, the correlation is positive and increases with risk aversion, and $\gamma = 5$ reproduces the correlation we observe in the data.

using horizons between 1 and 5 years. Figure 13 reports the 5th, 50th, and 95th percentiles of the predictive slope coefficients from the model simulations, together with their empirical counterparts, reproduced from Tables IV and V. Estimates for the levered consumption model and the co-integration model are in the upper and lower panels respectively and are very similar. In the left panels, we observe that disaster probability is a reliable predictor of excess returns. The distribution of the slope coefficients supports the positive values across all horizons, meaning that the model reproduces the equity return predictability we observe in the data. Indeed, the point estimates from Table V (also reported in Figure 13) display a similar pattern and belong to the interval produced by the model simulations at each horizon. In the right panels of Figure 13, we observe weaker evidence of consumption growth predictability, as in the data. Point estimates for the disaster probability are negative but close to zero, which is not surprising given the consumption dynamics of Eq. (10). The empirical counterparts from Table IV (also reported in Figure 13) exhibit a similar pattern.

Figure 13 also plots the median predictive slopes from the model simulations when realized disasters are turned off. In this case, the consumption growth predictability coefficients collapse to zero; indeed, disaster probability cannot predict consumption growth at all in the absence of realized disasters. In contrast, excess return slope coefficients are slightly larger in the absence of realized disasters and thus compare well with the postwar data coefficients (Table V).

Predictability results can be understood in light of the persistence of both consumption and dividend growth and excess returns. Tables A.XI and A.XII report these autocorrelations (from the first to the tenth yearly lag) from the model simulations described above. Table A.XI considers the levered consumption model where dividend growth dynamics are proportional to consumption growth dynamics. The model captures the lack of autocorrelation in consumption and gives rise to a negative autocorrelation of excess returns, quite in line with actual data. Thus, even if our estimates of disaster risk are somehow less persistent than calibrated disaster risk in Wachter (2013), the model leads to time-variation in risk premia in spite of unpredictable consumption growth. Table A.XII considers the cointegration model where dividend growth dynamics are enriched by the stationary variation of the dividend share of consumption. This specification of the model captures the additional empirical regularity of negative autocorrelation in dividend growth. Overall, although stylized, persistence predicted by the model accords well with the data and reassures on the model mechanism beyond the other model predictions discussed so far.

Finally, we also investigate the asset pricing predictions of a levered consumption model with CRRA preferences, such that the representative agent is indifferent regarding the timing

of uncertainty resolution. In turn, the main implication of introducing CRRA preferences is that the disaster probability π is no longer a priced state variable (i.e., the representative agent stops fearing the variation of π: $v_\pi = 0$). Nevertheless, the log equity premium still moves with π because changes in π induce changes in the disaster component of consumption risk, which increases with π. Table A.XIII in the Online Appendix reports the historical asset pricing moments and the distribution of their model counterparts. We keep all parameter values from the previous analysis of the levered consumption model. The model produces a risk-free rate (7.5%) and an equity premium (2.0%) that are, respectively, above and below their empirical counterparts. Changing the relative risk aversion coefficient can improve the fit of one moment, but only at the cost of deteriorating the fit of the other. In addition, the risk-free rate becomes almost twice more volatile than in the data.

D. Model-Implied Option Prices

In this section, we investigate model-implied option prices. Equity index options provide valuable insights as their prices indicate how investors perceive the likelihood of rare events. In particular, the prices of options for different strikes offer information about the market perception of equity crashes. Backus et al. (2011) connect the risk-neutral distribution of equity returns implied by options to the actual distribution of consumption growth, casting doubt on the plausibility of disaster risk models. Indeed, they document that *(i)* options prices imply smaller disaster probabilities than have been estimated from macroeconomic data and *(ii)* a consumption-based model provides a poor description of implied volatility when calibrated on macroeconomic data on disasters (e.g., Barro (2006), Nakamura et al. (2013)). While Backus et al. (2011) makes use of a rare disaster model with a constant probability and CRRA preferences, more recently Seo and Wachter (2019) argues that a time-varying disaster model under recursive preferences can reconcile these tensions. They show that the Wachter (2013) model can in fact reproduce the implied volatility level and skew observed in the data. In what follows, we thus examine the option prices implied by our calibration, which is disciplined by our estimates for π and a more conservative disaster size distribution.

Figure 14 illustrates our results. We depict prices as implied volatilities and plot them against a measure of moneyness, with higher strike prices on the right side. The measure of moneyness is computed as the proportionate difference between the strike and price: (strike − price) / price. A value of zero corresponds to an at-the-money option (strike = price). We use options with both a 3-month and a 1-year maturity, with deviations from return lognormality (i.e., flat volatility smiles) that are more discernible at the shorter maturities.

The upper panels of Figure 14 compare the implied volatilities predicted by our model with the empirical evidence from S&P 500 index options. Implied volatilities typically average around 20% and decrease with moneyness, the so-called volatility skew, with a steeper slope at short maturity. The levered consumption model (see Eq. (16)) leads to a very good fit of the data at both 3-month and 1-year maturity. Instead, the co-integration model (see Eq. (17)) predicts lower implied volatilities with a steeper slope than in the data for in-the-money put options, in particular for the 3-month maturity. In both cases, the model helps filling the gap documented by Backus et al. (2011) between actual option prices and those generated by a CRRA consumption-based model with a constant probability of disasters. In particular, the volatility skew produced by the levered consumption model compares well with the skew in Seo and Wachter (2019). This points to the importance of a time-varying disaster probability and highlights *(i)* the plausibility of disaster risk estimated from macroeconomic data and *(ii)* the sensitivity with respect to the dividend dynamics.[23]

To further inspect the model mechanism, the lower panels of Figure 14 consider two alternative model specifications. Either we shut down time-variation in the disaster probability (and increase the relative risk aversion to 7.5 to keep a sizeable equity premium) or we shut down the preference for the early resolution of uncertainty (by adopting CRRA preferences). In both cases, we observe that the level of implied volatility drops substantially and the volatility skew is more pronounced than in actual data. Therefore, we conclude that the joint assumption of time-varying disaster risk and preference about the resolution of uncertainty, initially proposed by Wachter (2013), is key to generate model-implied option prices that conform well with the data.

E. Summary of the Calibration Analysis

Overall, our model supports the idea that macroeconomic tail risk is a key driver of the equity premium. In particular, the proposed calibration—where we embed our estimates of macroeconomic tail risk that make use of macroeconomic data only—indicates that the model can generate a large and volatile equity premium. Moreover, model-implied options can generate a realistic volatility skew, meaning that macroeconomic tail risk helps rationalizing how market

[23] As an additional exercise, we compare the sample moments of an empirical measure of the volatility skew—that is a measure of return non-gaussianity implied in option prices—with its model counterpart. The latter is obtained by plugging in the estimates of $\hat{\pi}^-$ over the same sample. Table A.XIV in the Online Appendix reports the results. In accord with the visual inspection of Figure 14, the levered consumption model provides a good fit of the mean and slightly underestimates the standard deviation of the volatility skew observed in actual data, whereas the co-integration model overestimates both moments. Results are quite insensitive to using $\hat{\pi}^-$ from either the current or the previous year. Since $\hat{\pi}^-$ is estimated at yearly frequency, the sample is too short for a proper empirical assessment, but these figures are quite reassuring about the model predictions concerning option prices.

participants price extreme events. We acknowledge that other aspects of the model are overly simplistic. In the levered consumption model, macroeconomic tail risk π_t is the only state variable, which means that all equilibrium quantities are perfectly correlated, and the model cannot explain the empirical wedge between risk and return. We emphasized in Section III.B that macroeconomic tail risk and the dividend-price ratio are positively but imperfectly correlated. The co-integration model actually reproduces this feature of the data, and still generates a sizeable equity premium and return predictability, but the fit of the volatility skew partially deteriorates.[24]

V. Conclusion

A major criticism of the rare disaster paradigm is its reliance on an elusive state variable, which has been referred to as "dark matter for economists" (Campbell, 2017; Cochrane, 2017; Chen et al., 2022). In this paper, we propose an approach to measure macroeconomic tail risk over the long run. Our estimates capture the predictable variation in the lower tail of consumption growth, making use of and summarizing the information content of a rich set of predictive variables. The framework does not impose structural assumptions and does not require asset pricing information, both of which are key to test the empirical validity of the rare disaster paradigm.

We find that the lower quantiles of consumption growth are predictable by lagged recessions at home and abroad, wars and related disasters, and asset prices, among other variables. By combining the information in our predictive variables, we obtain country crisis probabilities that capture the rich economic history of the 42 countries present in our sample. Reflecting the predictive ability of its components, the crisis probabilities increase in periods of economic and political distress and covary with asset prices, as implied by rare disaster models.

We calibrate a rare disaster model using our estimates of macro risk dynamics. In doing so, we are careful to avoid asset price information, a step that ensures our calibration has not been reverse-engineered to fit asset pricing data. The model generates a high and volatile equity premium for a representative agent whose coefficient of relative risk aversion is 5 and whose elasticity of intertemporal substitution is 1. These results support prior calibrations of rare disasters model set to match asset pricing moments.

Our approach is inspired by the long tradition in finance of estimating the equity premium

[24] Section II.E also showed that many variables beyond option prices have forecasting power for future crises, which we interpret as evidence of an imperfect correlation between macroeconomic tail risk and equity volatility. While the co-integration model still produces a unit correlation with volatility, a richer model, for instance, assuming conditional heteroskedasticity in the dynamics of s_t, could break this perfect correlation.

using predictive regressions. Similar forecasting models are typical in the context of forecasting recessions and other crises. We show that this type of forecasting model can measure economic "dark matter," such as disaster risk. We leave to further research the study of other similar objects, such as long-run growth (Bansal and Yaron, 2004), which our framework can accommodate as well. Finally, we note that rare disaster models can rationalize other regularities observed in the markets for bonds and currencies (Gabaix, 2012; Gourio, 2013; Farhi and Gabaix, 2016), and offer the possibility of connecting macroeconomic aggregates with asset prices in production economies (Gabaix, 2011; Gourio, 2012; Kilic and Wachter, 2018; Isoré and Szczerbowicz, 2017). An empirical evaluation of these models would be a fruitful area for further research.

Table I: Predictive Variables

Variable	Description	Mean	S.D.	Min	Max
	Macroeconomic Predictors				
Recession	One-year consumption growth rate < 0	0.282	0.450	0.000	1.000
Recession abroad	Distance-weighted average of foreign recessions	0.230	0.161	0.000	0.890
Consumption	Consumption growth	0.019	0.063	−0.557	0.487
Consumption (world)	GDP-weighted consumption growth	0.018	0.039	−0.213	0.116
	War and Political Predictors				
War	Interstate war	0.040	0.195	0.000	1.000
Civil war	Intrastate or non-state war	0.059	0.235	0.000	1.000
Political crisis	Political crisis	0.013	0.115	0.000	1.000
War/political crisis abroad	Distance-weighted average of foreign war/political crises	0.096	0.117	0.000	0.761
	Natural Disaster Predictors				
Natural disaster	Major earthquake, tsunami, volcano eruption, and the Great Influenza	0.070	0.255	0.000	1.000
Famine	Major famine	0.008	0.091	0.000	1.000
	Financial Conditions Predictors				
Banking crisis	Banking crisis	0.088	0.284	0.000	1.000
Currency crisis	Exchange rate depreciation $\geq 15\%$	0.144	0.351	0.000	1.000
Sovereign default	Sovereign default	0.093	0.291	0.000	1.000
Hyperinflation	Annual inflation rate $\geq 20\%$	0.099	0.299	0.000	1.000
Credit growth (world)	GDP-weighted 3-year change in total bank loans to GDP	0.019	0.042	−0.132	0.084
	Asset Price Predictors				
Dividend-price ratio (US)	S&P 500 dividend-price ratio	0.040	0.016	0.011	0.093
Stock volatility (US)	S&P 500 realized volatility	0.200	0.073	0.050	0.507
Yield curve (US)	10-year–1-year Treasury spread	0.001	0.017	−0.053	0.038

Notes. We report descriptive statistics for the predictive variables (see Section I.C and the Online Appendix for detailed descriptions of each variable).

Table II: Variance Decompositions for Macro Risk

	(1)	(2)	(3)	(4)	(5)
		Panel A: Predictor Categories			
Macro	0.37	–	0.61	0.51	0.39
War and Political	0.28	0.48	–	0.38	0.27
Natural Disasters	0.00	0.01	0.01	0.01	0.00
Financial conditions	0.28	0.36	0.31	–	0.28
Asset Prices	-0.02	0.06	0.03	0.05	–
Residual	0.09	0.10	0.04	0.05	0.05
N	3767	3767	3767	3767	3767
Countries	42	42	42	42	42
AUROC	0.67	0.67	0.65	0.64	0.67
$\sigma(\hat{\pi}') \times 100$	2.05	1.59	1.65	1.59	1.79
$\rho(\hat{\pi}, \hat{\pi}')$	1.00	0.92	0.90	0.88	0.99
		Panel B: Subsamples			
	Full sample	Pre-1945	Post-1945	OECD	Ex-OECD
Macro	0.37	-0.01	0.15	0.35	0.37
War and Political	0.28	0.56	0.06	0.26	0.30
Natural Disasters	0.00	0.01	0.01	0.02	0.00
Financial conditions	0.28	0.11	0.55	0.29	0.16
Asset Prices	-0.02	0.04	0.08	-0.01	0.07
Residual	0.09	0.28	0.15	0.09	0.09
N	3767	994	2773	2585	1182
Countries	42	31	42	25	17
AUROC	0.67	0.75	0.76	0.66	0.73
$\sigma(\hat{\pi}') \times 100$	2.05	5.24	1.65	1.88	1.23
$\rho(\hat{\pi}, \hat{\pi}')$	1.00	0.61	0.39	0.97	0.77
		Panel C: Horizons			
	$H = 3$	$H = 1$	$H = 2$	$H = 4$	$H = 5$
Macro	0.37	0.56	0.38	0.37	0.08
War and Political	0.28	0.10	0.20	0.06	0.02
Natural Disasters	0.00	0.03	0.02	0.01	0.02
Financial conditions	0.28	0.15	0.18	0.43	0.60
Asset Prices	-0.02	0.06	0.06	0.05	0.16
Residual	0.09	0.09	0.17	0.07	0.12
N	3767	3767	3767	3722	3674
Countries	42	42	42	42	42
AUROC	0.67	0.84	0.78	0.65	0.65
$\sigma(\hat{\pi}') \times 100$	2.05	3.07	2.82	1.82	1.51
$\rho(\hat{\pi}, \hat{\pi}')$	1.00	0.88	0.96	0.89	0.71

Notes. This table presents variance decompositions of macro risk on the five predictor categories and a residual (see Eq. (7)). The table also reports the AUROC statistics, the volatility of the corresponding macro estimates, and their pairwise correlations with the baseline model estimates. Panel A compares the results for the full sample with results for predictor subsets where each predictor category is excluded in turn. Panel B shows subsample results comparing pre- and post-WW2 data, OECD and non-OECD countries. Panel C gives results for estimates constructed using consumption risk cumulated from one to five years.

Table III: Out-of-Sample Crisis Predictability

Model	In Sample					Out of Sample				
H	1	2	3	4	5	1	2	3	4	5
				Panel A: Expanding Window						
Unpenalized	0.87***	0.81***	0.69***	0.71***	0.72***	0.83***	0.76***	0.68***	0.65***	0.67***
LASSO	0.86***	0.82***	0.67***	0.70***	0.72***	0.83***	0.75***	0.66***	0.65***	0.67***
Ridge	0.86***	0.82***	0.67***	0.70***	0.72***	0.83***	0.78***	0.67***	0.65***	0.67***
Elastic	0.86***	0.82***	0.67***	0.70***	0.72***	0.83***	0.77***	0.69***	0.64***	0.67***
Probit	0.87***	0.83***	0.68***	0.73***	0.74***	0.82***	0.71***	0.60***	0.60***	0.62***
				Panel B: Sample Split						
Unpenalized	0.85***	0.79***	0.70***	0.69***	0.69***	0.81***	0.73***	0.63***	0.57	0.57
LASSO	0.85***	0.73***	0.67***	0.67***	0.67***	0.81***	0.72***	0.60***	0.61***	0.57
Ridge	0.85***	0.79***	0.67***	0.67***	0.68***	0.81***	0.75***	0.62***	0.61***	0.59
Elastic	0.84***	0.78***	0.69***	0.67***	0.68***	0.80***	0.73***	0.63***	0.59	0.59
Probit	0.86***	0.80***	0.69***	0.71***	0.69***	0.76***	0.70***	0.59	0.58	0.57

Notes. This table reports Area under the Receiver Operating Characteristic (AUROC) statistics for the prediction of future crises based on $\hat{\pi}$. Crises are defined as 2-SD cumulative declines in consumption growth measured over $H = 1, \ldots, 5$ years. In each case, $\hat{\pi}$ is estimated on a training sample of a given horizon and tested out of sample. The table reports AUROC in both samples. We evaluate estimates constructed with quantile regressions model with no penalization and LASSO, Ridge, and Elastic net penalized regression models and with the probit model. Panel A gives results where $\hat{\pi}$ is constructed out-of-sample using an expanding window after a 50-year training period (1900-1924). The forecasts are compared over the 1925–2020 period. Panel B gives results for a variant where crises are identified out of sample. Panel C gives results for the full historical sample (1900-2020), for repeated estimations of the model trained on 21 countries picked randomly and tested on the remaining 21 countries. We report statistical significance based on bootstrapped AUROC confidence bands.
*, **, and *** indicate statistical significance at 10%, 5%, and 1% levels, respectively.

Table IV: Macro Risk Weakly Forecasts Consumption Growth

	Full sample (1876-2020)			Prewar (1876-1945)			Postwar (1945-2020)		
H	1	3	5	1	3	5	1	3	5
				Consumption growth					
				$\Delta c_{i,t+h}^{(H)} = a_i + b\hat{\pi}_{i,t} + u_{i,t+H}$					
b	-0.32^{***}	-0.46^{**}	-0.16	-0.34^{***}	-0.78^{***}	-0.41	-0.13	0.13	0.23
	(0.06)	(0.18)	(0.26)	(0.10)	(0.27)	(0.43)	(0.10)	(0.22)	(0.32)
N	3,835	3,835	3,745	1,052	995	933	2,783	2,840	2,812
R^2	0.022	0.012	0.001	0.021	0.032	0.005	0.005	0.001	0.003
				GDP growth					
				$\Delta y_{i,t+h}^{(H)} = a_i + b\hat{\pi}_{i,t} + u_{i,t+H}$					
b	-0.34^{***}	-0.67^{***}	-0.67^{**}	-0.40^{***}	-0.84^{***}	-0.67	-0.16	-0.37	-0.53^{*}
	(0.06)	(0.17)	(0.26)	(0.09)	(0.26)	(0.43)	(0.10)	(0.22)	(0.31)
N	3,830	3,828	3,747	1,052	995	936	2,778	2,833	2,811
R^2	0.028	0.028	0.016	0.030	0.038	0.012	0.007	0.009	0.012
				Crises					
				$\text{Crisis}_{i,t+h}^{(H)} = a_i + b\hat{\pi}_{i,t} + u_{i,t+H}$					
b	1.23^{***}	2.73^{***}	3.21^{***}	1.51^{***}	4.09^{***}	5.25^{***}	0.63^{*}	1.17	1.70
	(0.23)	(0.60)	(0.93)	(0.35)	(0.93)	(1.46)	(0.35)	(0.76)	(1.08)
N	3,835	3,835	3,745	1,052	995	933	2,783	2,840	2,812
R^2	0.046	0.058	0.046	0.054	0.107	0.088	0.013	0.013	0.018

Notes. This table presents slope coefficients, standard errors, and adjusted R^2 statistics for predictive panel regressions of cumulative consumption growth on crisis probabilities. Observations are over the sample of 42 countries, from 1900 to 2020. Results are reported for the full sample, as well as for subsamples covering the pre- and post-WW2 periods. Standard errors in parentheses are bootstrapped to account for cross-sectional correlation and time aggregation. The bootstrap simulates the estimated $\hat{\pi}$ when it is used as a regressor to account for estimation error.
*, **, and *** indicate statistical significance at 10%, 5%, and 1% levels, respectively.

Table V: Macro Risk Forecasts Excess Stock Returns

	Full sample (1876-2020)			Prewar (1876-1945)			Postwar (1945-2020)		
H	1	3	5	1	3	5	1	3	5
				$\sum_{h=1}^{H} r_{i,t+h} - r^f_{i,t+h} = a_i + b\hat{\pi}_{i,t} + u_{i,t+H}$					
b	0.91**	1.91*	2.32	0.98**	2.02	1.91	1.36**	2.35*	2.94
	(0.37)	(0.98)	(1.44)	(0.47)	(1.36)	(2.04)	(0.62)	(1.35)	(1.88)
N	1,753	1,707	1,661	641	601	561	1,112	1,106	1,100
R^2	0.014	0.020	0.019	0.032	0.036	0.018	0.013	0.019	0.026
				$\sum_{h=1}^{H} r_{i,t+h} - r^f_{i,t+h} = a_i + b\hat{\pi}^{-}_{i,t} + u_{i,t+H}$					
b	1.00***	2.31**	2.86*	1.02*	2.13	2.32	1.70**	3.31**	3.60*
	(0.39)	(1.02)	(1.55)	(0.51)	(1.33)	(2.16)	(0.72)	(1.45)	(1.98)
N	1,753	1,707	1,661	641	601	561	1,112	1,106	1,100
R^2	0.016	0.026	0.026	0.032	0.037	0.024	0.017	0.031	0.035
				$\sum_{h=1}^{H} r_{i,t+h} - r^f_{i,t+h} = a_i + a_t + b\hat{\pi}_{i,t} + u_{i,t+H}$					
b	0.71**	1.84***	2.85***	1.07**	1.83*	2.26*	0.42	1.89*	3.21**
	(0.30)	(0.69)	(1.00)	(0.42)	(0.95)	(1.40)	(0.41)	(0.97)	(1.40)
N	1,753	1,707	1,661	641	601	561	1,112	1,106	1,100
R^2	0.012	0.029	0.046	0.034	0.073	0.094	0.011	0.039	0.067

Notes. This table presents slope coefficients, standard errors, and adjusted R^2 statistics for predictive panel regressions of cumulative excess returns on crisis probabilities. Observations are over the sample of 42 countries, from 1900 to 2020. Results are reported for the full sample, as well as for subsamples covering the pre- and post-WW2 periods. Standard errors in parentheses are bootstrapped to account for cross-sectional correlation and time aggregation. The bootstrap simulates the estimated $\hat{\pi}$ when it is used as a regressor to account for estimation error.

*, **, and *** indicate statistical significance at 10%, 5%, and 1% levels, respectively.

Table VI: Excess Return Predictability: Inspecting the Mechanism

	Full sample (1876-2020)			Prewar (1876-1945)			Postwar (1945-2020)		
H	1	3	5	1	3	5	1	3	5
	$\sum_{h=1}^{H} r_{i,t+h} - r^f_{i,t+h} = a_i + b$ Recession$_{i,t} + u_{i,t+H}$								
b	0.02	-0.06^{*}	-0.10^{*}	0.02	0.03	0.01	0.04^{*}	-0.13^{**}	-0.17^{**}
	(0.01)	(0.03)	(0.05)	(0.02)	(0.04)	(0.06)	(0.02)	(0.05)	(0.07)
N	2,129	2,079	2,031	1,017	973	931	1,112	1,106	1,100
R^2	0.002	0.004	0.008	0.003	0.001	0.000	0.004	0.013	0.015
	$\sum_{h=1}^{H} r_{i,t+h} - r^f_{i,t+h} = a_i + b\hat{\pi}_{i,t} + c$ Recession$_{i,t} + u_{i,t+H}$								
b	0.91^{**}	2.09^{**}	2.45^{*}	1.00^{**}	2.01	1.92	1.27^{*}	2.43^{*}	3.03^{*}
	(0.37)	(0.94)	(1.43)	(0.49)	(1.26)	(2.01)	(0.67)	(1.32)	(1.78)
c	-0.00	-0.09^{**}	-0.13^{**}	-0.00	0.00	-0.02	0.01	-0.13^{***}	-0.18^{**}
	(0.02)	(0.04)	(0.06)	(0.02)	(0.04)	(0.07)	(0.02)	(0.05)	(0.07)
N	1,753	1,707	1,661	641	601	561	1,112	1,106	1,100
R^2	0.014	0.028	0.032	0.032	0.036	0.019	0.014	0.034	0.043
	$\sum_{h=1}^{H} r_{i,t+h} - r^f_{i,t+h} = a_i + b$FinCrisis$_{i,t} + u_{i,t+H}$								
b	0.02	-0.17^{***}	-0.28^{***}	0.04	-0.17^{***}	-0.23^{**}	0.01	-0.16^{*}	-0.32^{**}
	(0.02)	(0.05)	(0.08)	(0.03)	(0.06)	(0.09)	(0.03)	(0.09)	(0.13)
N	2,129	2,079	2,031	1,017	973	931	1,112	1,106	1,100
R^2	0.001	0.016	0.029	0.004	0.022	0.026	0.000	0.012	0.033
	$\sum_{h=1}^{H} r_{i,t+h} - r^f_{i,t+h} = a_i + b\hat{\pi}_{i,t} + c$ FinCrisis$_{i,t} + u_{i,t+H}$								
b	0.97^{**}	2.08^{*}	2.27	0.96^{*}	1.95	1.96	2.16^{**}	2.82^{**}	2.82
	(0.44)	(1.01)	(1.50)	(0.49)	(1.34)	(2.05)	(0.97)	(1.34)	(1.81)
c	-0.02	-0.21^{***}	-0.31^{***}	0.01	-0.23^{***}	-0.31^{***}	-0.09^{*}	-0.21^{**}	-0.31^{**}
	(0.03)	(0.06)	(0.09)	(0.05)	(0.08)	(0.11)	(0.05)	(0.09)	(0.13)
N	1,753	1,707	1,661	641	601	561	1,112	1,106	1,100
R^2	0.015	0.042	0.051	0.032	0.068	0.059	0.020	0.039	0.056

Notes. This table presents slope coefficients, standard errors, and adjusted R^2 statistics for predictive panel regressions of cumulative excess returns on crisis probabilities, realized recessions, and financial crises. Observations are over the sample of 42 countries, from 1900 to 2020. Results are reported for the full sample, as well as for subsamples covering the pre- and post-WW2 periods. Standard errors in parentheses are bootstrapped to account for cross-sectional correlation, time aggregation, and generated predictor. The bootstrap simulates the estimated $\hat{\pi}$ when it is used as a regressor.
*, **, and *** indicate statistical significance at 10%, 5%, and 1% levels, respectively.

Table VII: Estimation of the Disaster Risk Model

	Estimate	Standard Error
Normal times:		
Average growth in consumption μ	0.023	0.022
Volatility of consumption growth σ	0.028	0.006
Disaster probability:		
Long-term disaster probability $\bar{\pi}$	0.034	0.044
Persistence ρ	0.738	0.017
Volatility parameter ν	0.219	0.014
Disaster size:		
Shape parameter β	0.066	0.017
Shifting parameter θ	0.056	-

Notes. The table presents GMM estimates (one-lag Newey and West (1987) standard errors in parentheses) of consumption dynamics in the time-varying disaster model. The log consumption growth evolves according to

$$\Delta c_t = \mu + \sigma\varepsilon_t + v_t,$$

where μ and σ are constants, ε_t is a standard normal random variable, $v_t = J_t \mathbf{1}_{\Delta n_t > 0}$, and Δn_t follows a Poisson distribution with time-varying probability π_t:

$$\pi_t - \bar{\pi} = \rho(\pi_{t-1} - \bar{\pi}) + \nu\sqrt{\pi_{t-1}}u_t.$$

Finally, J_t follows a shifted negative exponential distribution that takes a minimum value of $-\theta$ and parameter β, such that its mean and variance are $-(\theta + \beta)$ and β^2.

Table VIII: Asset Pricing Moments: Levered Consumption Model

	Data		Model				
	1871-2015	1946-2015	2.5%	5%	50%	95%	97.5%
Panel A: Unrestricted Model							
Average risk-free rate	2.17	1.26	1.57	1.64	1.84	2.02	2.05
Standard deviation of risk-free rate	4.51	3.14	1.14	1.18	1.46	1.84	1.92
Average excess return	6.00	6.82	5.58	5.66	6.21	6.80	6.94
Standard deviation of excess return	18.5	16.5	20.5	20.9	23.1	25.7	26.5
Average dividend yield	4.13	3.33	3.16	3.19	3.37	3.68	3.77
Standard deviation of dividend yield	1.57	1.42	0.98	1.05	1.62	3.52	4.42
Autocorrelation of dividend yield	78.4	90.0	48.3	52.7	65.0	75.6	77.8
Panel B: Restricted Model (No Realized Disasters)							
Average risk-free rate	2.17	1.26	1.57	1.61	1.84	2.03	2.05
Standard deviation of risk-free rate	4.51	3.14	1.12	1.16	1.46	1.82	1.92
Average excess return	6.00	6.82	6.85	6.94	7.44	8.06	8.18
Standard deviation of excess return	18.5	16.5	19.5	19.8	22.2	24.6	25.1
Average dividend yield	4.13	3.33	3.15	3.18	3.37	3.68	3.74
Standard deviation of dividend yield	1.57	1.42	0.98	1.03	1.61	3.32	4.11
Autocorrelation of dividend yield	78.4	90.0	50.9	53.8	65.5	75.9	77.6

Notes. Panel A reports unconditional moment statistics from S&P 500 real returns and three-month U.S. Treasury real rates, as well as percentiles of the same moments from model simulations. Panel B reports the same quantities from a restricted model in which disasters do not realize. Parameters are from Table VII. Additional parameters are $\gamma = 5, \delta = 99\%$ and $\phi = 2.6$.

Table IX: Asset Pricing Moments: Co-Integration Model

	Data		Model				
	1871-2015	1946-2015	2.5%	5%	50%	95%	97.5%
Panel A: Unrestricted Model							
Average risk-free rate	2.17	1.26	1.57	1.63	1.83	2.01	2.04
Standard deviation of risk-free rate	4.51	3.14	1.13	1.19	1.46	1.83	1.92
Average excess return	6.00	6.82	4.67	4.75	5.46	6.15	6.28
Standard deviation of excess return	18.5	16.5	19.3	19.6	21.6	23.7	24.0
Average dividend growth	3.61	6.06	1.50	1.53	1.82	2.07	2.11
Standard deviation of dividend growth	12.5	6.4	30.4	30.7	32.4	34.4	34.9
Average dividend yield	4.13	3.33	4.92	5.01	5.67	6.36	6.46
Standard deviation of dividend yield	1.57	1.42	2.56	2.70	3.39	4.78	5.42
Autocorrelation of dividend yield	78.4	90.0	67.5	70.8	78.5	83.7	84.7
Corr. div. yield and disaster probability	45.2	23.2	30.4	32.6	50.3	64.3	66.1
Panel B: Restricted Model (No Realized Disasters)							
Average risk-free rate	2.17	1.26	1.57	1.63	1.83	2.00	2.03
Standard deviation of risk-free rate	4.51	3.14	1.16	1.19	1.47	1.85	1.92
Average excess return	6.00	6.82	5.75	5.89	6.58	7.43	7.64
Standard deviation of excess return	18.5	16.5	18.8	19.0	20.8	23.1	23.6
Average dividend growth	3.61	6.06	2.09	2.11	2.30	2.50	2.53
Standard deviation of dividend growth	12.5	6.4	28.1	28.3	29.5	30.6	30.8
Average dividend yield	4.13	3.33	5.60	5.71	6.38	7.22	7.41
Standard deviation of dividend yield	1.57	1.42	2.86	3.00	3.95	6.45	7.81
Autocorrelation of dividend yield	78.4	90.0	65.3	69.3	78.5	83.7	84.4
Corr. div. yield and disaster probability	45.2	23.2	40.1	43.0	59.3	72.2	74.5

Notes. Panel A reports unconditional moment statistics from S&P 500 real returns and three-month U.S. Treasury real rates as well as percentiles of the same moments from model simulations of the co-integration model. Panel B reports the same quantities from a restricted model in which disasters do not realize. Parameters are from Table VII. Additional parameters are $\gamma = 5, \delta = 99\%, \phi = 0.90, \eta = 0.30,$ and $\kappa = 4$.

Table X: Premium and Return Variance Decomposition

	Levered Consumption	Co-Integration
Equity Premium:		
Normal Risk	0.12	0.06
Time-varying Disaster Probability	0.74	0.45
Disaster Severity	0.14	0.49
Return Variance:		
Normal Risk	0.10	0.24
Time-varying Disaster Probability	0.82	0.68
Disaster Severity	0.08	0.08

Notes. This table reports the relative contribution of the three components of the equity premium and the return variance at the steady state for both the levered consumption model and the cointegration model.

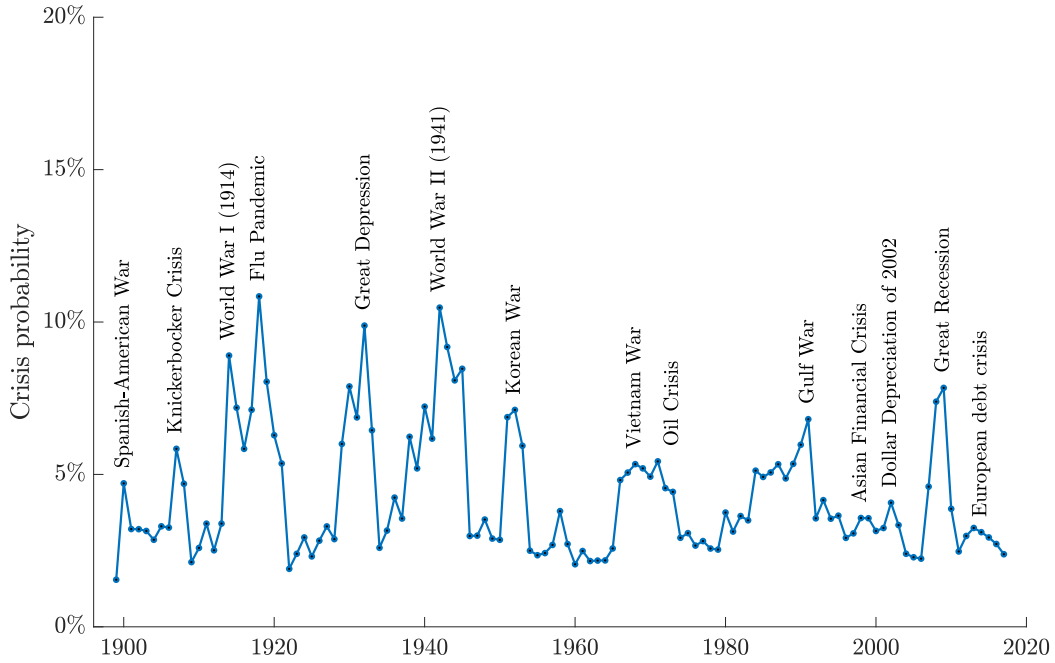

Figure 1: U.S. Macroeconomic Risk

Notes. This figure shows the probability of a severe macroeconomic crisis, $\hat{\pi}$, in the United States. A crisis is defined as a two-standard deviation drop in three-year consumption growth below the average growth rate. According to this definition, the United States experienced crises in 1918, 1919, 1930, 1931, 2007, and 2008.

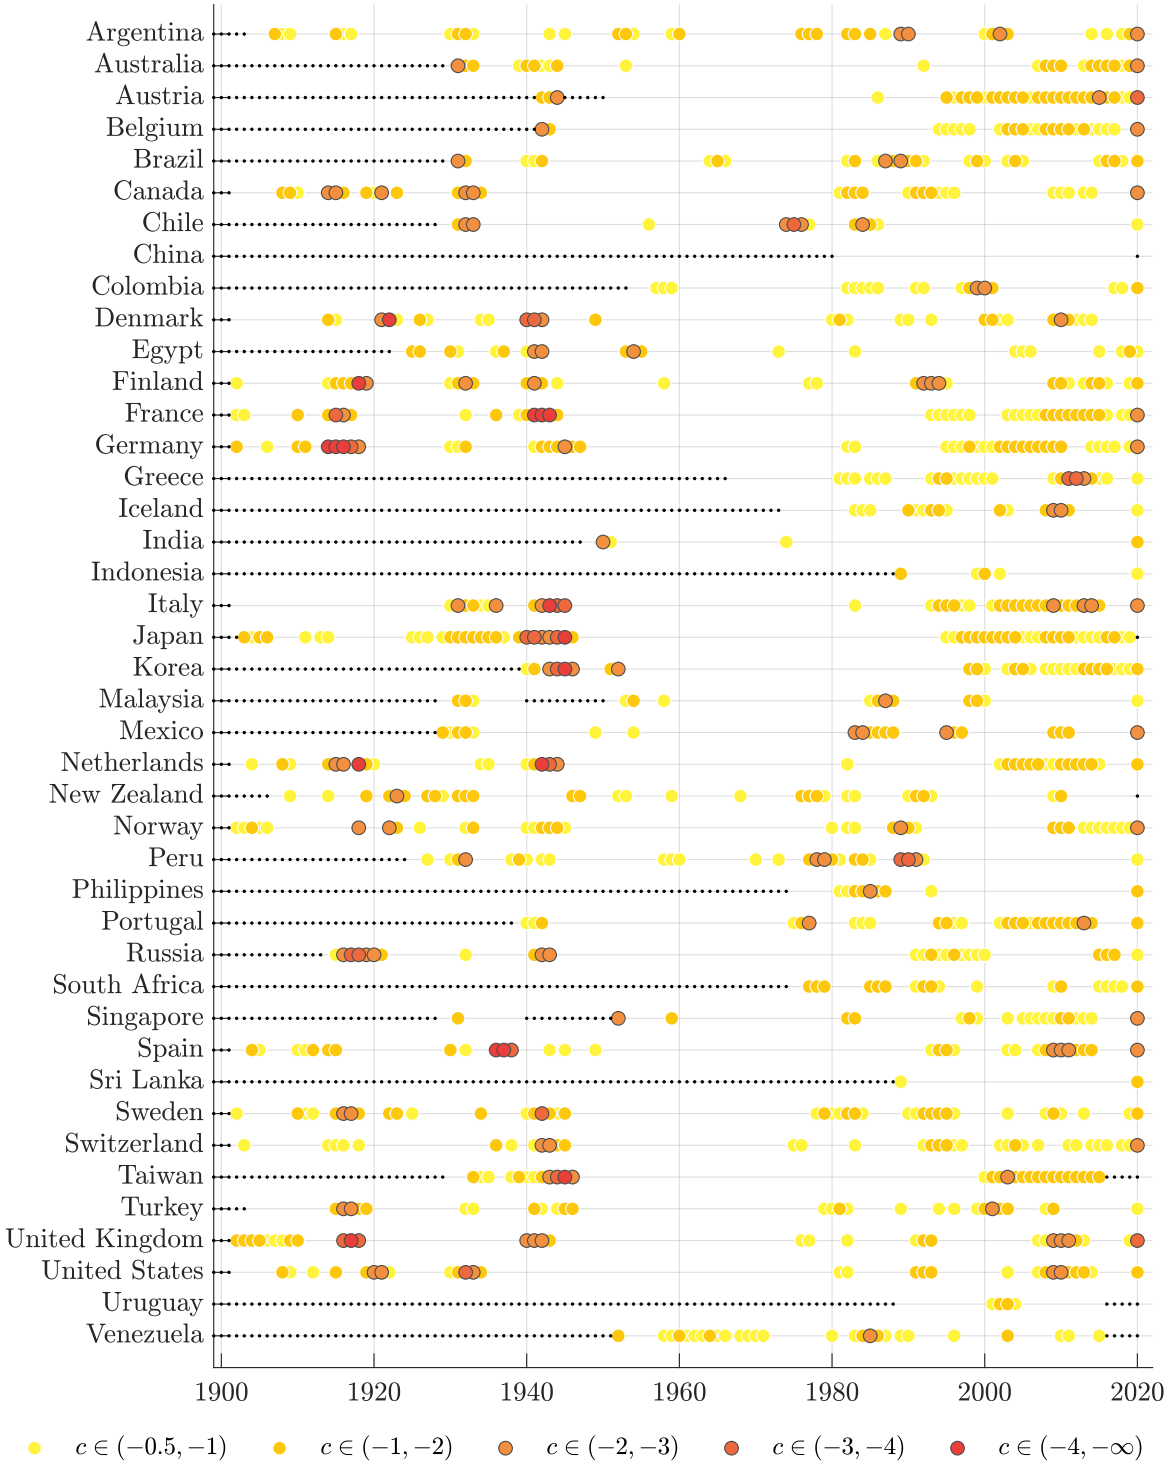

Figure 2: Timeline

Notes. This figure shows a heatmap of three-year consumption growth. Consumption growth in each country is standardized to have a zero mean and unit standard deviation. Only realizations less than -0.5 are shown. 2-SD macroeconomic crises are highlighted in grey. Small black dots indicate data unavailability.

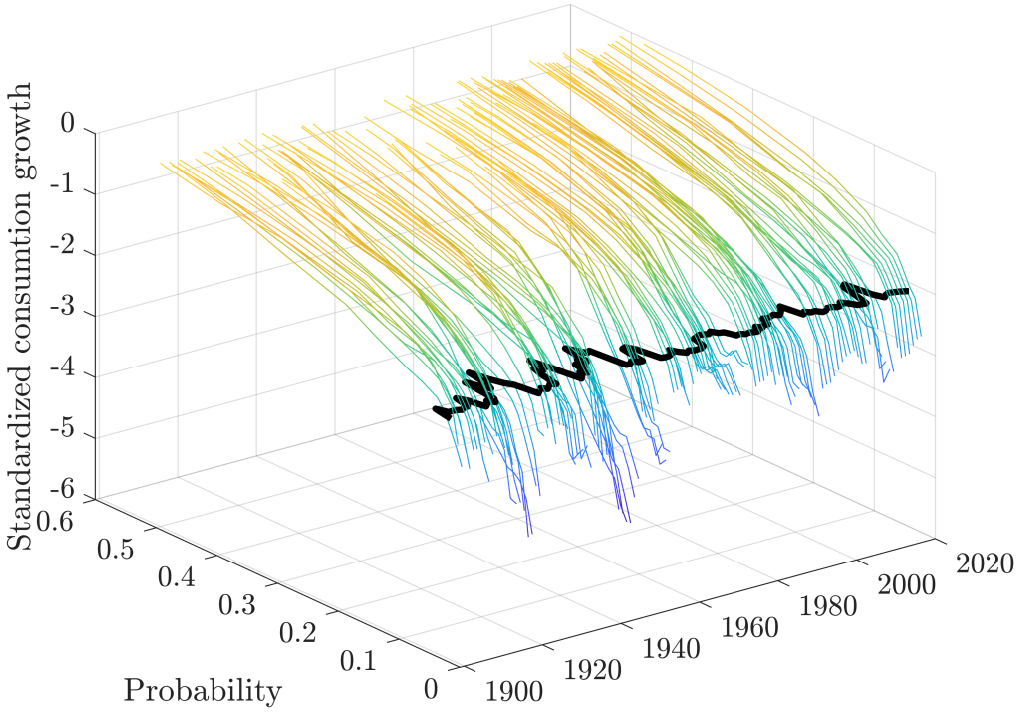

Figure 3: U.S. Consumption Growth Quantiles

Notes. This figure shows the conditional distribution of three-year consumption growth in the United States. The solid black line shows macro risk, $\hat{\pi}$, which is interpolated as the probability of a 2-SD decline in consumption and is depicted in Figure 1.

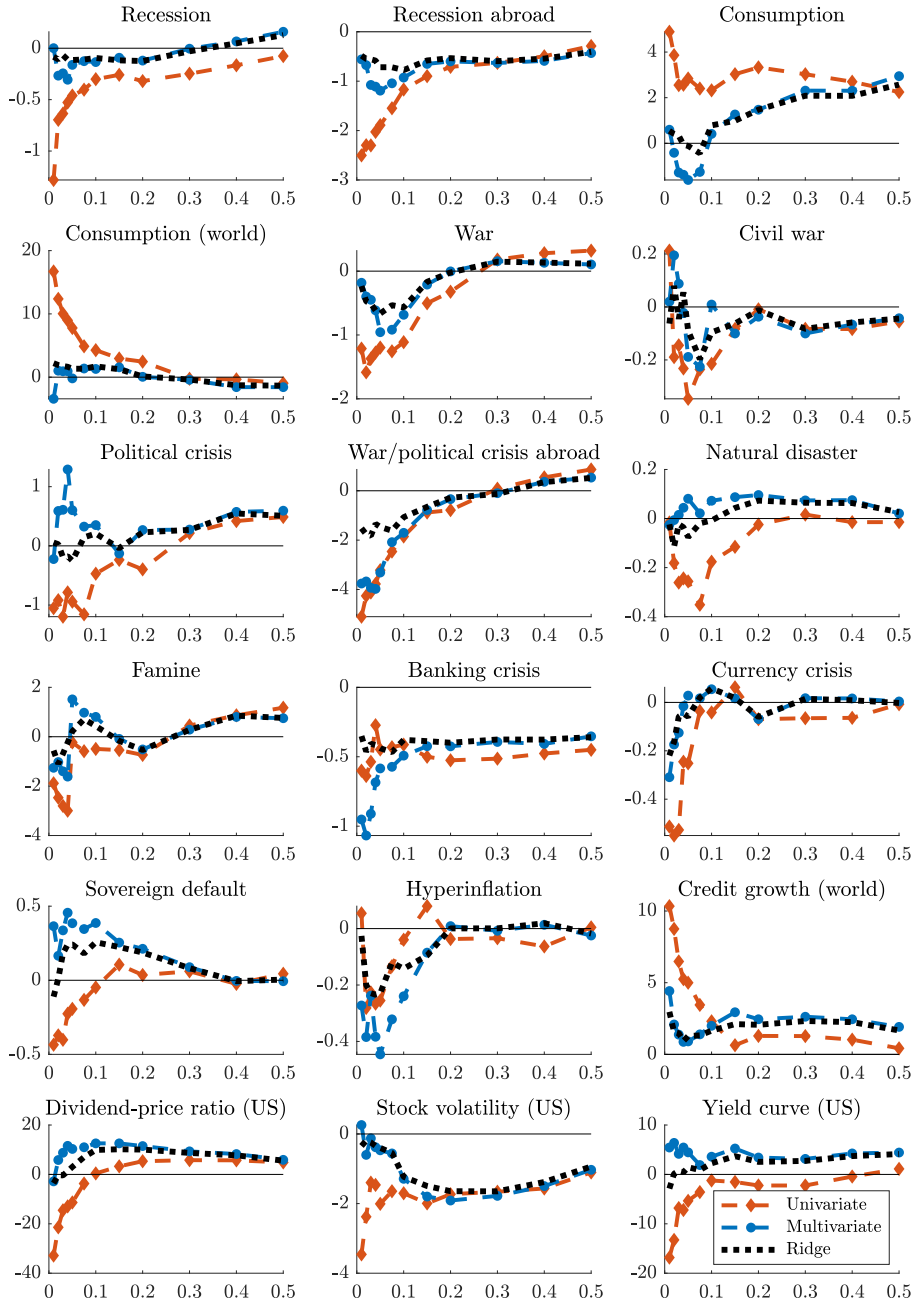

Figure 4: Quantile Regression Estimates

This figure shows quantile regression coefficient estimates for three-year consumption growth on the set of predictive variables. Each plot reports quantile slope coefficient against the .01, .02, .03, .04, .05, .075, .1, .15, .2, .3, .4, and .5 quantiles. The red dots report univariate regression estimates, together with their 95% confidence bands. The blue lines report the multivariate regression estimates. The dotted black lines report multivariate ridge estimates.

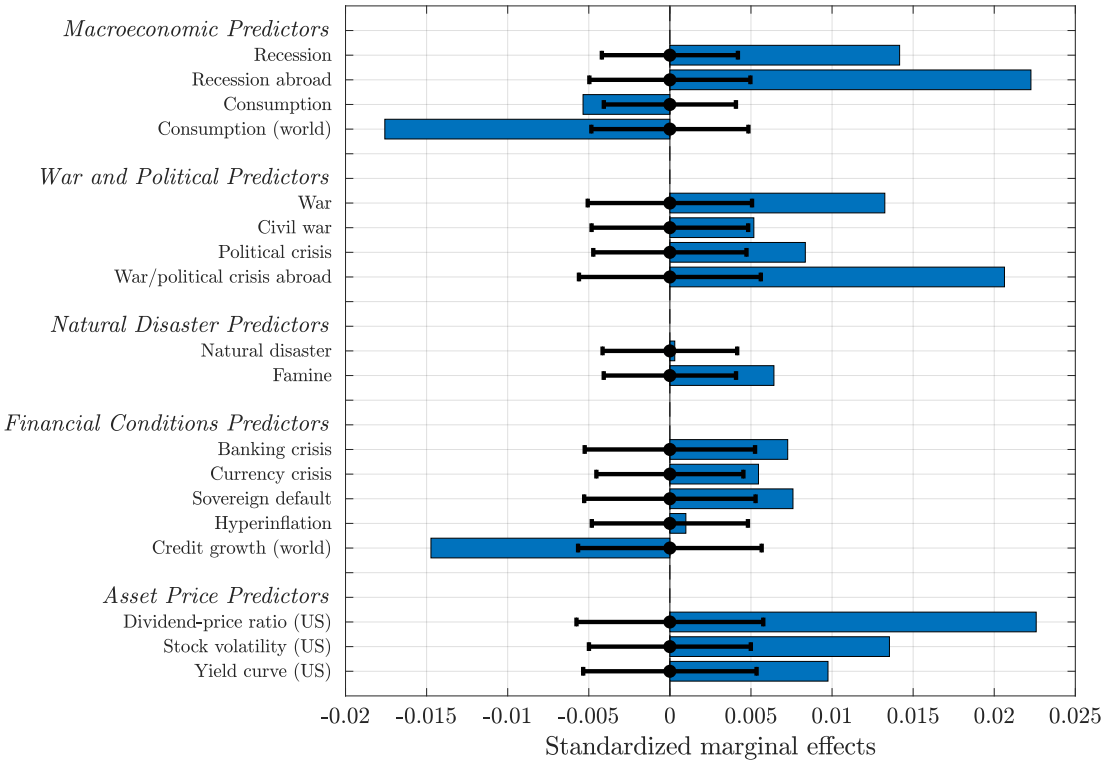

Figure 5: Marginal effects for Univariate Quantile Regressions

Notes. This figure shows marginal effects calculated from the QR model based on univariate regressions. The marginal effect is defined as the impact of a one standard deviation increase in a predictive variable on the unconditional probability of a macroeconomic crisis. The horizontal error bars depict 90% confidence intervals for the null of no predictability.

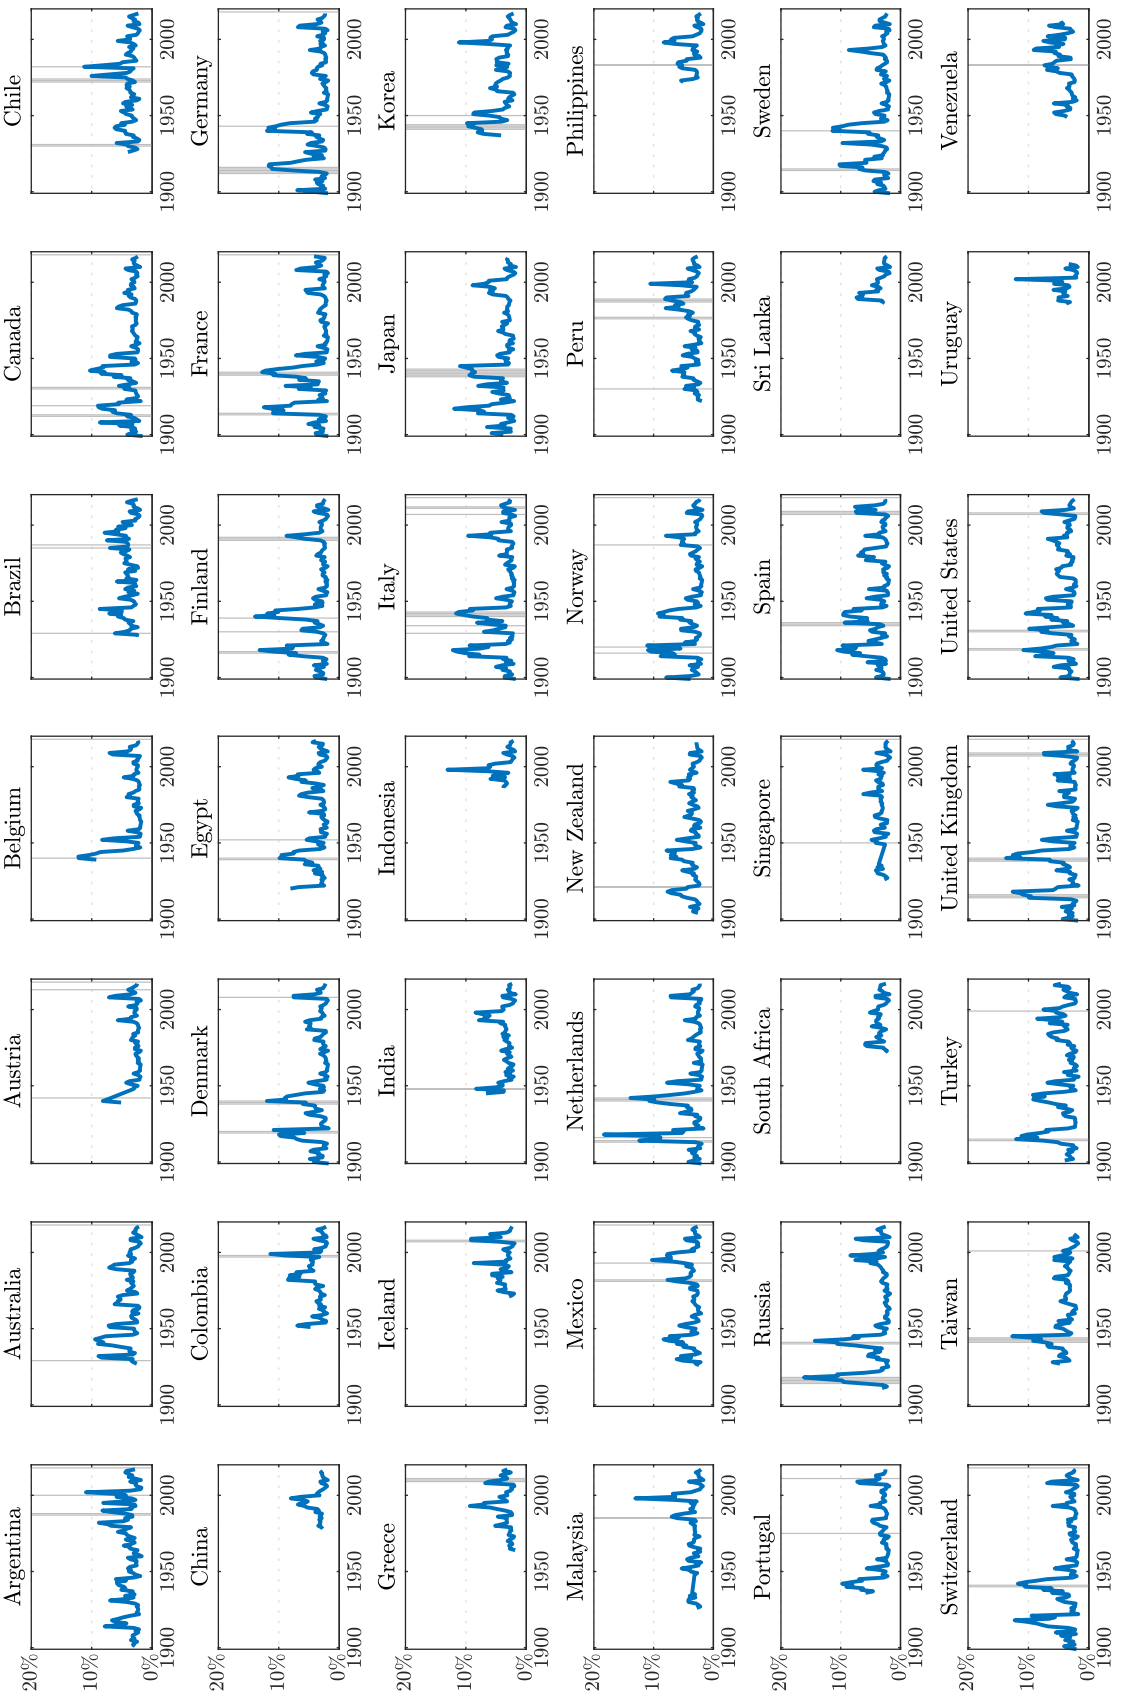

Figure 6: Country Crisis Probabilities, $\hat{\pi}_i$

Notes. The shaded areas represent realized crisis-years.

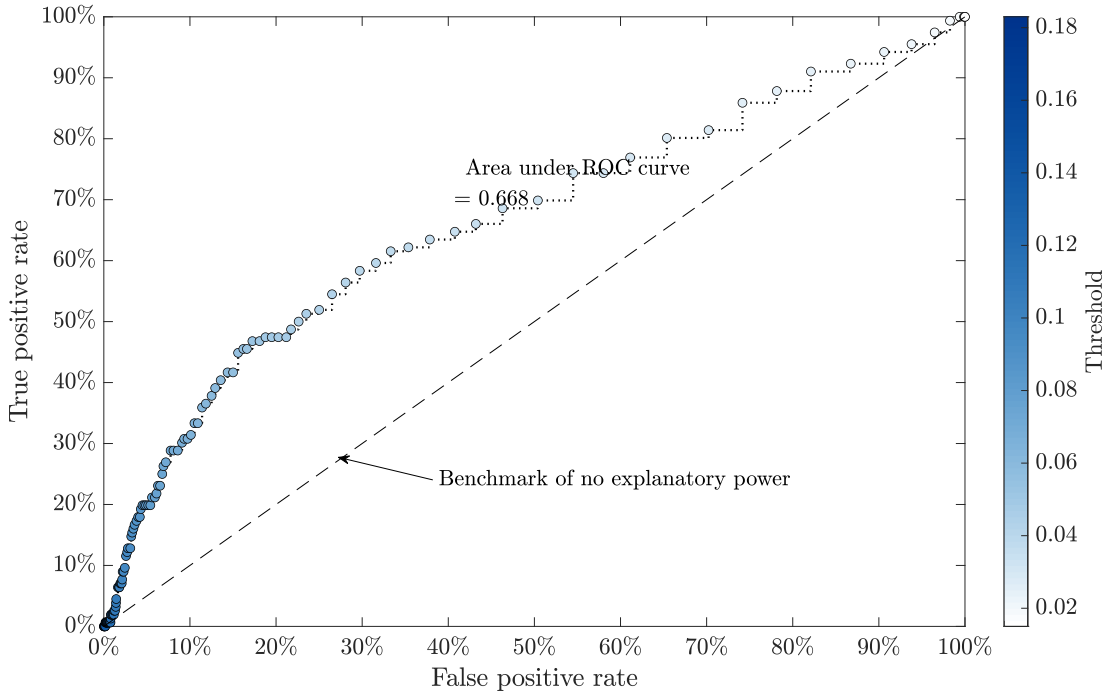

Figure 7: Receiver Operating Characteristic (ROC) Curve

Notes. The ROC curve assess the accuracy of binary forecasts (in this case in-sample $\hat{\pi}_i$ estimates). Each dot corresponds to a given threshold rule, with darker dots representing higher thresholds. The x-axis indicates how often the model wrongly predict a macroeconomic crisis in $t+1$; the y-axis gives the fraction of predicted disaster among all crises that did happen.

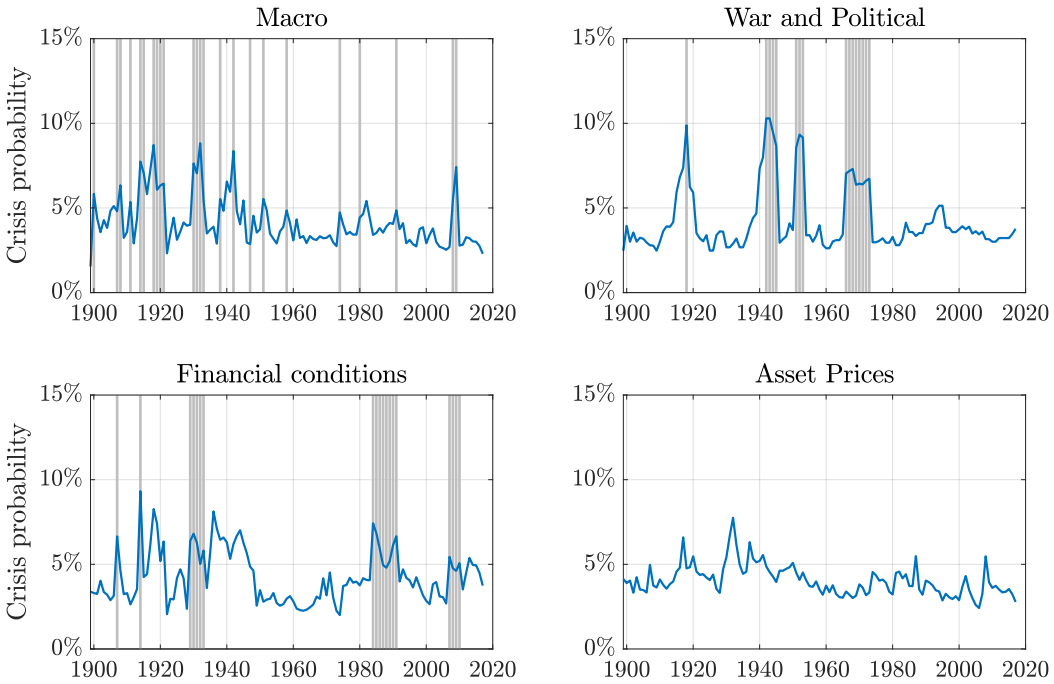

Figure 8: US Macroeconomic Risk by Predictor Categories

Notes. This figure shows macro risk estimates for four out of five predictive variables groups. The predictive variables associated with each group are listed in Table I. Shaded areas highlight important predictors in each category: recessions (Macro), wars (War and Political), and banking crises (Financial Conditions).

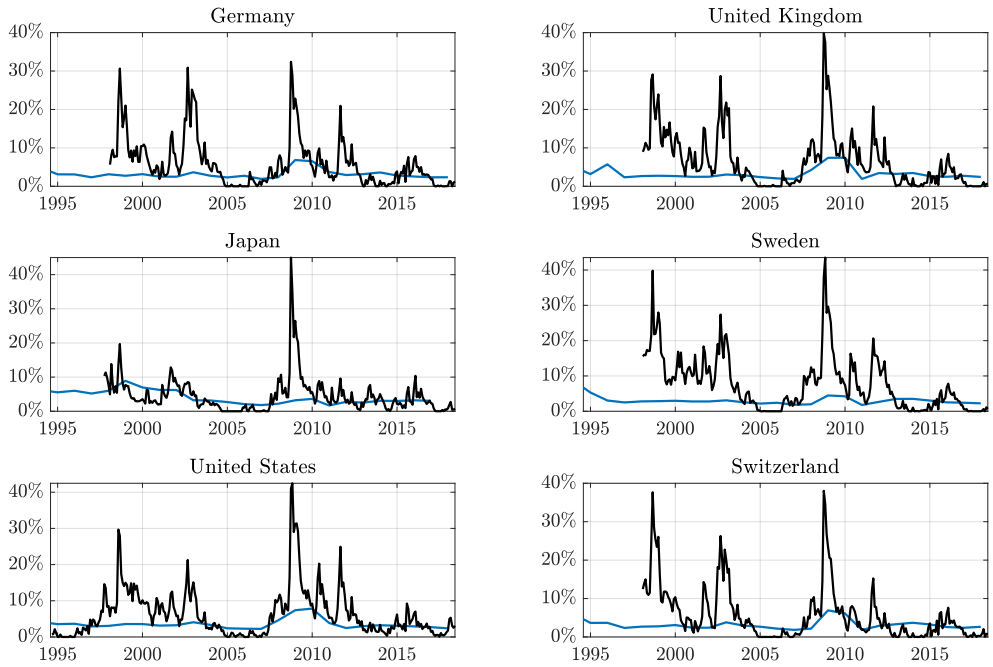

Figure 9: Comparison with Option-Based Estimates

Notes. This figure compares macro risk (blue) and option-based estimates (black lines), which are constructed following Barro and Liao (2021) using over-the-counter options prices for the following equity-market indices: S&P 500 (the United States), FTSE (the United Kingdom), DAX (Germany), Nikkei (Japan), OMX (Sweden), and SMI (Switzerland).

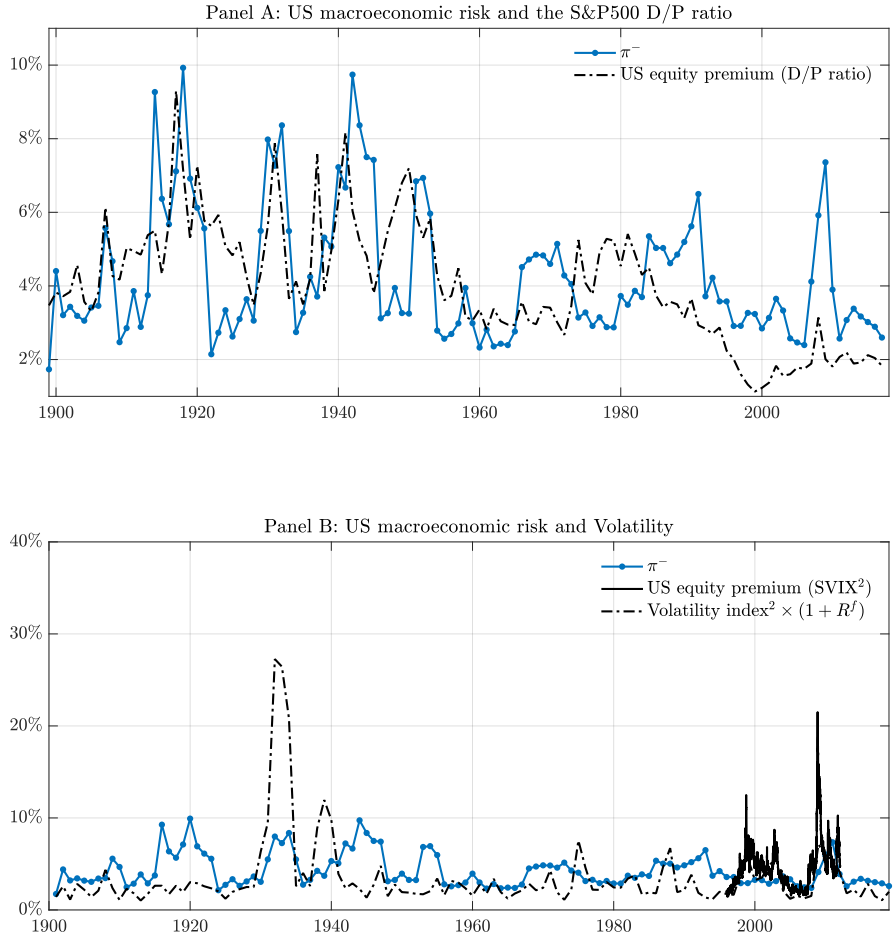

Figure 10: Macro Risk and the Equity Premium

Notes. This figure plots the U.S. macro risk estimates constructed without asset price information, $\hat{\pi}^-$, against two equity premium proxies: the S&P 500 dividend-price ratio and the Martin (2017) bound on the equity premium, $SVIX^2$. We extend the latter to our longer historical sample, substituting our realized volatility index squared for the risk-neutral volatility in the true index, multiplied by the real risk-free rate.

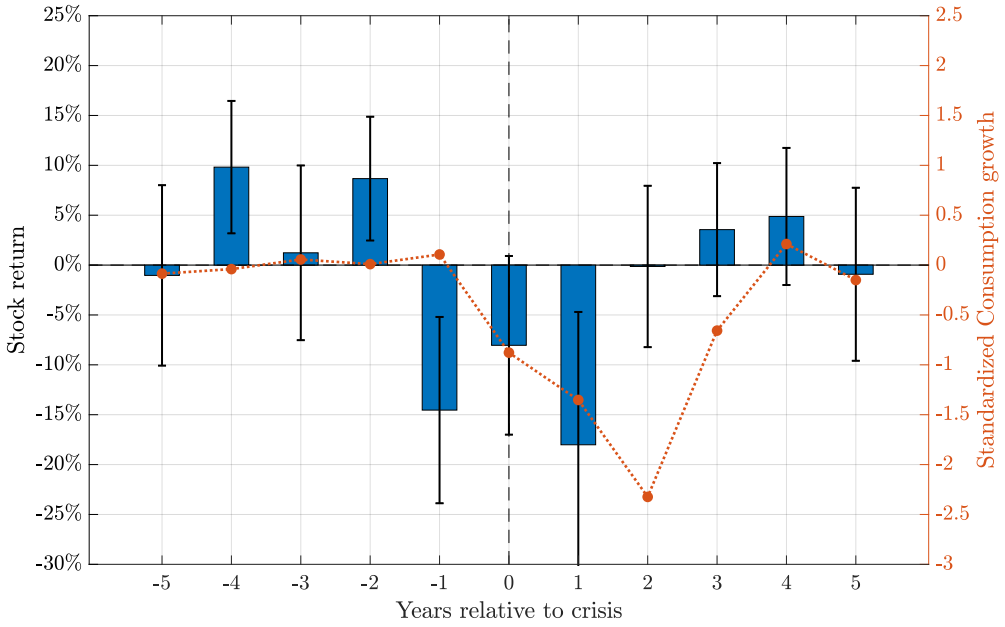

Figure 11: Macroeconomic Crises and Stock Returns

Notes. This figure reports the local currency average excess stock return and consumption growth around a macroeconomic crisis, obtained as the regression coefficients in Eq. (9).

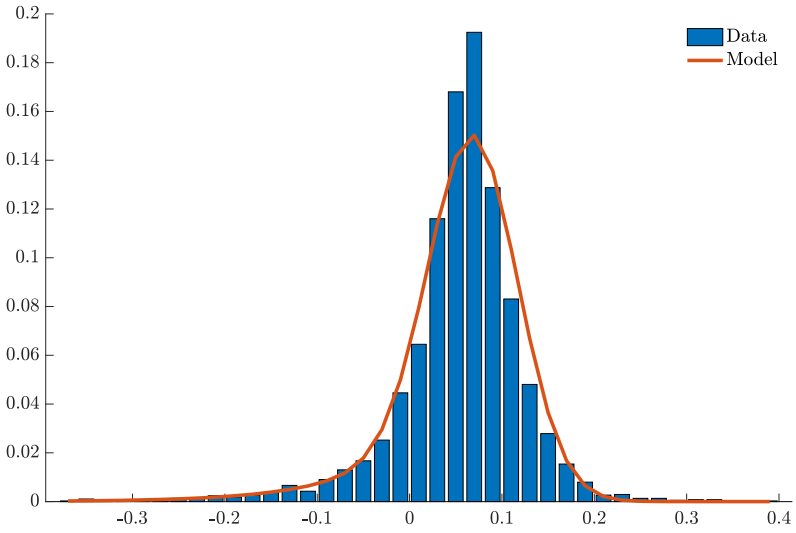

Figure 12: Consumption distribution

Notes. This figure shows the distribution of three-year consumption growth in the data together with the model-implied distribution. International consumption series are rescaled to have the same mean and standard deviation as consumption growth in the United States.

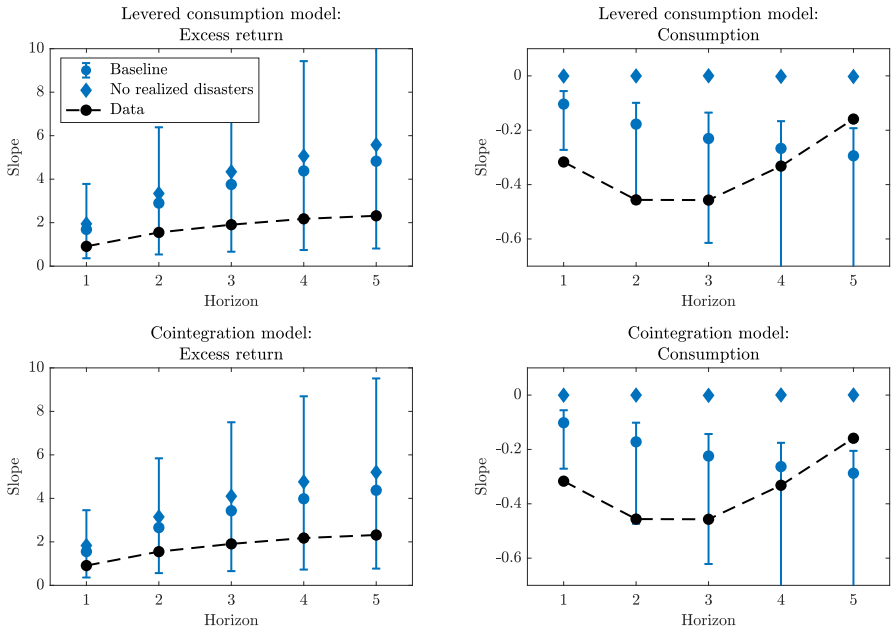

Figure 13: Model-Implied Predictability

Notes. This figure shows predictive slopes from the regression of either future cumulative excess returns (left) or consumption growth (right) on either actual or model-implied macro risk as a function of the horizon. Model-implied slopes are the median from the simulations of both the baseline model (we also report the 5th and 95th percentiles) and the restricted model in which disasters are not realized. Slopes with actual data correspond to point estimates from panel regressions (Tables IV and V).

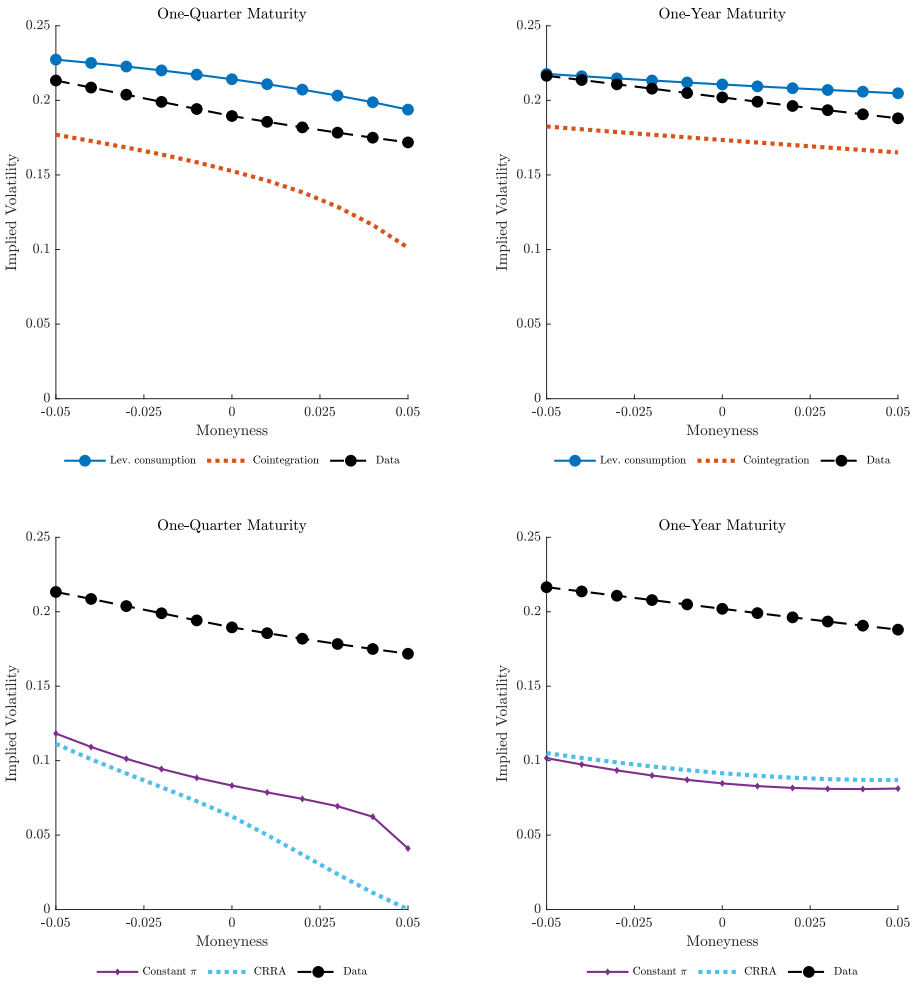

Figure 14: Model-Implied Option Prices

Notes. This figure shows the model implied volatility for put options with either one-quarter or one-year maturity as a function of relative moneyness, (strike - price) / price, in comparison with the typical evidence from S&P 500 index options. Option prices are computed by Monte Carlo simulations. The upper panels report the implied volatility for the levered consumption model and the co-integration model. The lower panels report the implied volatility for a levered consumption model with either a constant probability of disasters (and an increased relative risk aversion $\gamma = 7.5$) or CRRA preferences.

REFERENCES

Abel, Andrew B., 1999, Risk premia and term premia in general equilibrium, *Journal of Monetary Economics* 43, 3–33.

Adrian, Tobias, Nina Boyarchenko, and Domenico Giannone, 2019, Vulnerable growth, *American Economic Review* 109, 1263–1289.

Backus, David, Mikhail Chernov, and Ian Martin, 2011, Disasters Implied by Equity Index Options, *Journal of Finance* 66, 1969–2012.

Bansal, Ravi, Dana Kiku, and Amir Yaron, 2012, An empirical evaluation of the long-run risks model for asset prices, *Critical Finance Review* 1, 183–221.

Bansal, Ravi, and Amir Yaron, 2004, Risks for the long run: A potential resolution of asset pricing puzzles, *Journal of Finance* 59, 1481–1509.

Baron, Matthew, Luc A. Laeven, Julien Pénasse, and Yevhenii Usenko, 2021, Investing in Crises, ECB Woking Paper No 2548.

Baron, Matthew, and Wei Xiong, 2017, Credit expansion and neglected crash risk, *Quarterly Journal of Economics* 132, 713–764.

Barro, Robert J., 2006, Rare Disasters and Asset Markets in the Twentieth Century, *Quarterly Journal of Economics* 121, 823–866.

Barro, Robert J., and Tao Jin, 2011, On the Size Distribution of Macroeconomic Disasters, *Econometrica* 79, 1567–1589.

Barro, Robert J., and Gordon Y. Liao, 2021, Rare disaster probability and options pricing, *Journal of Financial Economics* 139, 750–769.

Barro, Robert J., and José F. Ursúa, 2008, Macroeconomic Crises since 1870, *Brookings Papers on Economic Activity* 2008, 255–350.

Barro, Robert J., and José F. Ursúa, Barro-Ursúa Macroeconomic Data, 2010.

Barro, Robert J., and José F. Ursúa, 2017, Stock-market crashes and depressions, *Research in Economics* 71, 384–398.

Beeler, Jason, and John Y. Campbell, 2012, The long-run risks model and aggregate asset prices: an empirical assessment, *Critical Finance Review* 1, 141–182.

Belo, Frederico, Pierre Collin-Dufresne, and Robert S. Goldstein, 2015, Dividend Dynamics and the Term Structure of Dividend Strips, *Journal of Finance* 70, 1115–1160.

Belsley, David A., Edwin. Kuh, and Roy E. Welsch, 2004, *Regression diagnostics: identifying influential data and sources of collinearity* (Wiley, New York).

Berkman, Henk, Ben Jacobsen, and John B. Lee, 2011, Time-Varying Rare Disaster Risk and Stock Returns, *Journal of Financial Economics* 101, 313–332.

Bloom, Nicholas, 2009, The Impact of Uncertainty Shocks, *Econometrica* 77, 623–685.

Bollerslev, Tim, and Viktor Todorov, 2011, Tails, Fears, and Risk Premia, *Journal of Finance* 66, 2165–2211.

Branger, Nicole, Holger Kraft, and Christoph Meinerding, 2016, The Dynamics of Crises and the Equity Premium, *Review of Financial Studies* 29, 232–270.

Campbell, John Y., 2003, Consumption-based asset pricing, in George M. Constantinides, Milton Harris, and Rene M. Stulz, eds., *Handbook of the Economics of Finance*, 803–888 (North-Holland, Armsterdam).

Campbell, John Y., 2017, *Financial decisions and markets: a course in asset pricing* (Princeton University Press).

Campbell, John Y., and John H. Cochrane, 1999, By force of habit: A consumption-based explanation of aggregate stock market behavior, *Journal of Political Economy* 107, 205–251.

Campbell, John Y., Andrew W. Lo, and A. Craig MacKinlay, 1997, *The econometrics of financial markets* (Princeton University Press).

Chen, Hui, Winston Wei Dou, and Leonid Kogan, 2022, Measuring the 'Dark Matter' in Asset Pricing Models, *Journal of Finance* (Forthcoming).

Cheng, Xu, Winston Wei Dou, and Zhipeng Liao, 2022, Macro-Finance Decoupling: Robust Evaluations of Macro Asset Pricing Models, *Econometrica* 90, 685–713.

Cochrane, John H., 2011, Presidential address: Discount rates, *Journal of Finance* 66, 1047–1108.

Cochrane, John H., 2017, Macro-finance, *Review of Finance* 21, 945–985.

Collin-Dufresne, Pierre, Michael Johannes, and Lars A. Lochstoer, 2016, Parameter learning in general equilibrium: The asset pricing implications, *American Economic Review* 106, 664–698.

Constantinides, George M., 2008, Discussion of "Macroeconomic crises since 1870", *Brookings Papers on Economic Activity* 336–350.

Donaldson, John, and Rajnish Mehra, 2008, Risk-Based Explanations of the Equity Premium, in Rajnish Mehra, ed., *Handbook of the Equity Risk Premium*, 37–99 (Elsevier, Amsterdam).

Drechsler, Itamar, and Amir Yaron, 2011, What's vol got to do with it, *Review of Financial Studies* 24, 1–45.

Epstein, Larry G., and Stanley E. Zin, 1989, Substitution, Risk Aversion, and the Temporal Behavior of Consumption Growth and Asset Returns: A Theoretical Framework, *Econometrica* 57, 937–969.

Estrella, Arturo, and Frederic S Mishkin, 1998, Predicting U.S. Recessions : Financial Variables as Leading Indicators, *Review of Economics and Statistics* 80, 45–61.

Farhi, Emmanuel, and Xavier Gabaix, 2016, Rare Disasters and Exchange Rates, *Quarterly Journal of Economics* 1–52.

Fasano, G., and A. Franceschini, 1987, A multidimensional version of the Kolmogorov–Smirnov test, *Monthly Notices of the Royal Astronomical Society* 225, 155–170.

Gabaix, Xavier, 2011, Disasterization: A simple way to fix the asset pricing properties of macroeconomic models, *American Economic Review* 101, 406–409.

Gabaix, Xavier, 2012, Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance, *Quarterly Journal of Economics* 127, 645–700.

Garcia-Ferrer, A., R. A. Highfield, F. Palm, and A. Zellner, 1987, Macroeconomic forecasting using pooled international data, *Journal of Business and Economic Statistics* 5, 53–67.

Ghaderi, Mohammad, Mete Kilic, and Sang Byung Seo, 2022, Learning, slowly unfolding disasters, and asset prices, *Journal of Financial Economics* 143, 527–549.

Giglio, Stefano, Bryan Kelly, and Seth Pruitt, 2016, Systemic risk and the macroeconomy: An empirical evaluation, *Journal of Financial Economics* 119, 457–471.

Gourio, Francois, 2008, Disasters and recoveries, *American Economic Review* 98, 68–73.

Gourio, Francois, 2012, Disaster Risk and Business Cycles, *American Economic Review* 102, 2734–2766.

Gourio, Francois, 2013, Credit Risk and Disaster Risk, *American Economic Journal: Macroeconomics* 5, 1–34.

Hansen, Lars Peter, 2008, Generalized Method of Moments Estimation, *The New Palgrave Dictionary of Economics* .

Harvey, Campbell R., 1988, The real term structure and consumption growth, *Journal of Financial Economics* 22, 305–333.

Hasler, Michael, and Roberto Marfè, 2016, Disaster Recovery and the Term Structure of Divi-

dend Strips, *Journal of Financial Economics* 122, 116–134.

Isoré, Marlène, and Urszula Szczerbowicz, 2017, Disaster risk and preference shifts in a New Keynesian model, *Journal of Economic Dynamics and Control* 79, 97–125.

Johannes, Michael, Lars Lochstoer, and Yiqun Mou, 2016, Learning About Consumption Dynamics, *Journal of Finance* 71, 551–600.

Jordà, Òscar, Moritz Schularick, and Alan M. Taylor, 2017, Macrofinancial history and the new business cycle facts, *NBER Macroeconomics Annual* 31, 213–263.

Julliard, Christian, and Anisha Ghosh, 2012, Can rare events explain the equity premium puzzle?, *Review of Financial Studies* 25, 3037–3076.

Kilic, Mete, and Jessica A. Wachter, 2018, Risk, unemployment, and the stock market: A rare-event-based explanation of labor market volatility, *Review of Financial Studies* 31, 4762–4814.

Koenker, Roger, and Gilbert Bassett, 1978, Regression Quantiles, *Econometrica* 46, 33.

Koulovatianos, Christos, and Volker Wieland, 2011, Asset Pricing under Rational Learning about Rare Disasters, CEPR Discussion Paper No. DP8514.

Lettau, Martin, Sydney C. Ludvigson, and Jessica A. Wachter, 2008, The Declining Equity Premium: What Role Does Macroeconomic Risk Play?, *Review of Financial Studies* 21, 1653–1687.

Lewis, Karen K., and Edith X. Liu, 2016, Disaster Risk and Asset Returns: An International Perspective, *Journal of International Economics* .

Longstaff, Francis A., and Monika Piazzesi, 2004, Corporate earnings and the equity premium, *Journal of Financial Economics* 74, 401–421.

Lucas, Robert Jr., 1978, Asset prices in an exchange economy, *Econometrica* 46, 1429–1445.

Manela, Asaf, and Alan Moreira, 2017, News implied volatility and disaster concerns, *Journal of Financial Economics* 123, 137–162.

Marfè, Roberto, 2017, Income Insurance and the Equilibrium Term Structure of Equity, *Journal of Finance* 72, 2073–2130.

Martin, Ian, 2013a, The Lucas Orchard, *Econometrica* 81, 55–111.

Martin, Ian, 2017, What is the Expected Return on the Market?, *Quarterly Journal of Economics* 132, 367–433.

Martin, Ian W R, 2013b, Consumption-Based asset pricing with higher cumulants, *Review of Economic Studies* 80, 745–773.

Mehra, Rajnish, and Edward C. Prescott, 1985, The equity premium: A puzzle, *Journal of monetary Economics* 15, 145–161.

Mian, Atif R., Amir Sufi, and Emil Verner, 2017, Household debt and business cycles worldwide, *Quarterly Journal of Economics* 1755–1817.

Nakamura, Emi, Jón Steinsson, Robert J. Barro, and José F. Ursúa, 2013, Crises and recoveries in an empirical model of consumption disasters, *American Economic Journal: Macroeconomics* 5, 35–74.

Newey, Whitney K., and Kenneth D. West, 1987, A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix, *Econometrica* 55, 703–708.

Orlik, Anna, and Laura Veldkamp, 2014, Understanding Uncertainty Shocks and the Role of Black Swans, NBER Working Paper No. 20445.

Reinhart, Carmen M., and Kenneth S. Rogoff, 2009, *This Time is Different : A Panoramic View of Eight Centuries of* (Princeton University Press, Princeton, NJ).

Rietz, Thomas A., 1988, The equity risk premium: A solution, *Journal of Monetary Economics* 22, 117–131.

Sarkees, Meredith Reid, and Frank Wayman, 2010, *Resort to War: 1816 - 2007* (CQ Press, Washington DC).

Schwert, G. William, 1989, Why Does Stock Market Volatility Change Over Time?, *Journal of Finance* 44, 1115–1153.

Seo, Sang Byung, and Jessica A. Wachter, 2019, Option prices in a model with stochastic disaster risk, *Management Science* 65, 3449–3469.

Siriwardane, Emil, 2015, The Probability of Rare Disasters : Estimation and Implications, Harvard Business School Working Paper, No. 16-061.

Tibshirani, Robert, 1996, Regression Shrinkage and Selection Via the Lasso, *Journal of the Royal Statistical Society: Series B (Methodological)* 58, 267–288.

van Binsbergen, Jules H., and Ralph S. J. Koijen, 2010, Predictive regressions: A present-value approach, *Journal of Finance* 65, 1439–1471.

Wachter, Jessica A., 2013, Can Time-Varying Risk of Rare Disasters Explain Aggregate Stock Market Volatility?, *Journal of Finance* 68, 987–1035.

Weitzman, Martin L., 2007, Subjective Expectations and Asset-Return Puzzles, *American Economic Review* 97, 1102–1130.

Zou, Hui, and Trevor Hastie, 2005, Regularization and variable selection via the elastic net, *Journal of the Royal Statistical Society. Series B: Statistical Methodology* 67, 301–320.

Online Appendix for Measuring Macroeconomic Tail Risk

Roberto Marfè, Julien Pénasse

January 17, 2024

Appendix A. Data Appendix

Geographic distances

Our "Disaster abroad" variable is distance-weighted. We measure the distance between two countries as the population-weighted average distance in kilometers between large cities' of each country pair (Mayer and Zignago, 2011).

GDP

We primarily work with consumption and GDP expressed per capita, but we sometimes need GDP levels to value-weight series. We thus obtain population data from the Maddison database (Inklaar et al., 2018), with the exception of pre-1950 Icelandic population data, which we downloaded from populstat.info (accessed on November 29, 2019). We interpolate the remaining small portion of the data that is missing (Shape-preserving piecewise cubic interpolation).

Wars.

The data for war and civil wars come from Sarkees and Wayman (2010), which extends from 1816 to 2007. "Wars" correspond to the events listed as interstate wars, while "Civil wars" correspond to the events classified as either intrastate or nonstate wars. We use Wikipedia pages listing wars involving each country in our sample to extend the data from 2008 to 2015.[1]

Political Crises.

We use data from the Center for Systemic Peace (CSP): Integrated Network for Societal Conflict Research. The *constraint on the executive* variable is constructed by the Polity IV project by coding the authority characteristics of governments across the globe from 1800 to 2016. The scale ranges from 1 (weak constraints on executive power) to 7 (strong constraints on executive power). The variable measures the extent of institutional constraints on executive power. We define political crises as a four or more decline in constraints on executive, or a collapse or interruption in authority (coded separately as "-66" and "-77" in the Polity IV data set).

Sovereign Defaults.

We utilize external default dummies assembled by Carmen Reinhart and Kenneth Rogoff (RR, see Reinhart and Rogoff (2009)). The data cover the period 1800–2014. We use Wikipedia's list of sovereign debt crises to complete our data up to 2015.[2]

Hyperinflation.

Our hyperinflation dates primarily come from RR, who define inflation crises as periods with annual inflation rate exceeding 20%. The data cover the period 1800–2010. We extend the sample to 2015 using inflation data from the World Development Indicators as well as Global Financial Data.

Currency crises.

We use currency crises dates from RR, which are defined as annual depreciation against the U.S. dollar (or the relevant anchor currency) of 15% or more. The data cover the period 1800–2010; we use exchange rate data from WDI and the Federal Reserve Bank of St. Louis to calculate crisis dates until 2015.

[1] For example, https://en.wikipedia.org/wiki/List_of_wars_involving_Argentina (as of November 28, 2017). This yields a single new civil war: the Sinai Insurgency (Egypt, 2011–2015), and extended periods for two other civil wars, namely, the Second Sri Lanka Tamil (Sri Lanka, 2006–2009) and unrest in Colombia (1989–2015).

[2] Accessed on November 28, 2017, at https://en.wikipedia.org/wiki/List_of_sovereign_debt_crises. In our sample of countries, we find one default in 2015 (Greece).

Financial Crises.

We use banking crisis dates constructed by RR, which covers the period 1800–2010. Banking crises correspond to events characterized by substantial bank runs and closure, merging, takeover, or large-scale government assistance of an important financial institution. We use Wikipedia's list of banking crises to complete our data to 2015.[3]

Natural Disasters.

We obtain data on international earthquakes, tsunamis, and volcano eruptions from the National Centers for Environmental Information. We focus on major events, the highest grade in event classifications. This corresponds to events that have caused at least 1,000 deaths, destroyed at least 1,000 houses or equivalent damages in monetary terms. When the database does not report precise numbers, we use the highest descriptive score (4), which corresponds to natural disasters of commensurable intensity. We also obtain the dates and location of major famines from the website Our World in Data (accessed on November 11, 2019).

Financial data.

The data come from Global Financial Data (GFD). We compute local currency excess returns by subtracting the continuously compounded 3-month short-term interest rate from the total equity return (in logs). Returns are calculated from the end of the previous year to the end of the current year. To eliminate the influence of outliers, we drop the 99.5th and 0.5th percentiles of stock excess returns. The former is GFD's country Treasury-bill yields. Our world variables are based on that U.S. series, which has the longest time coverage. Our baseline global dividend-price ratio is the Standard & Poor's 500 D/P ratio computed by GFD, which starts in 1871. Our volatility index is a composite of multiple time series. From 1986 onward, we use the CBOE VXO index of percentage implied volatility, on a hypothetical at the money S&P100 option 30 days to expiration. Following Bloom (2009), we use actual returns volatilities to extend the volatility series back in time. We estimate the annual standard deviation of stock returns using the daily returns to the Standard and Poor's S&P 500 from 1928 through 1985. The estimates from 1885 through 1927 use daily returns on the Dow Jones composite portfolio, as in Schwert (1989). Finally, since daily stock return data are not available before 1885, we compute the volatility through Schwert (1989)'s generalization of the Officer (1973) moving standard deviation estimator using monthly S&P 500 returns. The volatility indices are normalized to have the same mean and variance as the VXO index when they overlap from 1986 onward. Finally, we use the one-year interest rates in the United States from Robert Shiller to proxy for the short-term risk-free rate. The term spread subtracts the one-year interest rate to the 10-year Treasury bond yield.

[3] Accessed on November 28, 2017, at https://en.wikipedia.org/wiki/List_of_banking_crises. In our sample of countries, we find one default in 2015 (Greece). We find no evidence of sovereign defaults for Iceland and Switzerland over our sample period. We extend the Spanish and Icelandic financial crises from 2008 to 2012.

Table A.I: Predictor Variables: Correlation Matrix

	Recession	Recession abroad	Consumption	Consumption (world)	War	Civil war	Political crisis	War/political crisis abroad	Natural disaster	Famine	Banking crisis	Currency crisis	Sovereign default	Hyperinflation	Credit growth (world)	Dividend-price ratio (US)	Stock volatility (US)	Yield curve (US)
Recession	1.00																	
Recession abroad	0.32	1.00																
Consumption	-0.64	-0.27	1.00															
Consumption (world)	-0.19	-0.41	0.20	1.00														
War	0.09	0.17	-0.11	-0.19	1.00													
Civil war	-0.01	-0.05	0.01	-0.00	0.04	1.00												
Political crisis	0.10	0.19	-0.13	-0.09	0.09	0.04	1.00											
War/political crisis abroad	0.13	0.38	-0.13	-0.38	0.35	0.02	0.27	1.00										
Natural disaster	-0.01	0.01	0.01	-0.16	0.07	0.08	0.01	0.14	1.00									
Famine	0.05	0.05	-0.03	-0.02	0.13	0.03	0.07	0.10	0.06	1.00								
Banking crisis	0.09	0.08	-0.08	0.01	-0.06	0.08	-0.02	-0.05	0.04	-0.03	1.00							
Currency crisis	0.09	0.06	-0.06	-0.01	0.02	0.09	0.07	0.03	0.07	0.08	0.09	1.00						
Sovereign default	0.08	0.12	-0.05	-0.07	0.07	0.11	0.13	0.11	0.02	0.23	0.02	0.23	1.00					
Hyperinflation	0.08	0.09	-0.04	-0.08	0.05	0.10	0.06	0.13	0.08	0.11	0.06	0.45	0.26	1.00				
Credit growth (world)	-0.08	-0.23	0.07	0.29	-0.18	-0.02	-0.14	-0.47	-0.17	-0.08	0.07	-0.03	-0.09	-0.07	1.00			
Dividend-price ratio (US)	0.20	0.40	-0.09	-0.43	0.17	-0.05	0.11	0.29	0.03	0.05	-0.03	0.04	0.11	0.08	-0.24	1.00		
Stock volatility (US)	0.10	0.23	-0.07	-0.21	-0.02	-0.01	-0.02	-0.03	0.04	-0.03	0.10	0.04	0.05	0.01	0.08	0.17	1.00	
Yield curve (US)	0.01	0.18	-0.03	0.09	0.02	0.02	0.04	0.15	-0.03	0.03	0.04	0.04	0.09	-0.02	-0.17	-0.19	0.11	1.00

Notes. The average absolute correlation between variables equals 0.11

Table A.II: Predictor Variables: Pairwise Correlations Across Countries

	Mean	Min	25%	75%	Max
Recession	0.09	-0.27	0.01	0.19	0.53
Recession abroad	0.74	0.21	0.64	0.90	0.99
Consumption	0.08	-0.54	-0.05	0.20	0.78
Consumption (world)	1.00	1.00	1.00	1.00	1.00
War	0.22	-0.13	-0.03	0.42	1.00
Civil war	0.03	-0.51	-0.07	0.09	1.00
Political crisis	0.12	-0.07	-0.03	-0.01	1.00
War/political crisis abroad	0.78	0.02	0.70	0.92	1.00
Natural disaster	0.46	-0.24	0.20	0.77	1.00
Famine	0.21	-0.05	-0.03	0.55	1.00
Banking crisis	0.13	-0.23	-0.06	0.27	0.89
Currency crisis	0.07	-0.25	-0.05	0.17	0.66
Sovereign default	0.12	-0.33	-0.09	0.30	0.82
Hyperinflation	0.10	-0.35	-0.04	0.21	1.00
Credit growth (world)	1.00	1.00	1.00	1.00	1.00
Dividend-price ratio (US)	1.00	1.00	1.00	1.00	1.00
Stock volatility (US)	1.00	1.00	1.00	1.00	1.00
Yield curve (US)	1.00	1.00	1.00	1.00	1.00

Appendix B. Quantile regression approach

We search for the vector β that minimizes the quantile loss function

$$\hat{B} = \arg\min_{B} \sum_{q=Q} \sum_{t=1}^{T-h} \rho_{\tau_q}(\Delta c_{i,t+h} - X_{i,t}B) + \lambda_1 \sum_{j=1}^{N} |b_j| + \lambda_2 \sum_{j=1}^{N} b_j^2, \tag{B1}$$

where $\rho_{\tau_q}(x) = x\,(\tau_q - I_{x<0})$ is the quantile loss function and λ_1 and λ_2 are parameters. These parameters penalize the number of parameters and thus allow us to include shrinkage in the quantile estimation.

Setting $\lambda_1 = 0$ and $\lambda_2 = 0$ yields the standard quantile regression (QR) estimator of B. Setting $\lambda_1 \neq 0$ and $\lambda_2 = 0$ yields the quantile equivalent of the LASSO estimator ('least absolute shrinkage selection operator') of Tibshirani (1996). As is well known, the LASSO estimator imposes a sparsity assumption in that it uses a L_1 penalty, leading some coefficient estimates to be exactly zero for large enough values of λ_1. Setting $\lambda_1 = 0$ and $\lambda_2 \neq 0$ yields the ridge estimator that shrinks the estimates of b_j towards zero. As the ridge estimator uses a L_2 penalty, the ridge estimates will almost never be zero exactly. As a result, uninformative predictors can still inflate the forecast error variance. However, the ridge tends to dominate the LASSO in settings with many correlated regressors. Finally, setting $\lambda_1 \neq 0$ and $\lambda_2 \neq 0$ yields the elastic net estimator of Zou and Hastie (2005).

We estimate $\hat{B}$ using the .01, .02, .03, .04, .05, .075, .1, .15, .2, .3, .4 and .5 quantiles. These quantiles always cover the 2-SD threshold in our sample. We use ten-fold cross-validation to select the penalization parameters that minimize the out-of-sample quantile loss function.

In a second step, we smooth the conditional quantile estimates $X_{i,t}\hat{B}$ using a log-linear function. We fit the function $\hat{Q}_{X_{i,t}}(\tau \mid X_{i,t}) \approx g_{i,t}(\tau) = \log(a_{1,t} + b_{i,t}\tau)$. This gives us an estimate of macro risk, $\hat{\pi}_{i,t} = g_{i,t}^{-1}(-k)$. we fit the function $\hat{Q}_{X_{i,t}}(\tau \mid X_{i,t}) \approx g_{i,t}(\tau) = \log(a_{1,t} + b_{i,t}\tau)$, where g has a log-linear form. We obtain $\hat{\pi}_{i,t} = g_{i,t}^{-1}(-k)$. Finally, we interpolate $\hat{\pi}_{i,t}$ as $\hat{\pi}_{i,t} = g_{i,t}^{-1}(-2)$.

Table A.III: Marginal Effects: Excluding Predictor Categories

	(1)	(2)	(3)	(4)	(5)	(6)
Recession	0.006**		−0.004	0.000	0.010***	−0.002
Recession abroad	0.008***		0.011***	0.010***	0.007**	0.012***
Consumption	0.008***		0.005*	0.006**	−0.001	−0.006**
Consumption (world)	−0.002		−0.002	−0.003	0.008***	0.007**
War	0.003	−0.000		0.008***	−0.002	0.001
Civil war	0.006**	0.014***		−0.002	0.009***	0.010***
Political crisis	0.002	0.009***		0.006**	0.003	0.002
War/political crisis abroad	0.011***	−0.002		0.002	−0.003	0.002
Natural disaster	−0.003	0.007**	0.007**		0.002	−0.001
Famine	0.003	0.002	−0.004		0.010***	−0.004*
Banking crisis	0.006**	0.012***	0.006**	0.010***		0.001
Currency crisis	0.002	0.002	0.002	0.001		0.012***
Sovereign default	0.001	0.003	0.004	−0.002		0.009***
Hyperinflation	−0.002	−0.002	0.003	0.002		0.003
Credit growth (world)	0.000	−0.003	−0.001	0.008***		0.005*
Dividend-price ratio (US)	0.010***	0.003	−0.002	0.006**	0.006*	
Stock volatility (US)	0.006**	0.003	0.011***	0.003	0.002	
Yield curve (US)	−0.003	0.005*	0.008***	0.004	0.007***	

Notes. This table shows the standardized marginal effects for the multivariate quantile regression. The marginal effect is defined as the impact of a one standard deviation increase in a predictive variable on the unconditional probability of a macroeconomic crisis.
*, **, and *** indicate statistical significance at 10%, 5%, and 1% levels, respectively, based on bootstrapped p-values.

Table A.IV: Marginal Effects: Subsamples

	Full sample	Pre-1945	Post-1945	OECD	Ex-OECD
Recession	0.006**	0.004	0.006**	0.007*	0.006
Recession abroad	0.008***	0.006	0.005**	0.009**	0.008**
Consumption	0.008***	0.012***	0.008***	0.006*	0.013***
Consumption (world)	−0.002	0.001	0.006***	−0.000	−0.005
War	0.003	0.005	−0.002	0.001	0.005
Civil war	0.006**	0.003	0.007***	0.004	0.008**
Political crisis	0.002	−0.004	0.004*	0.001	−0.001
War/political crisis abroad	0.011***	0.033***	−0.004	0.012**	0.005
Natural disaster	−0.003	−0.004	−0.002	−0.002	−0.005
Famine	0.003	0.004	0.000	0.002	0.004
Banking crisis	0.006**	0.002	0.007**	0.006*	0.004
Currency crisis	0.002	0.004	0.002	0.004	−0.003
Sovereign default	0.001	−0.002	−0.000	0.004	0.001
Hyperinflation	−0.002	−0.008**	0.001	−0.003	0.003
Credit growth (world)	0.000	0.005	0.005*	0.000	−0.002
Dividend-price ratio (US)	0.010***	−0.002	0.006*	0.009**	0.009*
Stock volatility (US)	0.006**	−0.000	0.001	0.004	0.005
Yield curve (US)	−0.003	0.006	−0.004	−0.001	−0.005

Notes. This table shows the standardized marginal effects for the multivariate quantile regression. The marginal effect is defined as the impact of a one standard deviation increase in a predictive variable on the unconditional probability of a macroeconomic crisis.
*, **, and *** indicate statistical significance at 10%, 5%, and 1% levels, respectively, based on bootstrapped p-values.

Table A.V: Marginal Effects: Horizons

	$H = 3$	$H = 1$	$H = 2$	$H = 4$	$H = 5$
Recession	0.006^{**}	0.008^{**}	0.007^{**}	0.001	0.002
Recession abroad	0.008^{***}	0.006	0.008^{**}	0.008^{***}	0.008^{***}
Consumption	0.008^{***}	0.003	0.004	0.002	0.003
Consumption (world)	−0.002	−0.001	−0.000	−0.004	−0.001
War	0.003	0.003	0.003	0.002	0.003
Civil war	0.006^{**}	0.001	0.003	0.006^{**}	0.006^{**}
Political crisis	0.002	0.002	0.002	−0.001	−0.003
War/political crisis abroad	0.011^{***}	0.008^{**}	0.011^{***}	0.003	−0.001
Natural disaster	−0.003	−0.000	−0.002	−0.002	−0.002
Famine	0.003	0.005^{**}	0.004^{*}	−0.000	−0.002
Banking crisis	0.006^{**}	0.004	0.006^{*}	0.005^{*}	0.004
Currency crisis	0.002	0.004	0.001	0.003	0.002
Sovereign default	0.001	0.000	0.000	−0.000	−0.002
Hyperinflation	−0.002	0.000	−0.001	−0.002	0.001
Credit growth (world)	0.000	−0.001	−0.001	−0.005	-0.011^{***}
Dividend-price ratio (US)	0.010^{***}	0.012^{***}	0.011^{***}	0.012^{***}	0.011^{***}
Stock volatility (US)	0.006^{**}	0.006^{*}	0.007^{**}	0.005^{*}	0.004
Yield curve (US)	−0.003	−0.002	−0.004	−0.002	−0.002

Notes. This table shows the standardized marginal effects for the multivariate quantile regression. The marginal effect is defined as the impact of a one standard deviation increase in a predictive variable on the unconditional probability of a macroeconomic crisis.
*, **, and *** indicate statistical significance at 10%, 5%, and 1% levels, respectively, based on bootstrapped p-values.

Appendix C. Crisis events

Table A.VI lists the 156 crises in our sample.

Table A.VII tests whether the autocorrelation of realized crises has changed after 1945. To do so, we construct a crisis indicator $y_{i,t}$, which is equal to one if the one-year normalized consumption growth is below a given cutoff and zero otherwise. To get an idea of persistence across the left side of the consumption distribution, we consider several cutoffs from 1 to 2.5. We regress $y_{i,t}$ on its value the previous year as well as on its past value interacted with a dummy variable that is equal to one if the observation is after 1945. A negative coefficient indicates a decrease in persistence after 1945. The table indicates that none of the interaction coefficients is statistically significant. Regardless of the cutoff value, crises have not become significantly less persistent after 1945. This indicates that measurement errors are unlikely to exaggerate crisis predictability.

Table A.VI: List of Crises

Year	Disasters	Year	Disasters
1912	Canada (−0.12), Germany (−0.17)	1974	Chile (−0.46)
1913	Canada (−0.14), France (−0.24), Germany (−0.28), Netherlands (−0.15)	1975	Portugal (−0.08)
1914	France (−0.19), Germany (−0.38), Netherlands (−0.18), Russia (−0.19), Sweden (−0.06), Turkey (−0.37), United Kingdom (−0.08)	1976	Peru (−0.14)
1915	Germany (−0.31), Russia (−0.36), Sweden (−0.07), Turkey (−0.56), United Kingdom (−0.16)	1977	Peru (−0.19)
1916	Finland (−0.36), Germany (−0.27), Netherlands (−0.43), Norway (−0.08), Russia (−0.73), United Kingdom (−0.18)	1981	Mexico (−0.08)
1917	Finland (−0.18), Russia (−0.68)	1982	Chile (−0.29), Mexico (−0.12)
1918	Russia (−0.72), United States (−0.11)	1983	Philippines (−0.08), Venezuela (−0.29)
1919	Canada (−0.22), Denmark (−0.06), United States (−0.15)	1985	Brazil (−0.11), Malaysia (−0.16)
1920	Denmark (−0.22), Norway (−0.07)	1987	Argentina (−0.13), Brazil (−0.12), Norway (−0.05), Peru (−0.25)
1921	New Zealand (−0.19)	1988	Argentina (−0.17), Peru (−0.35)
1929	Australia (−0.22), Brazil (−0.22), Italy (−0.06)	1989	Peru (−0.22)
1930	Canada (−0.22), Chile (−0.47), Finland (−0.20), Peru (−0.15), United States (−0.21)	1990	Finland (−0.11)
1931	Canada (−0.20), Chile (−0.40), United States (−0.17)	1991	Finland (−0.14)
1934	Italy (−0.06), Spain (−0.49)	1992	Finland (−0.07)
1935	Spain (−0.61)	1993	Mexico (−0.10)
1936	Spain (−0.56)	1997	Colombia (−0.10)
1938	Denmark (−0.23), Japan (−0.20), United Kingdom (−0.11)	1998	Colombia (−0.10)
1939	Denmark (−0.29), Egypt (−0.16), Finland (−0.26), France (−0.49), Japan (−0.22), United Kingdom (−0.15)	1999	Turkey (−0.12)
1940	Belgium (−0.72), Denmark (−0.22), Egypt (−0.18), France (−0.65), Italy (−0.10), Japan (−0.23), Netherlands (−0.63), Russia (−0.77), Sweden (−0.17), Switzerland (−0.13), United Kingdom (−0.16)	2000	Argentina (−0.26)
1941	France (−0.75), Italy (−0.24), Japan (−0.18), Korea (−0.09), Netherlands (−0.62), Russia (−0.77), Switzerland (−0.12), Taiwan (−0.23)	2001	Taiwan (0.02)
1942	Austria (−0.42), Italy (−0.18), Japan (−0.34), Korea (−0.23), Netherlands (−0.53), Taiwan (−0.51)	2007	Iceland (−0.28), Italy (−0.05), Spain (−0.09), United Kingdom (−0.02), United States (−0.02)
1943	Germany (−0.30), Italy (−0.24), Japan (−0.71), Korea (−0.47), Taiwan (−0.86)	2008	Denmark (−0.06), Iceland (−0.29), Spain (−0.09), United Kingdom (−0.04), United States (−0.02)
1944	Korea (−0.30), Taiwan (−0.54)	2009	Greece (−0.14), Spain (−0.09), United Kingdom (−0.04)
1948	India (−0.19)	2010	Greece (−0.23)
1950	Korea (−0.46), Singapore (−0.20)	2011	Greece (−0.19), Italy (−0.08), Portugal (−0.09)
1952	Egypt (−0.18)	2012	Italy (−0.08)
1972	Chile (−0.26)	2013	Austria (−0.02)
1973	Chile (−0.47)	2018	Argentina (−0.26), Australia (−0.03), Austria (−0.09), Belgium (−0.07), Canada (−0.06), France (−0.05), Germany (−0.04), Italy (−0.08), Mexico (−0.11), Norway (−0.06), Singapore (−0.09), Spain (−0.12), Switzerland (−0.04), United Kingdom (−0.11)

Notes. This table lists crises based on consumption data in our panel of 42 countries from 1900-2020. Disaster sizes are indicated in parentheses.

Table A.VII: Persistence in Realized Crises Before and After 1945

Cutoff:	1	1.5	2	2.5
$y_{i,t-1}$	0.243***	0.238***	0.209***	0.239**
	(0.04)	(0.05)	(0.07)	(0.09)
$y_{i,t-1} \times 1_{t>1945}$	0.067	−0.091	−0.096	−0.143
	(0.07)	(0.08)	(0.09)	(0.11)
R^2	0.081	0.038	0.029	0.036
N	4,762	4,762	4,762	4,762

Notes. This table reports regression results for crisis indicators, constructed based on one-year normalized consumption growth falling below cutoffs ranging from 1 to 2.5. The indicators are regressed on its value from the previous year and its past value interacted with a post-1945 dummy variable. A negative coefficient indicates decreased persistence after 1945.

Appendix D. Macro Risk: Alternative Estimates and Figures

In the baseline model, we estimate $\hat{\pi}$ using quantile regression with a Ridge penalization. A crisis is defined as a 2 standard-deviation drop in three-year consumption growth, and 0 otherwise. We present alternative estimates below.

Excluding Asset Price Predictors

In the spirit of Mehra and Prescott (1985), and subsequent literature, we estimate our model without relying on asset pricing data. Figure A.I shows macro risk estimates excluding asset price predictors, $\hat{\pi}^{-}$, which we use in the estimation of the disaster risk model parameters.

Horizon And Threshold Sensitivities

Figure A.II shows sensitivity results for U.S. macro risk with respect to the forecasting horizon.
Figure A.III shows sensitivity results for U.S. macro risk with respect to the crisis threshold.

Unpenalized And Probit Models

Figure A.IV shows US estimates for unpenalized regressions as well as LASSO and Elastic Net penalizations.
Figure A.V shows US estimates constructed using a probit regression model.

Heterogeneous Loadings Model

Our baseline model imposes the homogeneity restriction that all countries have the same loading on individual predictors. Figure A.VI shows an estimate for an alternative model where the countries' exposures to global variables are allowed to be heterogeneous. To do so, we write down an alternative model in which the countries' exposures to global variables are allowed to differ. The conditional quantiles $Q_{X_{i,t}}(\tau|X_{i,t})$ are thus given by:

$$Q_{X_{i,t}}(\tau|X_{i,t}) = X_{i,t}^{L} B^{L} + X_{t}^{G} B_{i}^{G}, \tag{D1}$$

where we partitioned the predictor vector $X_{i,t}$ into local $X_{i,t}^{L}$ and global components X_{t}^{G} (i.e., world consumption growth, global credit growth, and the three asset price predictors). The model above can be cast into our common slope model by expanding the predictor set to include interactions with country fixed effects. We find that this alternative model generates very similar estimates in-sample, but has limited out-of-sample forecasting power, likely due to overfitting of the individual loading parameters.

Subsample Sensitivities and GDP-Based Estimates

Figure A.VII shows US estimates constructed using the pre- and post-WW2 samples.
Figure A.VIII shows US estimates training the model using OECD and non-OECD countries.
Figure A.IX compares U.S. macro risk estimates and the cross-country average.
Figure A.X shows estimates based on GDP data.

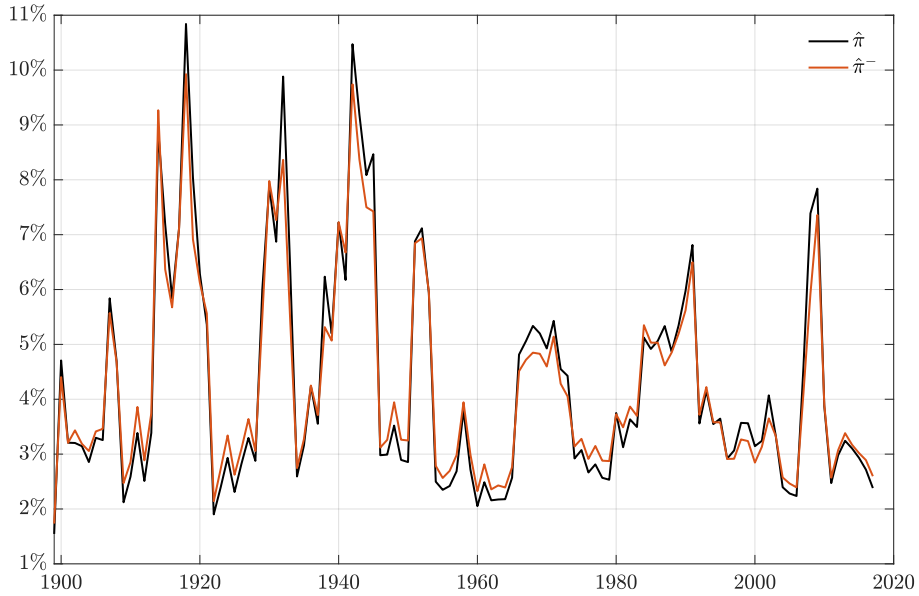

Figure A.I: Excluding Asset Price Predictors

Notes. This figure compares our baseline U.S. macro risk based with estimates excluding asset price predictors, $\hat{\pi}^-$.

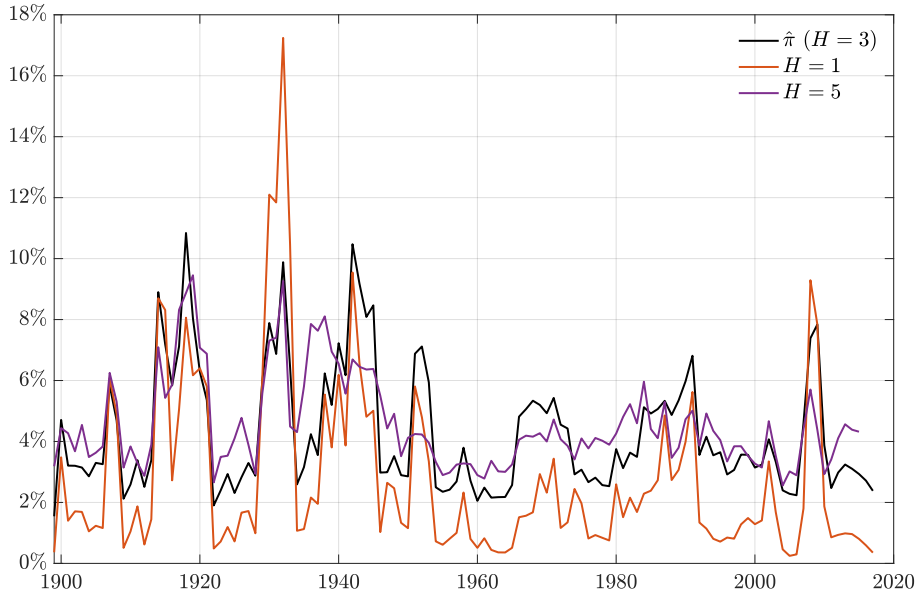

Figure A.II: Horizon sensitivity

Notes. This figure compares U.S. macro risk against the probabilities of macroeconomic crises one and five years ahead.

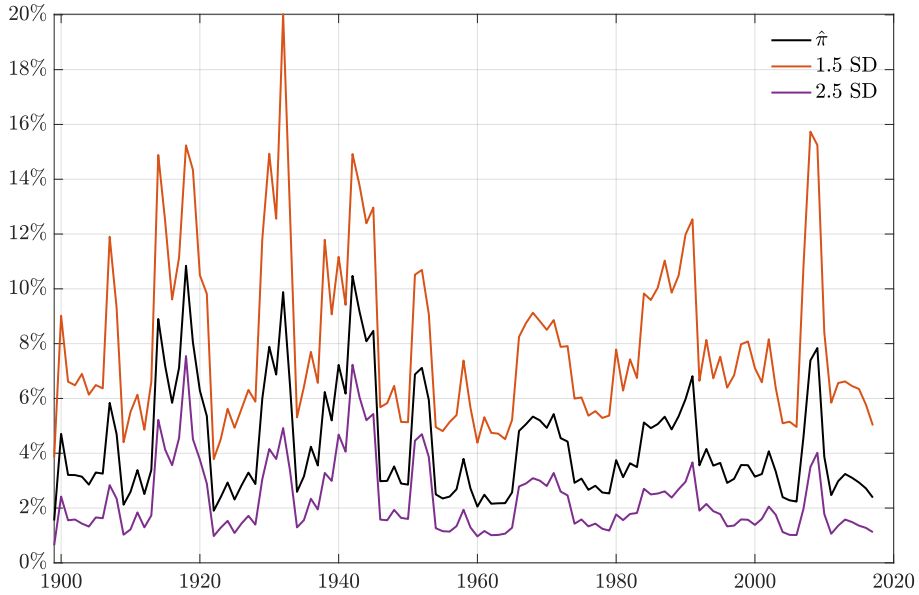

Figure A.III: Threshold Sensitivity

Notes. This figure compares our baseline U.S. macro risk estimate to alternative estimates constructed using crisis thresholds of 1.5, and 2.5 standard deviations from average consumption growth rates.

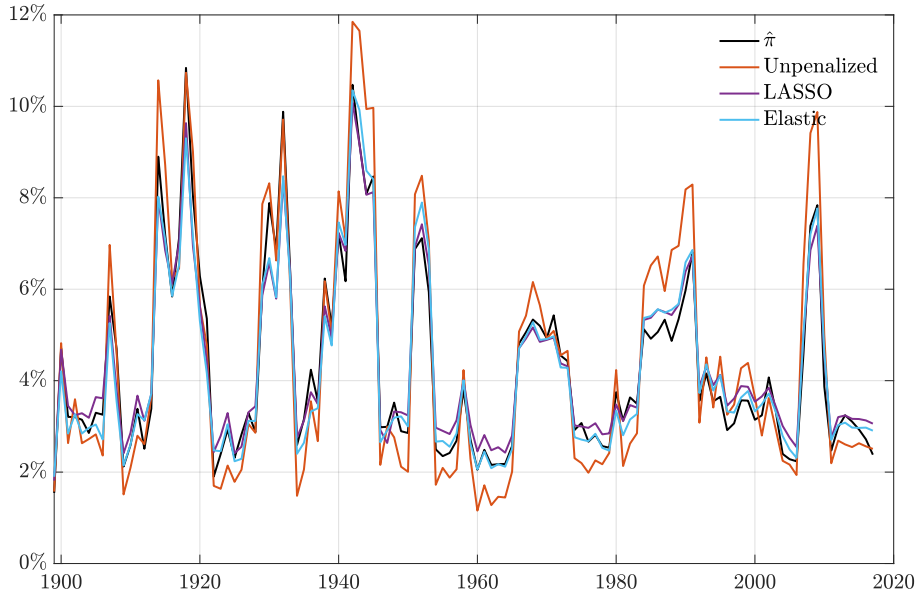

Figure A.IV: Penalized Regression Estimates

Notes. This figure compares U.S. macro risk against penalized regression estimates.

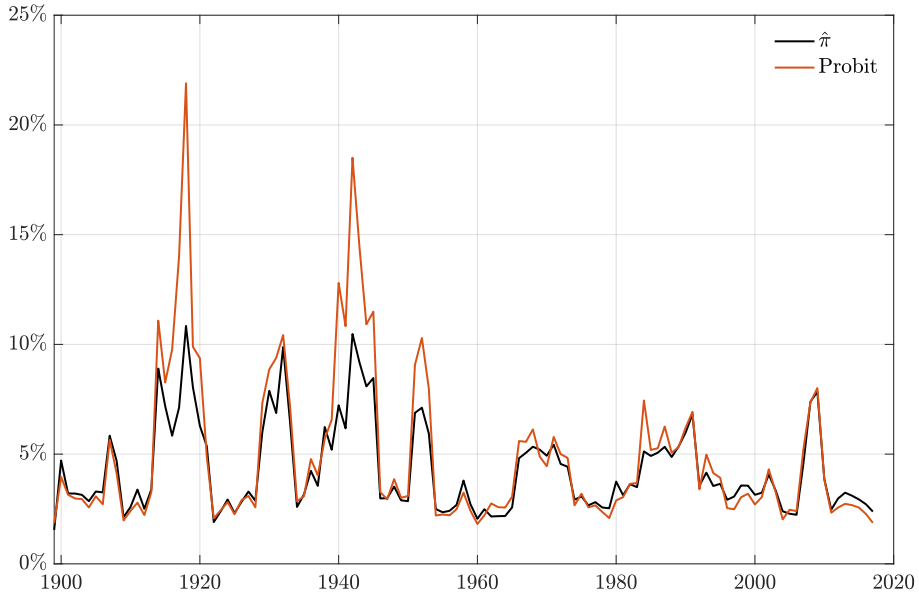

Figure A.V: Probit Estimates

Notes. This figure compares U.S. macro risk against probit estimates.

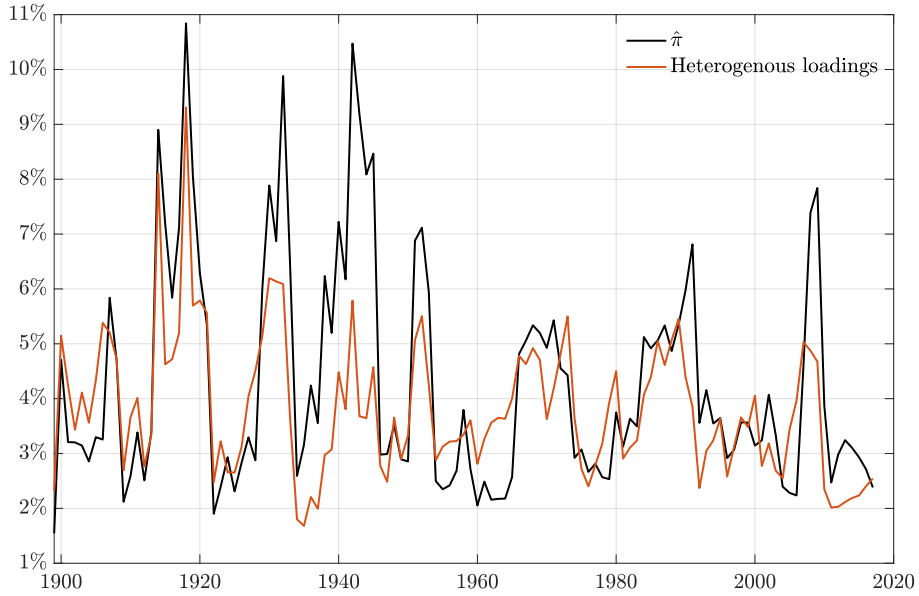

Figure A.VI: Relaxing the Homogeneity Restriction

Notes. This figure compares our baseline U.S. macro risk estimate to an alternative model where the countries' exposures to global variables are allowed to be heterogeneous. The figure shows estimates based on a Ridge penalization.

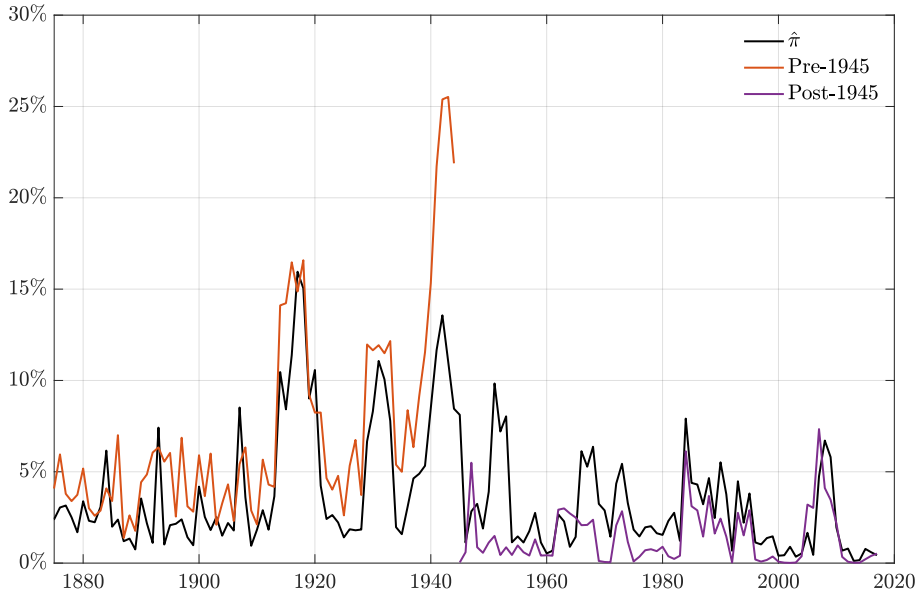

Figure A.VII: Subsamples: before/after World War 2

Notes. This figure compares our baseline U.S. macro risk estimate to alternative estimates constructed from separate regressions for the pre- and post-WW2 samples. Marginal effect estimates are reported in Table A.IV.

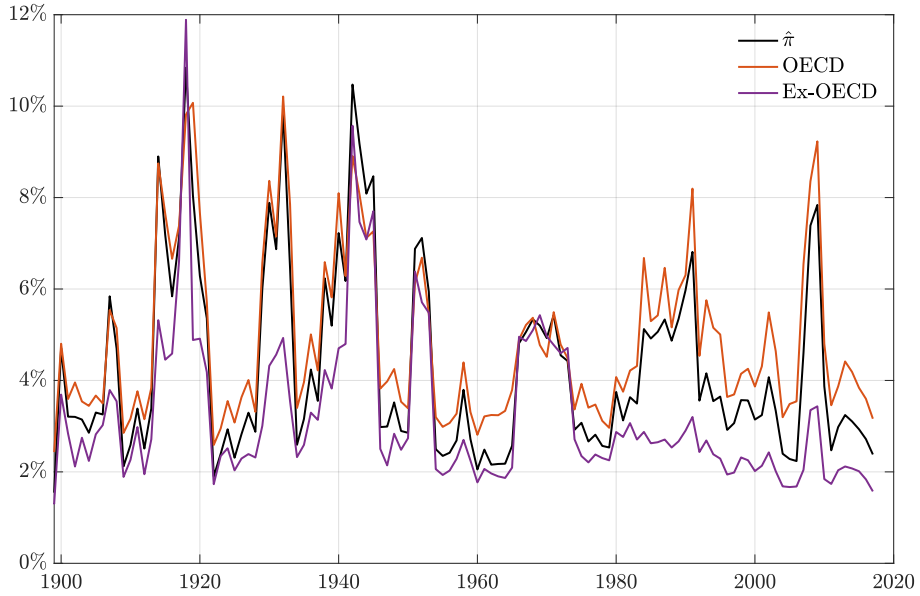

Figure A.VIII: Subsamples: OECD versus non-OECD countries

Notes. This figure compares our baseline U.S. macro risk estimate to alternative estimates constructed from separate regressions for the OECD sample only and another excluding the OECD sample. The corresponding estimates of marginal effect are reported in Table A.IV.

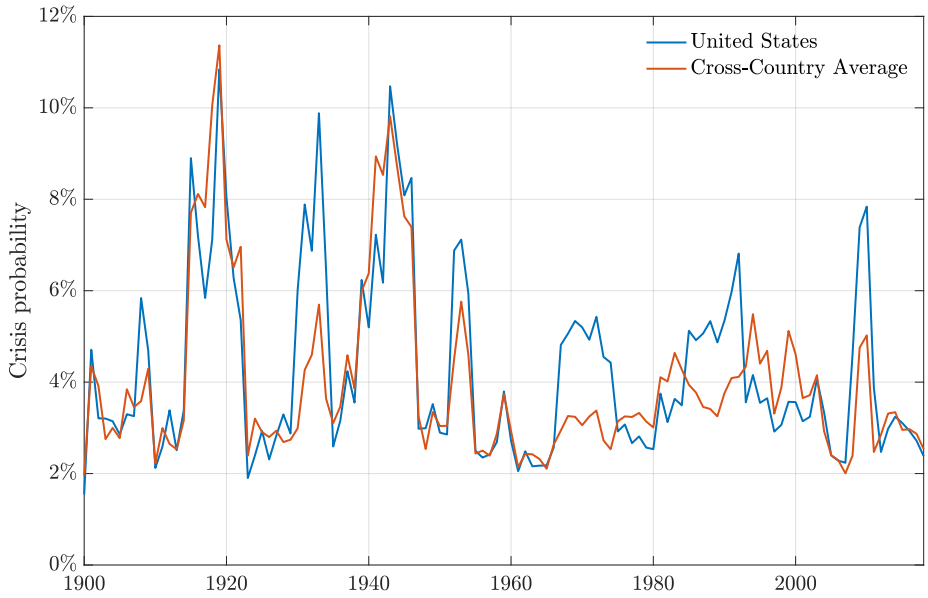

Figure A.IX: US Estimates vs. Cross-Country Average

Notes. This figure compares U.S. macro risk against the arithmetic average of individual country estimates.

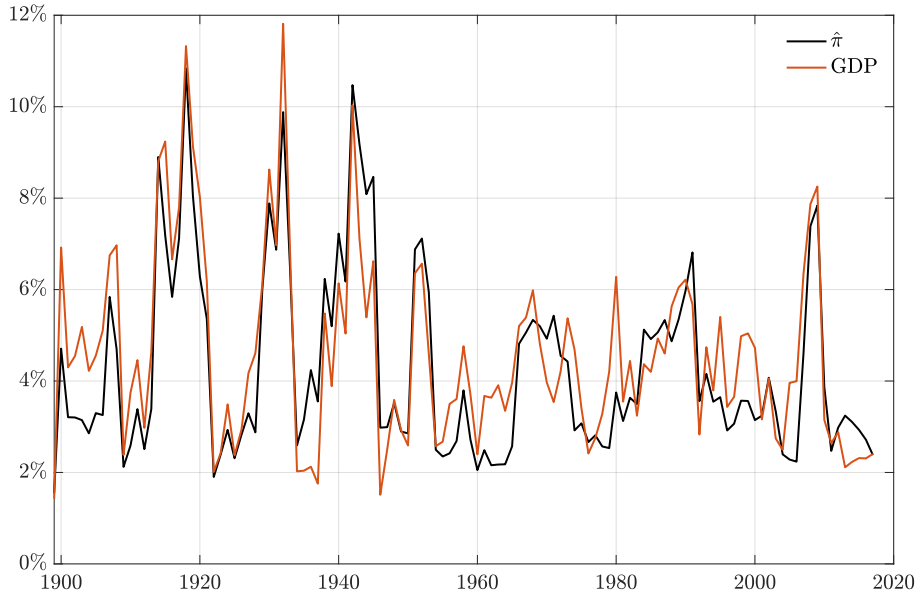

Figure A.X: Crisis definition: consumption vs. GDP crises

Notes. This figure compares our baseline U.S. macro risk based on consumption data to an alternative estimate based on GDP data.

Local And Global Risks Decomposition

A potential concern is that our econometric setup does not distinguish between local and global crises, which could be predicted by different sets of variables. To address this concern, we use a probit approach as it allows us to work with realized crises that we can classify as global or local. We assume a country is in a *global* crisis if normalized consumption falls by more than two standard deviations and more than 20% of countries experience 2-SD crises at the same time. Other 2-SD crises in which less than 20% of countries are affected are considered local crises. During a global crisis, if a country's consumption does not exceed two standard deviations, that country is not considered to be in crisis. This allows us to maintain the same crisis frequency, and thus this approach remains consistent with our baseline method. Roughly, this rule corresponds to treating the two world wars and the Covid crisis as global crises and the remaining crises as local crises. Empirically, global crises and local crises are uncorrelated (corr. $= -0.02$).

Equipped with global and local crises, we construct the two crisis probabilities of interest. We forecast global crises with global variables and local crises with local variables. (We find that not imposing this assumption yields similar results). Figure A.XI displays a stacked area plot that decomposes global, $\hat{\pi}^G$, and local $\hat{\pi}^L$ crisis probabilities for the United States. We note that World War II, the Great Depression following 1929, the 2008 crisis, and to a lesser extent World War I, are the periods with the highest global risks ($\hat{\pi}^G$). In contrast, US-specific macro risk ($\hat{\pi}^L$) dominates in the second half of the twentieth century. We also find that the correlation between the two probabilities is positive but low (corr. $= 0.19$).

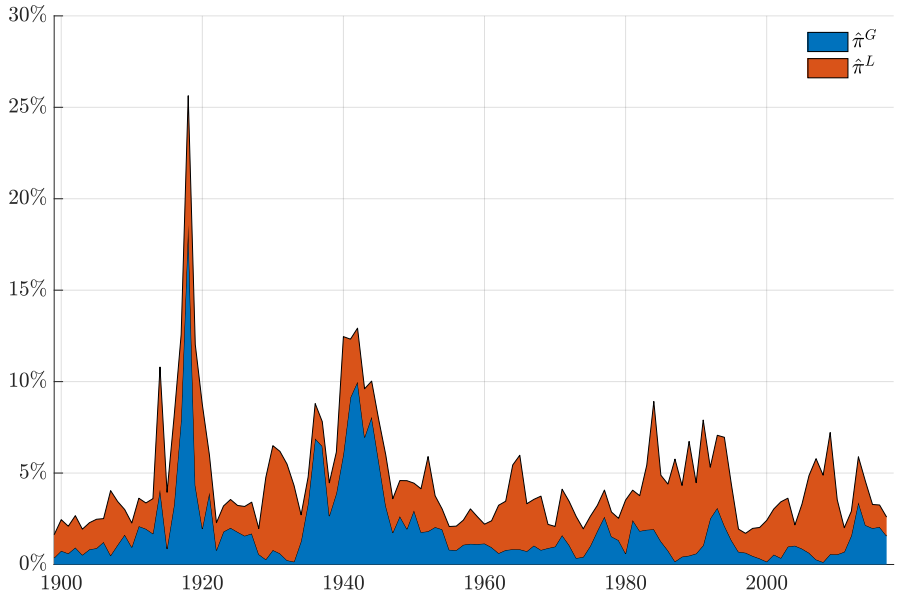

Figure A.XI: Local and Global Risks Decomposition

Notes. This figure presents a stacked area plot decomposing US macro risk into global ($\hat{\pi}^G$) and local ($\hat{\pi}^L$) components. The decomposition utilizes a probit model, distinguishing crises as either local or global. Global crises are defined by instances where more than 20% of countries are in a crisis, characterized by a consumption drop of more than two standard deviations. Crises affecting less than 20% of countries are classified as local.

We next ask if our baseline approach (where we pool local and global predictors to estimate by quantile regressions consumption growth conditional distributions) could be biased. By construction, there is no overlap between local and global crises, which means that the sum of the two crisis probabilities ($\hat{\pi}^w + \hat{\pi}^j$) should be comparable with $\hat{\pi}$. Figure A.XII compares the

sum of the two probabilities and $\hat{\pi}$. We see that the two probabilities are to a large extent similar and highly correlated; this is true for the US but also globally (corr. = 0.80). This result is not sensitive to the choice of decomposition. We have tried with alternative thresholds defining a global crisis at 10% and 40%, for example, and the sum of global and local probabilities remain similar to our baseline $\hat{\pi}$. We therefore conclude that the potential bias is not materially important in our baseline setup.

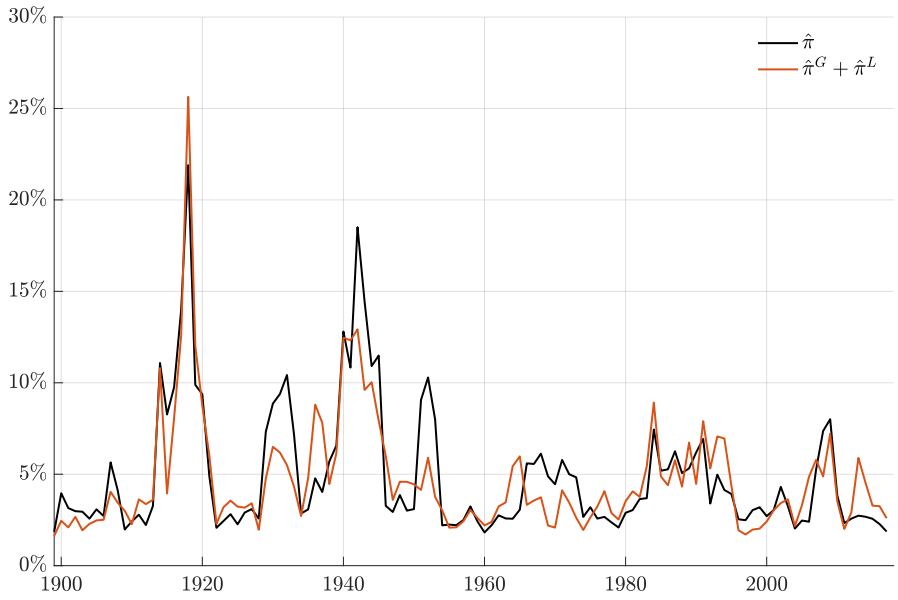

Figure A.XII: Robustness to Local and Global Risks Distinction

Notes. This figure compares our the probit U.S. macro risk estimate to the sum of global and local risk estimates constructed by separating global and local crises.

Two-Stage Quantile Regression Approach

We next consider an alternative approach where we replace binary crisis indicators (wars, natural disasters, political crises, etc.) with crisis *probabilities*, constructed in a preliminary stage. In the first stage, we construct the time t probability of such non-macro crises between $t+1$ and $t+3$, which we use to forecast consumption quantiles in the second stage. This approach gives us an alternative variable for each type of non-macro crisis that we can include in $X_{i,t}$. Our predictive set contains $j = 1 \dots 10$ binary indicators encoding non-macro crises. Let us denote by $q_t^{(j)}$ the probability of a type j crisis within three years ($t+1$ to $t+3$). For each j, we estimate $q_t^{(j)}$ using a probit regression with our baseline set of predictors, $X_{i,t}$ and where the dependent variable equals one if there is at least one type j crisis between $t+1$ and $t+3$. To avoid overfitting, we estimate the model using a LASSO penalization tuned with 10-fold cross-validation.

We present the estimation results of non-macro crises in Figure A.XIII for the United States. The blue lines indicate the penalization estimates and the dashed red lines indicate the non-penalized estimates. We also highlight in gray the realization of crises. The average pairwise correlation between the indicators and the probabilities is 0.60. This correlation is unsurprising. Events such as financial crises and wars tend to persist over time. Consequently, the best predictor of a crisis is the realization of a crisis state in the past (see e.g., Mueller and Rauh (2022) for a similar observation in the context of predicting wars).

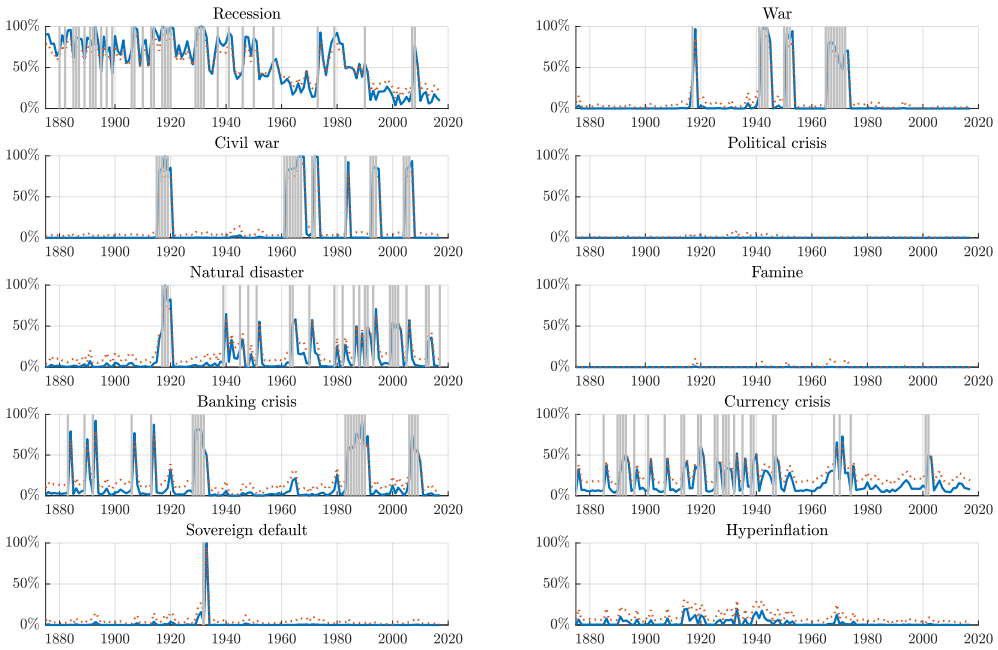

Figure A.XIII: Probabilities of Non-Macro Crises in the United States

Notes. This figure presents the estimation results of non-macro crises in the United States. It shows crisis probability estimates for each of the ten non-macro crises, such as wars, natural disasters, and political crises. Each probability is constructed with a probit model, where the dependent variable is an indicator equal to one for crises occurring between $t+1$ and $t+3$. The blue lines represent 10-fold cross-validation LASSO estimates, while the dashed red lines depict non-penalized estimates. Areas shaded in gray signify the actual realization of crises.

Equipped with these new predictors, we can then re-estimate the main model for $\hat{\pi}$. This alternative specification replaces the binary indicators represented in gray in Figure A.XIII with

the crisis probabilities $q_{i,t}^{(j)}$ represented in blue. We present the estimates $\tilde{\pi}$ for the US in Figure A.XIV, compared as usual to our baseline $\hat{\pi}$. These estimates are very similar: this is true for the US as for the rest of the countries in our sample (corr. = 0.92).

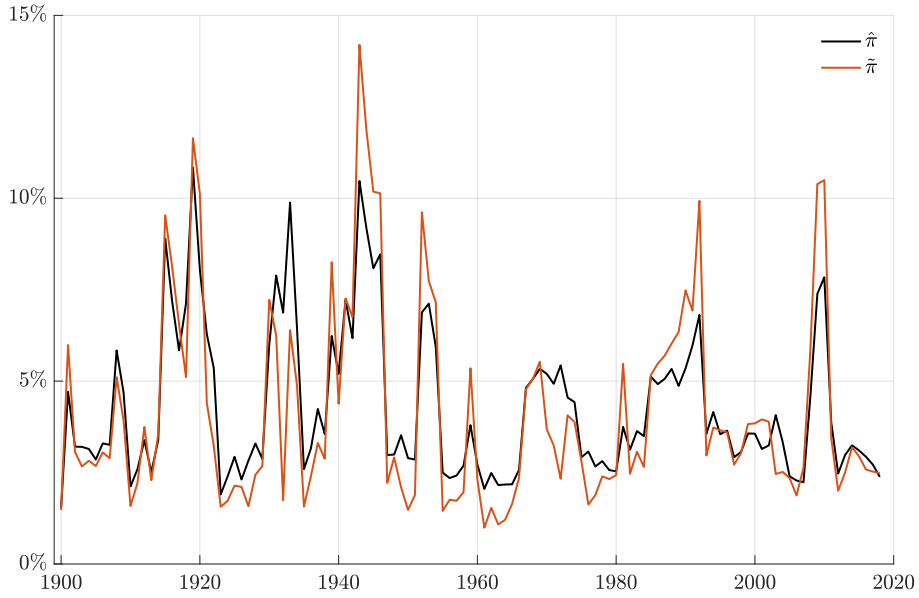

Figure A.XIV: Forecasting using non-macro crisis probability estimates

Notes. This figure compares our baseline U.S. macro risk estimate to an alternative estimate constructed from replacing non-macro crisis indicators with non-macro probabilities constructed from probit estimates.

Macro Risk and Market Volatility

Figure A.XV compares the U.S. option-implied crisis probability in Barro and Liao (2021) together with the volatility index. The volatility index is an element in our predictor set (see Section I.C), which we show here scaled to be comparable with Barro and Liao's option-implied probability.

Figure A.XVI compares $\hat{\pi}$ to an alternative estimate constructed solely from the volatility index.

Figure A.XVII compares $\hat{\pi}$ to Manela and Moreira (2017)'s NVIX.

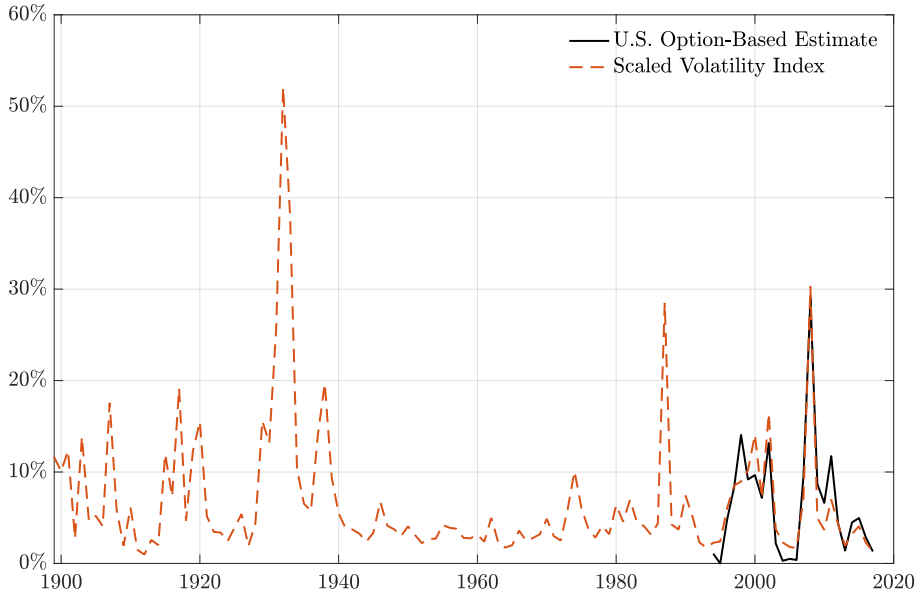

Figure A.XV: Option-Based Estimates and the Volatility Index

Notes. This figure shows the U.S. option-implied crisis probability in Barro and Liao (2021) together with the volatility index.

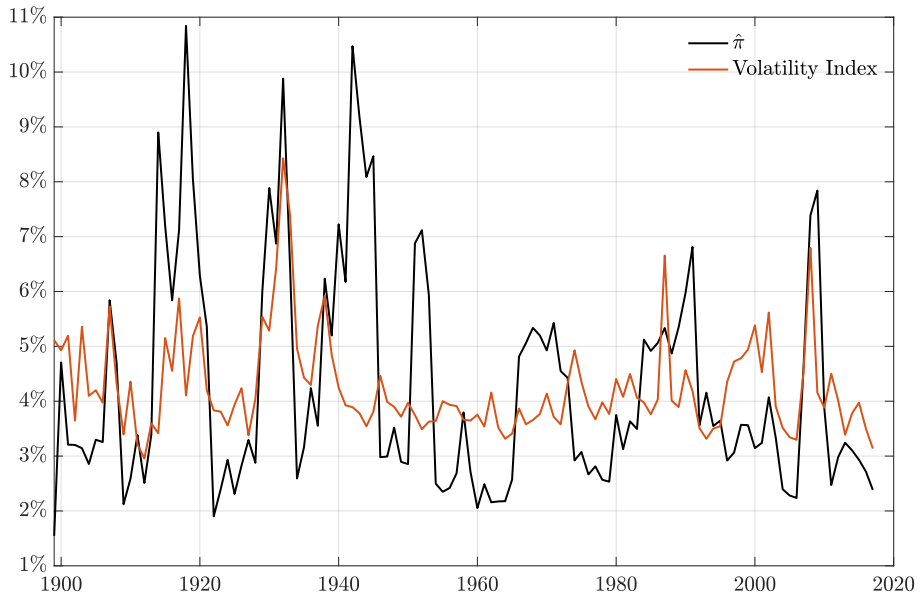

Figure A.XVI: Baseline vs. Volatility Index

Notes. This figure compares our baseline U.S. macro risk estimate to an alternative estimate constructed using the volatility index.

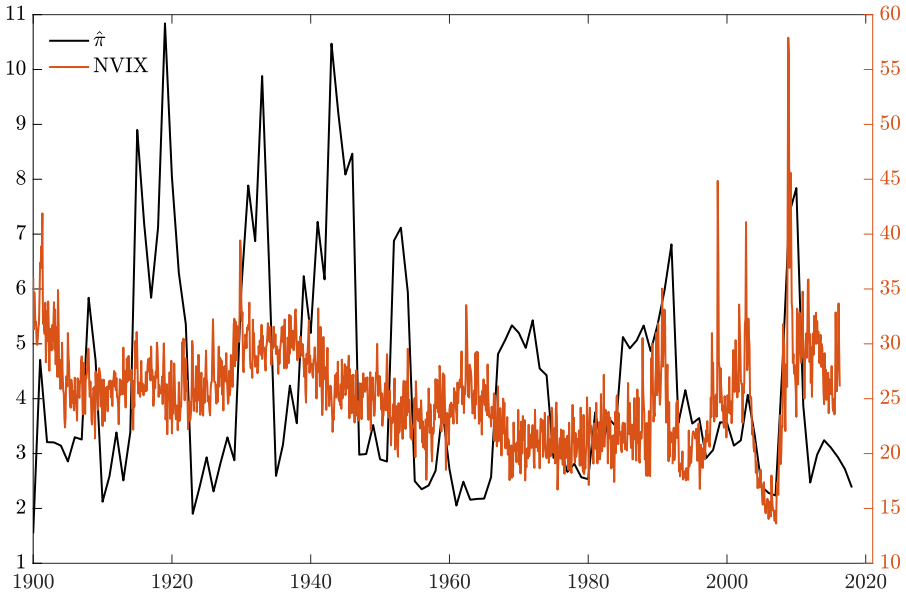

Figure A.XVII: Macro Risk and News Implied Volatility

Notes. This figure compares our baseline U.S. macro risk estimate to the News implied volatility index (NVIX) proposed by Manela and Moreira (2017). The NVIX series spans the 1889–2007 to 2016–2003 period and was downloaded from Asaf Manela's website in February 2022.

Appendix E. Comparison with Berkman et al. (2011)

Berkman et al. (2011) uses international political crises such as wars and other violent acts to proxy for macroeconomic tail risk. Specifically, they use the number of international crises in a given month as their main proxy. Figure A.XVIII shows the number of international crises used in Berkman et al. (2011) together with $\hat{\pi}$ for the United States. The correlation between the two series is positive (0.15), but there are also marked differences between the two measures over several episodes. For instance, there are relatively few international political crises in the interwar periods, whereas $\hat{\pi}$ rises to about 10% (a fivefold increase) in the wake of the 1929 crash. In contrast, $\hat{\pi}$ declines at the beginning of the Great Moderation in the 1980s, while political risk spikes after the Soviet intervention in Afghanistan in 1979, a major crisis at the time.

To better understand the two series, we compare the raw datasets. Berkman et al. (2011) relies on a database of major international political crises from the International Crisis Behavior project (ICB), which covers the 1918–2006 period. Our wars and political crises data comes from Sarkees and Wayman (2010) and from the Center for Systemic Peace (CSP), which we extended to 2019 using online sources (see Appendix A). Both datasets have strengths and weaknesses. The ICB database is sampled monthly and crises are defined on a narrative-based approach. As noted in Berkman et al. (2011), this data is best suited for event studies since the timing of the crises are likely to be closely aligned with the news to which investors react. Our yearly dataset has the advantage of starting earlier, which allows us to cover the First World War and the late nineteenth-century crises in our analysis.

Figure A.XIX shows the number of wars and political crises in the paper and in the ICB dataset. Following Berkman et al. (2011), we calculate the number of political crises as the sum of crises that start or are ongoing in a given month and plot the average number of international crises per month. Once again, the two series are positively but fairly weakly correlated (the correlation is 0.15). To a large extent, our wars and political crises data reflect the number of countries involved in armed conflicts. Therefore, we see peaks in the number of crises around the two world wars, as well as the Korean and Vietnam Wars. The ICB dataset primarily focuses on perceived crises, and thus includes many crises that did not translate into armed conflicts. The largest number of crises is actually in 1919, a year marked by substantial international tension in the wake of the First World War. The number of recorded crises during the Second World War is modest in comparison. Likewise, in the postwar period, the ICB data reflect crises such as the Berlin Blockade in July 1948, the Berlin crisis in 1961, the Cuban Missile Crisis in 1962, and many other crises, that did not always translate into major conflicts.

Berkman et al. (2011) shows that political crises have forecasting power for future crises. We next compare the forecasting power of ICB crises with wars and political crises in our dataset. To do so, we map the number of crises to a disaster probability using our quantile regression approach. The AUROC is positive but relatively low (0.53), similar to the AUROC for individual predictors such as options. Interestingly, the number of wars and political crises in Figure A.XIX delivers a higher AUROC of 0.66. This exercise confirms that the crises in Berkman et al. (2011) are best suited for event studies than for constructing an objective measure of disaster risk.

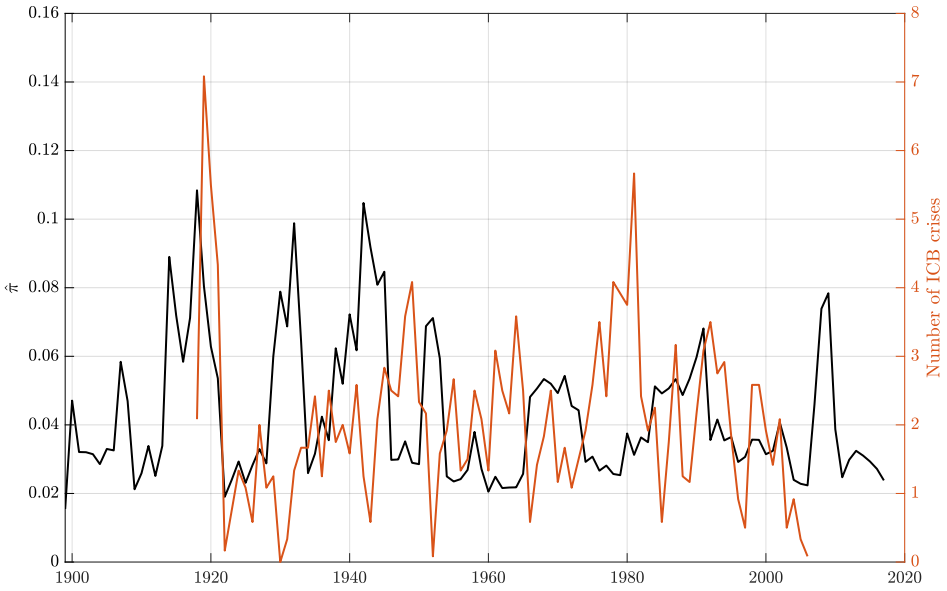

Figure A.XVIII: Comparison with Berkman et al. (2011)

Notes. This figure compares π for the United States with the number of wars and political crises in the paper and in the ICB dataset used in Berkman et al. (2011).

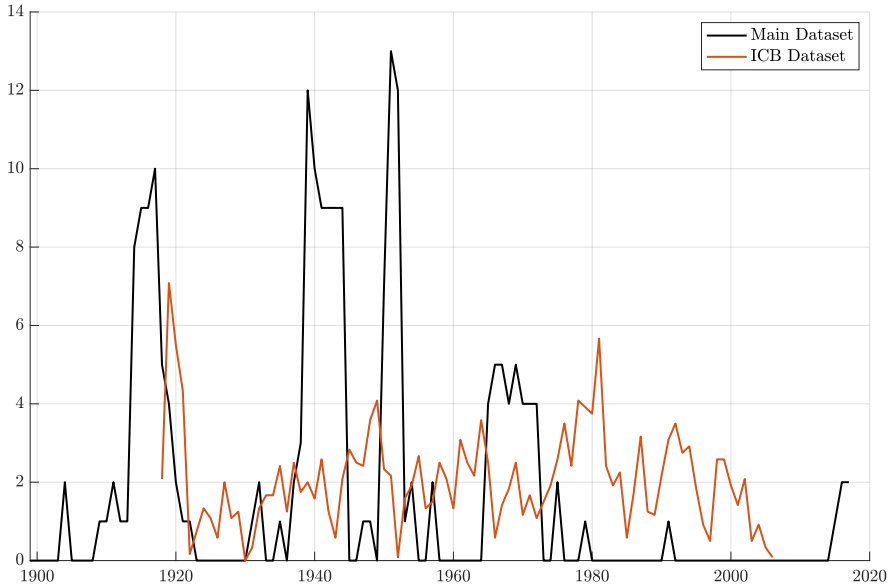

Figure A.XIX: Wars and Political Crises Comparison

Notes. This figure shows the number of wars and political crises in the paper and in the ICB dataset used in Berkman et al. (2011).

Appendix F. Understanding the Marginal Contribution of Credit Growth

There is ample evidence in the literature that credit growth tends to be associated with worse macroeconomic outcomes. However, in our setting, we find credit growth to have little incremental forecasting power when conditioning on the full set of predictor variables. As a result, our estimates are essentially insensitive to credit variables; Figure A.XX below illustrates this insensitivity for the US. We use global credit growth to maximize sample size, but using local credit growth (like is common in the literature) or excluding credit growth deliver very similar estimates, with pairwise correlations close to unity across specifications. This is a general pattern in the paper: we use many predictors, such that individual predictors tend to be redundant.

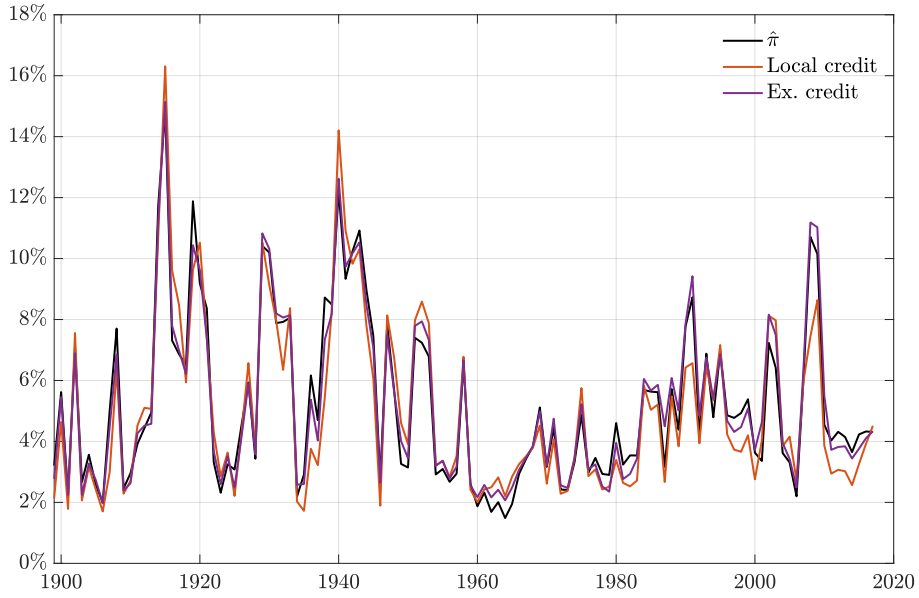

Figure A.XX: Macro Risk Sensitivity to Credit Growth Variables

Notes. This figure compares $\hat{\pi}$ with two alternative specifications constructed either using local credit growth for the subset of countries where the data are available or excluding credit growth.

We explain further why macro risk is insensitive to credit growth below. In particular, we document that the predictive power of local credit growth is relatively weak in the long Jordà-Schularick-Taylor (JST) sample (Jordà et al., 2017). We also show below that our macro risk estimates $\hat{\pi}$ and credit growth *jointly* predict equity returns.

We first revisit the negative association between credit and consumption and GDP documented in the literature (Jordà et al., 2017; Mian et al., 2017). In particular, Mian et al. (2017) (MSV) presents regression results for output growth that are close to the quantile regressions we use in the paper. MSV shows that lagged 3-year credit growth forecasts 3-year GDP growth in an unbalanced panel of 30 countries from 1960 to 2012. We replicate this result using the same time period as MSV in Table A.VIII.

Table A.VIII examines the relationship between lagged 3-year credit growth and future GDP growth, consumption growth, as well as consumption crises. We begin with the MSV specification using their replication dataset and continue with the JST dataset following your suggestion. The first column replicates MSV's baseline result using their replication data. We confirm that the growth of the debt-to-GDP ratio negatively predicts future GDP growth (the

coefficient is identical to the one reported in column 1 from Table 3 in MSV). MSV uses GDP growth, while we use per capita GDP growth, following Barro and Ursúa (2008). In the second column, we replicate this result using our per capita GDP series and the debt-to-GDP ratio from the JST dataset we use in the paper. These two specifications include country fixed effects as per MSV, whereas we do not use fixed effects (since consumption growth is standardized). The third column shows that MSV's result is not sensitive to the inclusion of fixed effects. The fourth column replaces the GDP growth rate with the growth rate of standardized consumption, which is the main variable we use in the paper. The fifth column uses the *global* credit growth predictor from our paper. Again, we document a negative association. Finally, the last two columns use the same set of two predictive variables but replace the dependent variable with a binary indicator that equals one when a crisis occurs, that is, when standardized consumption is below −2. We observe a positive and significant coefficient in both columns, meaning that both local and global credit growth are associated with increased macroeconomic tail risk.

Table A.VIII: Credit Expansion and Future Growth (MSV Sample, 1960–2012)

	$\Delta_3 y_{it+3}^{\text{MSV}}$	$\Delta_3 y_{it+3}$	$\Delta_3 y_{it+3}$	$\Delta_3 c_{it+3}$	$\Delta_3 c_{it+3}$	$1_{\Delta_3 c_{it+3}<-2}$	$1_{\Delta_3 c_{it+3}<-2}$
$\Delta_3 d_{it-1}^{\text{MSV}}$	-0.119^{***}	–	–	–	–	–	–
	(0.03)	–	–	–	–	–	–
$\Delta_3 d_{it-1}^{\text{JST}}$	–	-0.233^{***}	-0.214^{***}	-3.035^{***}	–	0.456^{**}	–
	–	(0.04)	(0.04)	(0.59)	–	(0.20)	–
$\Delta_3 d_{t-1}$	–	–	–	–	-5.851^{**}	–	0.804^{**}
	–	–	–	–	(2.49)	–	(0.34)
FE	✓	✓					
R^2	0.087	0.178	0.177	0.163	0.043	0.065	0.019
N	695	570	570	570	717	570	717

Notes. This table presents the results of a replication and robustness analysis based on the original findings by Mian et al. (2017) (MSV). The sample period is 1960 to 2012. The table showcases various specifications for predicting three-year future output growth (denoted $\Delta_3 y$), standardized consumption growth (denoted $\Delta_3 c$), and crisis indicators (denoted $1_{\Delta_3 c_{it+3}<-2}$). The independent variables are debt-to-GDP ratios constructed at the country or global level. Country variables are denoted with a subscript i. The MSV and JST superscripts in the first column and the first two rows indicate that we use data from MSV or JST, respectively. The first two columns include country fixed effects as per MSV.

MSV uses a shorter sample (1960 to 2012) than ours. We show the results of the same regression using our sample (1875 to 2020) in Table A.IX. We observe that the predictive power of all variables is weaker for both GDP and consumption growth (columns 1-4). This is consistent with the evidence in Jordà et al. (2017) that the negative effects of credit growth are mostly present in the post–World War II period. We note that the OLS coefficient for global credit growth in column 4 is positive but insignificant. Similarly, the univariate 50% quantile regression coefficient we report in Figure 4 is positive but also insignificant. Regarding the predictability of realized crises (columns 5-6), we observe that the coefficient is either insignificant or has the opposite sign, meaning that credit growth is associated with *lower* tail risk.[4] This is again in line with the quantile regression coefficients reported in Figure 4. The corresponding multivariate marginal effect (not shown) is approximately zero and insignificant. In conclusion,

[4]These results are inconsistent with Baron et al. (2023) who find that credit growth is associated with more tail risk. The difference is likely due to the different sample and credit growth variable. Baron et al. (2023) standardize credit growth, censor credit growth at zero, and exclude observations around the two World Wars. Following this approach delivers a positive but insignificant coefficient.

the relationship between credit growth and our macro risk estimates, $\hat{\pi}$, is generally weak in our long sample. This explains why the inclusion of either local or global credit growth, or the exclusion of credit growth altogether, does not significantly affect our estimation.

Table A.IX: Credit Expansion and Future Growth (Main Sample, 1875–2020)

	$\Delta_3 y_{it+3}$	$\Delta_3 y_{it+3}$	$\Delta_3 c_{it+3}$	$\Delta_3 c_{it+3}$	$1_{\Delta_3 c_{it+3}<-2}$	$1_{\Delta_3 c_{it+3}<-2}$
$\Delta_3 d_{it-1}^{\text{JST}}$	-0.069^{**}	-0.068^{**}	-1.887^{***}	–	0.003	–
	(0.03)	(0.03)	(0.60)	–	(0.12)	–
$\Delta_3 d_{t-1}$	–	–	–	-0.427	–	-0.378^{*}
	–	–	–	(1.31)	–	(0.20)
FE	✓					
R^2	0.005	0.005	0.029	0.000	0.000	0.006
N	1,819	1,819	1,819	3,764	1,819	3,764

Notes. This table extends the replication and robustness analysis introduced in Table A.VIII above, using the sample period specified in the paper (1875 to 2020).

We also show $\hat{\pi}$ and credit growth jointly predict returns. Table A.X presents regression results for stock returns. In the top panel, we use local credit growth. Similarly to Baron and Xiong (2017), we observe a negative association between credit growth and future returns. In the bottom panel, we look at a bivariate specification in which we include credit growth and our baseline $\hat{\pi}$. We observe that both variables approximately retain their predictive power. In other words, the evidence that macro risk is associated with higher future excess returns is not inconsistent with the negative association with credit growth documented in the literature.

Table A.X: Credit Expansion and Future Returns

	Full sample (1876-2020)			Prewar (1876-1945)			Postwar (1945-2020)		
H	1	3	5	1	3	5	1	3	5
	$\sum_{h=1}^{H} r_{i,t+h} - r_{i,t+h}^{f} = a_i + b\Delta_3 d_{it-1}^{\text{JST}} + u_{i,t+H}$								
b	-0.27^{***}	-0.67^{***}	-0.84^{***}	-0.22^{*}	-0.79^{**}	-1.40^{***}	-0.37^{***}	-0.75^{**}	-0.66
	(0.08)	(0.22)	(0.33)	(0.12)	(0.34)	(0.48)	(0.11)	(0.30)	(0.42)
N	1,921	1,900	1,878	843	825	807	1,078	1,075	1,071
R^2	0.012	0.023	0.024	0.011	0.040	0.092	0.018	0.023	0.012
	$\sum_{h=1}^{H} r_{i,t+h} - r_{i,t+h}^{f} = a_i + b\hat{\pi}_{i,t} + c\Delta_3 d_{it-1}^{\text{JST}} + u_{i,t+H}$								
b	0.86^{**}	1.87^{*}	2.36^{*}	1.03^{**}	1.71	1.06	1.40^{**}	2.49^{*}	3.93^{**}
	(0.36)	(0.97)	(1.36)	(0.46)	(1.35)	(2.11)	(0.65)	(1.30)	(1.75)
c	-0.21^{***}	-0.49^{**}	-0.59^{*}	-0.05	-0.55^{*}	-1.29^{**}	-0.37^{***}	-0.59^{**}	-0.29
	(0.08)	(0.23)	(0.33)	(0.12)	(0.33)	(0.49)	(0.11)	(0.29)	(0.45)
N	1,661	1,628	1,594	583	553	523	1,078	1,075	1,071
R^2	0.024	0.038	0.037	0.038	0.063	0.105	0.031	0.042	0.048

Notes. This table presents slope coefficients, standard errors, and adjusted R^2 statistics for predictive panel regressions of cumulative excess returns on credit growth, following Baron and Xiong (2017). Observations are over the sample of 42 countries, from 1900 to 2020. Results are reported for the full sample, as well as for subsamples covering the pre- and post-WW2 periods. Bootstrapped standard errors are shown in parentheses.
*, **, and *** indicate statistical significance at 10%, 5%, and 1% levels, respectively.

Appendix G. The Levered Consumption Model

The dynamics of consumption belong to the affine class and are given by

$$\begin{aligned}\Delta c_t &= \mu + \sigma\varepsilon_t + v_t,\\ \pi_t &= (1-\rho)\bar{\pi} + \rho\pi_{t-1} + \nu\sqrt{\pi_{t-1}}u_t.\end{aligned}$$

Therefore, the following expectation has exponential affine solution (Drechsler and Yaron, 2011):

$$\mathbb{E}_t[e^{u_1\Delta c_{t+1}+u_2\pi_{t+1}}] = e^{g_0(u)+g_1(u)'[\Delta c_t,\pi_t]'}, \qquad u=(u_1,u_2)'\in\mathbb{R}^2,$$

where

$$\begin{aligned}g_0(u) &= \big(\mu-\tfrac{\sigma^2}{2}\big)u_1 + \bar{\pi}(1-\rho)u_2 + \tfrac{1}{2}\sigma^2u_1^2,\\ g_1(u) &= \big[0, \tfrac{1}{2}\nu^2u_2^2 + \rho u_2 + \varphi(u_1) - 1\big]'.\end{aligned}$$

The representative agent has recursive utility of the form

$$V_t = \big[(1-\delta)C_t^{1-1/\psi} + \delta(\mathbb{E}_t[V_{t+1}^{1-\gamma}])^{\frac{1-1/\psi}{1-\gamma}}\big]^{1/(1-1/\psi)}.$$

Normalized utility obtains if we take the limit as $\psi\to 1$, divide by C_t, rearrange, and then take the logarithm:

$$vc_t = \frac{\delta}{1-\gamma}\log(\mathbb{E}_t[e^{(1-\gamma)(\Delta c_{t+1}+vc_{t+1})}]).$$

Hypothesizing an affine form for vc_t,

$$vc_t = v_0 + v_c\Delta c_t + v_\pi\pi_t,$$

we can substitute to obtain

$$v_0 + v_c\Delta c_t + v_\pi\pi_t = \frac{\delta}{1-\gamma}\log(\mathbb{E}_t[e^{(1-\gamma)(\Delta c_{t+1}+v_0+v_c\Delta c_{t+1}+v_\pi\pi_{t+1})}])$$

and then compute the expectation on the right-hand side:

$$\begin{aligned}v_0 + v_c\Delta c_t + v_\pi\pi_t = \frac{\delta}{1-\gamma}\big((1-\gamma)v_0 &+ g_0([(1-\gamma)(1+v_c),(1-\gamma)(v_\pi)]')\\ &+ g_1([(1-\gamma)(1+v_c),(1-\gamma)v_\pi]')'[\Delta c_t,\pi_t]'\big).\end{aligned}$$

Finally, we solve for the coefficients:

$$\begin{aligned}v_c &: v_c = 0,\\ v_\pi &: v_\pi = v_\pi\delta\rho + \delta(1-\gamma)v_\pi^2\nu^2/2 + \tfrac{\delta}{1-\gamma}(\varphi(1-\gamma)-1),\\ v_0 &: v_0 = v_0\delta + \mu\delta + v_\pi\delta(1-\rho)\bar{\pi} - \delta\gamma\sigma^2/2.\end{aligned}$$

To derive asset prices, we first solve for the stochastic discount factor:

$$\begin{aligned}M_{t+1} &= \delta\frac{e^{-\gamma\Delta c_{t+1}+(1-\gamma)vc_{t+1}}}{\mathbb{E}_t[e^{(1-\gamma)(\Delta c_{t+1}+vc_{t+1})}]}\\ &= \delta e^{-\gamma\Delta c_{t+1}+(1-\gamma)vc_{t+1}-(1-\gamma)v_0-g_0([1-\gamma,(1-\gamma)v_\pi]')-g_1([1-\gamma,(1-\gamma)v_\pi]')[\Delta c_t,\pi_t]'}.\end{aligned}$$

Since the risk-free rate satisfies $e^{-r_{f,t}} = \mathbb{E}_t[M_{t+1}]$, we obtain

$$\begin{aligned}e^{-r_{f,t}} = \delta e^{(1-\gamma)v_0+g_0([-\gamma,(1-\gamma)v_\pi]')+g_1([-\gamma,(1-\gamma)v_\pi]')[\Delta c_t,\pi_t]'}\\ \times e^{-(1-\gamma)v_0-g_0([1-\gamma,(1-\gamma)v_\pi]')-g_1([1-\gamma,(1-\gamma)v_\pi]')[\Delta c_t,\pi_t]'}\end{aligned}$$

and therefore

$$r_{f,t} = -\log\delta + \mu - \gamma\sigma^2 + \pi_t(\varphi(1-\gamma) - \varphi(-\gamma)).$$

Recall that dividends can be expressed as $\Delta d_t = \phi \Delta c_t$. Stock returns are therefore given by

$$\begin{aligned} r_{d,t+1} &= \log\tfrac{P_{t+1}+D_{t+1}}{P_t} = \log\tfrac{P_{t+1}/D_{t+1}+1}{P_t/D_t} + \log\tfrac{D_{t+1}}{D_t} \\ &= \log(e^{-d_{t+1}+p_{t+1}}+1) + d_t - p_t + \Delta d_{t+1} \\ &\approx k_0 - k_1(d_{t+1}-p_{t+1}) + d_t - p_t + \Delta d_{t+1} \end{aligned}$$

for some endogenous constants k_0 and k_1 to be derived later. We posit an affine form for the logarithm of the D/P ratio:

$$d_t - p_t = A_0 + A_c \Delta c_t + A_\pi \pi_t.$$

The Euler equation for the stock (i.e., the claim asset on D_t) is $1 = \mathbb{E}_t[M_{t+1} e^{r_{d,t+1}}]$. We now plug in M_{t+1}, the log-linearized $r_{d,t+1}$, and our affine guesses for $d_t - p_t$ and $d_{t+1} - p_{t+1}$:

$$\begin{aligned} 1 = \mathbb{E}_t[\delta e^{-\gamma\Delta c_{t+1} + (1-\gamma)v c_{t+1} - (1-\gamma)v_0 - g_0([1-\gamma,(1-\gamma)v_\pi]') - g_1([1-\gamma,(1-\gamma)v_\pi]')[\Delta c_t, \pi_t]'} \\ \times e^{k_0 - k_1(A_0 + A_c\Delta c_{t+1} + A_\pi \pi_{t+1}) + (A_0 + A_c \Delta c_t + A_\pi \pi_t) + \phi \Delta c_{t+1}}]. \end{aligned}$$

Then we rearrange terms and solve the expectation as follows:

$$\begin{aligned} 1 &= \delta e^{-g_0([1-\gamma,(1-\gamma)v_\pi]') - g_1([1-\gamma,(1-\gamma)v_\pi]')[\Delta c_t,\pi_t]' + k_0 - k_1 A_0 + (A_0 + A_c\Delta c_t + A_\pi \pi_t)} \\ &\quad \times \mathbb{E}_t[e^{-\gamma\Delta c_{t+1} + (1-\gamma)v_\pi \pi_{t+1} - k_1(A_c \Delta c_{t+1} + A_\pi \pi_{t+1}) + \phi \Delta c_{t+1}}] \\ &= \delta e^{-g_0([1-\gamma,(1-\gamma)v_\pi]') - g_1([1-\gamma,(1-\gamma)v_\pi]')[\Delta c_t,\pi_t]' + k_0 - k_1 A_0 + (A_0 + A_c\Delta c_t + A_\pi \pi_t)} \\ &\quad \times e^{g_0([\phi-\gamma-k_1A_c,(1-\gamma)v_\pi - k_1 A_\pi]') + g_1([\phi-\gamma-k_1A_c,(1-\gamma)v_\pi - k_1 A_\pi]')[\Delta c_t,\pi_t]'}. \end{aligned}$$

Finally, we solve for the coefficients of the D/P ratio and the log-linearization constants:

$$\begin{aligned} k_0 &= -k_1 \log(k_1) - (1-k_1)\log(1-k_1), \\ \log k_1 &= \log(1-k_1) - A_0 - A_c \mathbb{E}[\Delta c_t] - A_\pi \mathbb{E}[\pi_t]. \end{aligned}$$

The coefficients satisfy

$$\begin{aligned} A_c &: \quad A_c = 0, \\ A_\pi &: \quad A_\pi = g_1([1-\gamma,(1-\gamma)v_\pi]')[0,1]' - g_1([\phi-\gamma-k_1A_c,(1-\gamma)v_\pi - k_1A_\pi]')[0,1]', \\ A_0 &: \quad A_0 = -\log\delta - k_0 + k_1 A_0 - g_0([\phi-\gamma-k_1A_c,(1-\gamma)v_\pi - k_1A_\pi]') + g_0([1-\gamma,(1-\gamma)v_\pi]'). \end{aligned}$$

Note that one must solve simultaneously for k_1 and A_π and then for k_0 and A_0. Hence the equity premium is given by

$$\begin{aligned} &\log \mathbb{E}_t[e^{r_{d,t+1}}] - r_{f,t} \\ &\quad = \log(\mathbb{E}_t[e^{r_{d,t+1}}]\mathbb{E}_t[M_{t+1}]) \\ &\quad = \log \mathbb{E}_t[e^{r^J_{d,t+1}}] + \log \mathbb{E}_t[M^J_{t+1}] - \log \mathbb{E}_t[e^{r^J_{d,t+1}} M^J_{t+1}] - \operatorname{cov}_t[r^C_{d,t+1}, m^C_{t+1}] \\ &\quad = [\phi, -k_1 A_\pi] \begin{bmatrix} \sigma^2 & 0 \\ 0 & \nu^2 \pi_t \end{bmatrix} [\gamma, (\gamma-1)v_\pi]' + \pi_t[\varphi(\phi) + \varphi(-\gamma) - \varphi(\phi-\gamma) - 1], \end{aligned}$$

where the superscripts C and J denote (respectively) the normal and nonnormal components. The return variance can be expressed as

$$\operatorname{var}_t[r_{d,t+1}] = [\phi, -k_1 A_\pi] \begin{bmatrix} \sigma^2 & 0 \\ 0 & \nu^2 \pi_t \end{bmatrix} [\phi, -k_1 A_\pi]' + \phi^2 \big(\tfrac{\partial^2}{\partial u^2} \varphi(u)\big|_{u=0}\big) \pi_t.$$

Finally, the risk-neutral return variance satisfies

$$\operatorname{var}^{\mathbb{Q}}_t[r_{d,t+1}] = [\phi, -k_1 A_\pi] \begin{bmatrix} \sigma^2 & 0 \\ 0 & \nu^2 \pi_t \end{bmatrix} [\phi, -k_1 A_\pi]' + \phi^2 \big(\tfrac{\partial^2}{\partial u^2} \varphi(u)\big|_{u=-\gamma}\big) \pi_t$$

(Drechsler and Yaron, 2011).

Appendix H. The Co-Integration Model

The dynamics of consumption and dividends belong to the affine class and are given by

$$\begin{aligned}\Delta c_t &= \mu + \sigma\varepsilon_t + v_t,\\ \Delta d_t &= \Delta c_t + \Delta s_t,\\ s_t &= (1-\phi)\bar{s} + \phi s_{t-1} + \eta z_t + \tilde{v}_t,\\ \pi_t &= (1-\rho)\bar{\pi} + \rho\pi_{t-1} + \nu\sqrt{\pi_{t-1}}u_t,\end{aligned}$$

where ε_t, z_t, and u_t and mutually independent standard normal random variables and

$$\begin{aligned}v_t &= J_t \mathbf{1}_{\Delta n_t>0},\\ \tilde{v}_t &= \kappa J_t \mathbf{1}_{\Delta n_t>0},\end{aligned}$$

for $\kappa > 0$. Therefore, the following expectation has exponential affine solution (Drechsler and Yaron, 2011):

$$\mathbb{E}_t[e^{u_1\Delta c_{t+1}+u_2\pi_{t+1}+u_3 s_{t+1}}] = e^{g_0(u)+g_1(u)'[\Delta c_t,\pi_t,s_t]'}, \qquad u=(u_1,u_2,u_3)'\in\mathbb{R}^3,$$

where

$$\begin{aligned}g_0(u) &= \big(\mu - \tfrac{\sigma^2}{2}\big)u_1 + \bar{\pi}(1-\rho)u_2 + \bar{s}(1-\phi)u_3 + \tfrac{1}{2}\big(\sigma^2u_1^2 + \eta^2u_3^2\big),\\ g_1(u) &= \big[0, \tfrac{1}{2}\nu^2u_2^2 + \rho u_2 + \varphi(u_1) + \varphi(\kappa u_3) - 2, \phi u_3\big]'.\end{aligned}$$

The representative agent has recursive utility of the form

$$V_t = \big[(1-\delta)C_t^{1-1/\psi} + \delta(\mathbb{E}_t[V_{t+1}^{1-\gamma}])^{\frac{1-1/\psi}{1-\gamma}}\big]^{1/(1-1/\psi)}.$$

Normalized utility obtains if we take the limit as $\psi \to 1$, divide by C_t, rearrange, and then take the logarithm:

$$vc_t = \frac{\delta}{1-\gamma}\log(\mathbb{E}_t[e^{(1-\gamma)(\Delta c_{t+1}+vc_{t+1})}]).$$

Hypothesizing an affine form for vc_t,

$$vc_t = v_0 + v_c\Delta c_t + v_\pi\pi_t + v_s s_t,$$

we can substitute to obtain

$$v_0 + v_c\Delta c_t + v_\pi\pi_t + v_s s_t = \frac{\delta}{1-\gamma}\log(\mathbb{E}_t[e^{(1-\gamma)(\Delta c_{t+1}+v_0+v_c\Delta c_{t+1}+v_\pi\pi_{t+1}+v_s s_{t+1})}])$$

and then compute the expectation on the right-hand side:

$$\begin{aligned}v_0 + v_c\Delta c_t + v_\pi\pi_t + v_s s_t = \frac{\delta}{1-\gamma}\big((1-\gamma)v_0 &+ g_0([(1-\gamma)(1+v_c),(1-\gamma)v_\pi,(1-\gamma)v_s]')\\ &+ g_1([(1-\gamma)(1+v_c),(1-\gamma)v_\pi,(1-\gamma)v_s]')'[\Delta c_t,\pi_t,s_t]'\big).\end{aligned}$$

Finally, we solve for the coefficients:

$$\begin{aligned}v_c &: v_c = 0,\\ v_s &: v_s = 0,\\ v_\pi &: v_\pi = v_\pi\delta\rho + \delta(1-\gamma)v_\pi^2\nu^2/2 + \tfrac{\delta}{1-\gamma}(\varphi(1-\gamma)-1),\\ v_0 &: v_0 = v_0\delta + \mu\delta + v_\pi\delta(1-\rho)\bar{\pi} - \delta\gamma\sigma^2/2.\end{aligned}$$

To derive asset prices, we first solve for the stochastic discount factor:

$$\begin{aligned} M_{t+1} &= \delta \frac{e^{-\gamma \Delta c_{t+1} + (1-\gamma) v c_{t+1}}}{\mathbb{E}_t[e^{(1-\gamma)(\Delta c_{t+1} + v c_{t+1})}]} \\ &= \delta e^{-\gamma \Delta c_{t+1} + (1-\gamma) v c_{t+1} - (1-\gamma) v_0 - g_0([1-\gamma, (1-\gamma) v_\pi, (1-\gamma) v_s]') - g_1([1-\gamma, (1-\gamma) v_\pi, (1-\gamma) v_s]')[\Delta c_t, \pi_t, s_t]'}. \end{aligned}$$

Since the risk-free rate satisfies $e^{-r_{f,t}} = \mathbb{E}_t[M_{t+1}]$, we obtain

$$\begin{aligned} e^{-r_{f,t}} = \delta e^{(1-\gamma) v_0 + g_0([-\gamma, (1-\gamma) v_\pi, (1-\gamma) v_s]') + g_1([-\gamma, (1-\gamma) v_\pi, (1-\gamma) v_s]')[\Delta c_t, \pi_t, s_t]'} \\ \times e^{-(1-\gamma) v_0 - g_0([1-\gamma, (1-\gamma) v_\pi, (1-\gamma) v_s]') - g_1([1-\gamma, (1-\gamma) v_\pi, (1-\gamma) v_s]')[\Delta c_t, \pi_t, s_t]'} \end{aligned}$$

and therefore

$$r_{f,t} = -\log \delta + \mu - \gamma \sigma^2 + \pi_t(\varphi(1-\gamma) - \varphi(-\gamma)).$$

Stock returns are therefore given by

$$\begin{aligned} r_{d,t+1} &= \log \tfrac{P_{t+1} + D_{t+1}}{P_t} = \log \tfrac{P_{t+1}/D_{t+1} + 1}{P_t / D_t} + \log \tfrac{D_{t+1}}{D_t} \\ &= \log(e^{-d_{t+1} + p_{t+1}} + 1) + d_t - p_t + \Delta d_{t+1} \\ &\approx k_0 - k_1 (d_{t+1} - p_{t+1}) + d_t - p_t + \Delta d_{t+1} \end{aligned}$$

for some endogenous constants k_0 and k_1 to be derived later. We posit an affine form for the logarithm of the D/P ratio:

$$d_t - p_t = A_0 + A_c \Delta c_t + A_\pi \pi_t + A_s s_t.$$

The Euler equation for the stock (i.e., the claim asset on D_t) is $1 = \mathbb{E}_t[M_{t+1} e^{r_{d,t+1}}]$. We now plug in M_{t+1}, the log-linearized $r_{d,t+1}$, and our affine guesses for $d_t - p_t$ and $d_{t+1} - p_{t+1}$:

$$\begin{aligned} 1 = \mathbb{E}_t[\delta e^{-\gamma \Delta c_{t+1} + (1-\gamma) v c_{t+1} - (1-\gamma) v_0 - g_0([1-\gamma, (1-\gamma) v_\pi, (1-\gamma) v_s]') - g_1([1-\gamma, (1-\gamma) v_\pi, (1-\gamma) v_s]')[\Delta c_t, \pi_t, s_t]'} \\ \times e^{k_0 - k_1(A_0 + A_c \Delta c_{t+1} + A_\pi \pi_{t+1} + A_s s_{t+1}) + (A_0 + A_c \Delta c_t + A_\pi \pi_t + A_s s_t) + \Delta c_{t+1} + \Delta s_{t+1}}]. \end{aligned}$$

Then we rearrange terms and solve the expectation as follows:

$$\begin{aligned} 1 &= \delta e^{-g_0([1-\gamma, (1-\gamma) v_\pi, (1-\gamma) v_s]') - g_1([1-\gamma, (1-\gamma) v_\pi, (1-\gamma) v_s]')[\Delta c_t, \pi_t, s_t]' + k_0 - k_1 A_0 + (A_0 + A_c \Delta c_t + A_\pi \pi_t + A_s s_t)} \\ &\quad \times \mathbb{E}_t[e^{-\gamma \Delta c_{t+1} + (1-\gamma) v_\pi \pi_{t+1} + (1-\gamma) v_s s_{t+1} - k_1(A_c \Delta c_{t+1} + A_\pi \pi_{t+1} + A_s s_{t+1}) + \Delta c_{t+1} + \Delta s_{t+1}}] \\ &= \delta e^{-g_0([1-\gamma, (1-\gamma) v_\pi, (1-\gamma) v_s]') - g_1([1-\gamma, (1-\gamma) v_\pi, (1-\gamma) v_s]')[\Delta c_t, \pi_t, s_t]' + k_0 - k_1 A_0 + (A_0 + A_c \Delta c_t + A_\pi \pi_t + (A_s - 1) s_t)} \\ &\quad \times e^{g_0([1-\gamma-k_1 A_c, (1-\gamma) v_\pi - k_1 A_\pi, 1 + (1-\gamma) v_s - k_1 A_s]') + g_1([1-\gamma-k_1 A_c, (1-\gamma) v_\pi - k_1 A_\pi, 1 + (1-\gamma) v_s - k_1 A_s]')[\Delta c_t, \pi_t, s_t]'}. \end{aligned}$$

Finally, we solve for the coefficients of the D/P ratio and the log-linearization constants:

$$\begin{aligned} k_0 &= -k_1 \log(k_1) - (1 - k_1) \log(1 - k_1), \\ \log k_1 &= \log(1 - k_1) - A_0 - A_c \mathbb{E}[\Delta c_t] - A_\pi \mathbb{E}[\pi_t] - A_s \mathbb{E}[s_t]. \end{aligned}$$

The coefficients satisfy

$$\begin{aligned} A_c &: A_c = 0, \\ A_\pi &: A_\pi = g_1([1-\gamma, (1-\gamma) v_\pi, (1-\gamma) v_s]')[0,1,0]' \\ &\qquad - g_1([1-\gamma-k_1 A_c, (1-\gamma) v_\pi - k_1 A_\pi, 1 + (1-\gamma) v_s - k_1 A_s]')[0,1,0]', \\ A_s &: A_s = g_1([1-\gamma, (1-\gamma) v_\pi, (1-\gamma) v_s]')[0,0,1]' - 1 \\ &\qquad - g_1([1-\gamma-k_1 A_c, (1-\gamma) v_\pi - k_1 A_\pi, 1 + (1-\gamma) v_s - k_1 A_s]')[0,0,1]', \\ A_0 &: A_0 = -\log \delta - k_0 + k_1 A_0 - g_0([1-\gamma-k_1 A_c, (1-\gamma) v_\pi - k_1 A_\pi, 1 + (1-\gamma) v_s - k_1 A_s]') \\ &\qquad + g_0([1-\gamma, (1-\gamma) v_\pi, (1-\gamma) v_s]'). \end{aligned}$$

Hence the equity premium is given by

$$\begin{aligned}
&\log \mathbb{E}_t[e^{r_{d,t+1}}] - r_{f,t} \\
&\quad = \log(\mathbb{E}_t[e^{r_{d,t+1}}]\mathbb{E}_t[M_{t+1}]) \\
&\quad = \log \mathbb{E}_t[e^{r^J_{d,t+1}}] + \log \mathbb{E}_t[M^J_{t+1}] - \log \mathbb{E}_t[e^{r^J_{d,t+1}} M^J_{t+1}] - \mathrm{cov}_t[r^C_{d,t+1}, m^C_{t+1}] \\
&\quad = [1, -k_1 A_\pi, 1 - k_1 A_s] \begin{bmatrix} \sigma^2 & 0 & 0 \\ 0 & \nu^2 \pi_t & 0 \\ 0 & 0 & \eta^2 \end{bmatrix} [\gamma, (\gamma-1)v_\pi, (\gamma-1)v_s]' \\
&\qquad + \pi_t[\varphi(1 + \kappa(1 - k_1 A_s)) - \varphi(1 + \kappa(1 - k_1 A_s) - \gamma) + \varphi(-\gamma) - 1],
\end{aligned}$$

where the superscripts C and J denote (respectively) the normal and nonnormal components. The return variance can be expressed as

$$\begin{aligned}
\mathrm{var}_t[r_{d,t+1}] =& [1, -k_1 A_\pi, 1 - k_1 A_s] \begin{bmatrix} \sigma^2 & 0 & 0 \\ 0 & \nu^2 \pi_t & 0 \\ 0 & 0 & \eta^2 \end{bmatrix} [1, -k_1 A_\pi, 1 - k_1 A_s]' \\
&+ (1 + \kappa(1 - k_1 A_s))^2 \big(\tfrac{\partial^2}{\partial u^2} \varphi(u)\big|_{u=0}\big) \pi_t.
\end{aligned}$$

Finally, the risk-neutral return variance satisfies

$$\begin{aligned}
\mathrm{var}^{\mathbb{Q}}_t[r_{d,t+1}] =& [1, -k_1 A_\pi, 1 - k_1 A_s] \begin{bmatrix} \sigma^2 & 0 & 0 \\ 0 & \nu^2 \pi_t & 0 \\ 0 & 0 & \eta^2 \end{bmatrix} [1, -k_1 A_\pi, 1 - k_1 A_s]' \\
&+ (1 + \kappa(1 - k_1 A_s))^2 \big(\tfrac{\partial^2}{\partial u^2} \varphi(u)\big|_{u=-\gamma}\big) \pi_t,
\end{aligned}$$

(Drechsler and Yaron, 2011).

Appendix I. The Levered Consumption Model with CRRA Preferences

Under CRRA preferences of the represenative agent the stochastic discount factor simplifies to:

$$M_{t+1} = \delta e^{-\gamma \Delta c_{t+1}}.$$

Since the risk-free rate satisfies $e^{-r_{f,t}} = \mathbb{E}_t[M_{t+1}]$, we obtain

$$e^{-r_{f,t}} = \delta e^{g_0([-\gamma,0]') + g_1([-\gamma,0]')[\Delta c_t, \pi_t]'}$$

and therefore

$$r_{f,t} = -\log \delta + \gamma\mu - \gamma(1+\gamma)\sigma^2 - \pi_t \varphi(1-\gamma).$$

Recall that dividends can be expressed as $\Delta d_t = \phi \Delta c_t$. Stock returns are therefore given by

$$\begin{aligned}
r_{d,t+1} &= \log \tfrac{P_{t+1}+D_{t+1}}{P_t} = \log \tfrac{P_{t+1}/D_{t+1}+1}{P_t/D_t} + \log \tfrac{D_{t+1}}{D_t} \\
&= \log(e^{-d_{t+1}+p_{t+1}} + 1) + d_t - p_t + \Delta d_{t+1} \\
&\approx k_0 - k_1(d_{t+1} - p_{t+1}) + d_t - p_t + \Delta d_{t+1}
\end{aligned}$$

for some endogenous constants k_0 and k_1 to be derived later. We posit an affine form for the logarithm of the D/P ratio:

$$d_t - p_t = A_0 + A_c \Delta c_t + A_\pi \pi_t.$$

The Euler equation for the stock (i.e., the claim asset on D_t) is $1 = \mathbb{E}_t[M_{t+1} e^{r_{d,t+1}}]$. We now plug in M_{t+1}, the log-linearized $r_{d,t+1}$, and our affine guesses for $d_t - p_t$ and $d_{t+1} - p_{t+1}$:

$$\begin{aligned}
1 = \mathbb{E}_t[&\delta e^{-\gamma \Delta c_{t+1} - g_0([1-\gamma,0]') - g_1([1-\gamma,0]')[\Delta c_t, \pi_t]'} \\
&\times e^{k_0 - k_1(A_0 + A_c \Delta c_{t+1} + A_\pi \pi_{t+1}) + (A_0 + A_c \Delta c_t + A_\pi \pi_t) + \phi \Delta c_{t+1}}].
\end{aligned}$$

Then we rearrange terms and solve the expectation as follows:

$$\begin{aligned}1 &= \delta e^{-g_0([1-\gamma,0]')-g_1([1-\gamma,0]')[\Delta c_t,\pi_t]'+k_0-k_1A_0+(A_0+A_c\Delta c_t+A_\pi\pi_t)} \\ &\quad\times \mathbb{E}_t[e^{-\gamma\Delta c_{t+1}-k_1(A_c\Delta c_{t+1}+A_\pi\pi_{t+1})+\phi\Delta c_{t+1}}] \\ &= \delta e^{-g_0([1-\gamma,0]')-g_1([1-\gamma,0]')[\Delta c_t,\pi_t]'+k_0-k_1A_0+(A_0+A_c\Delta c_t+A_\pi\pi_t)} \\ &\quad\times e^{g_0([\phi-\gamma-k_1A_c,-k_1A_\pi]')+g_1([\phi-\gamma-k_1A_c,-k_1A_\pi]')[\Delta c_t,\pi_t]'}.\end{aligned}$$

Finally, we solve for the coefficients of the D/P ratio and the log-linearization constants:

$$\begin{aligned}k_0 &= -k_1\log(k_1)-(1-k_1)\log(1-k_1), \\ \log k_1 &= \log(1-k_1)-A_0-A_c\mathbb{E}[\Delta c_t]-A_\pi\mathbb{E}[\pi_t].\end{aligned}$$

The coefficients satisfy

$$\begin{aligned}A_c &: A_c = 0, \\ A_\pi &: A_\pi = g_1([1-\gamma,0]')[0,1]'-g_1([\phi-\gamma-k_1A_c,-k_1A_\pi]')[0,1]', \\ A_0 &: A_0 = -\log\delta-k_0+k_1A_0-g_0([\phi-\gamma-k_1A_c,-k_1A_\pi]')+g_0([1-\gamma,0]').\end{aligned}$$

Note that one must solve simultaneously for k_1 and A_π and then for k_0 and A_0. Hence the equity premium is given by

$$\begin{aligned}&\log\mathbb{E}_t[e^{r_{d,t+1}}]-r_{f,t} \\ &\quad= \log(\mathbb{E}_t[e^{r_{d,t+1}}]\mathbb{E}_t[M_{t+1}]) \\ &\quad= \log\mathbb{E}_t[e^{r^J_{d,t+1}}]+\log\mathbb{E}_t[M^J_{t+1}]-\log\mathbb{E}_t[e^{r^J_{d,t+1}}M^J_{t+1}]-\mathrm{cov}_t[r^C_{d,t+1},m^C_{t+1}] \\ &\quad= [\phi,-k_1A_\pi]\begin{bmatrix}\sigma^2 & 0\\ 0 & \nu^2\pi_t\end{bmatrix}[\gamma,0]'+\pi_t[\varphi(\phi)+\varphi(-\gamma)-\varphi(\phi-\gamma)-1],\end{aligned}$$

where the superscripts C and J denote (respectively) the normal and nonnormal components. The return variance can be expressed as

$$\mathrm{var}_t[r_{d,t+1}] = [\phi,-k_1A_\pi]\begin{bmatrix}\sigma^2 & 0\\ 0 & \nu^2\pi_t\end{bmatrix}[\phi,-k_1A_\pi]'+\phi^2\big(\tfrac{\partial^2}{\partial u^2}\varphi(u)\big|_{u=0}\big)\pi_t.$$

Finally, the risk-neutral return variance satisfies

$$\mathrm{var}^{\mathbb{Q}}_t[r_{d,t+1}] = [\phi,-k_1A_\pi]\begin{bmatrix}\sigma^2 & 0\\ 0 & \nu^2\pi_t\end{bmatrix}[\phi,-k_1A_\pi]'+\phi^2\big(\tfrac{\partial^2}{\partial u^2}\varphi(u)\big|_{u=-\gamma}\big)\pi_t$$

(Drechsler and Yaron, 2011).

Table A.XI: Autocorrelations: Levered Consumption Model

Lag	Data		Model				
	Estimate	SE	2.5%	5%	50%	95%	97.5%
			Panel A: Consumption				
1	0.10	0.12	−0.05	−0.04	0.02	0.08	0.10
2	0.10	0.10	−0.05	−0.04	0.01	0.07	0.09
3	−0.10	0.08	−0.05	−0.05	0.01	0.07	0.08
4	−0.01	0.13	−0.06	−0.05	0.01	0.06	0.07
5	−0.12	0.09	−0.06	−0.05	0.00	0.06	0.07
6	0.09	0.08	−0.06	−0.05	0.00	0.06	0.07
7	−0.07	0.07	−0.05	−0.05	0.00	0.06	0.07
8	−0.17	0.07	−0.06	−0.05	0.00	0.05	0.06
9	−0.07	0.08	−0.06	−0.05	0.00	0.06	0.07
10	0.09	0.08	−0.06	−0.05	0.00	0.05	0.07
$\sum$ 1^{st} *to* 5^{th} *lag*	−0.03				0.05		
			Panel B: Dividends				
1	0.30	0.12	−0.05	−0.04	0.02	0.08	0.10
2	−0.06	0.10	−0.05	−0.04	0.01	0.07	0.09
3	−0.20	0.07	−0.05	−0.05	0.01	0.07	0.08
4	−0.19	0.13	−0.06	−0.05	0.01	0.06	0.07
5	−0.18	0.08	−0.06	−0.05	0.00	0.06	0.07
6	−0.01	0.09	−0.06	−0.05	0.00	0.06	0.07
7	0.05	0.07	−0.05	−0.05	0.00	0.06	0.07
8	0.07	0.08	−0.06	−0.05	0.00	0.05	0.06
9	0.07	0.06	−0.06	−0.05	0.00	0.06	0.07
10	0.10	0.06	−0.06	−0.05	0.00	0.05	0.07
$\sum$ 1^{st} *to* 5^{th} *lag*	−0.33				0.05		
			Panel C: Excess Returns				
1	0.02	0.09	−0.19	−0.18	−0.11	−0.03	−0.02
2	−0.19	0.06	−0.16	−0.15	−0.08	−0.01	0.00
3	0.08	0.10	−0.14	−0.12	−0.06	0.01	0.02
4	−0.04	0.13	−0.12	−0.11	−0.04	0.02	0.03
5	−0.18	0.07	−0.10	−0.09	−0.03	0.03	0.04
6	0.08	0.08	−0.09	−0.08	−0.02	0.04	0.06
7	0.14	0.07	−0.08	−0.07	−0.02	0.04	0.05
8	−0.04	0.09	−0.08	−0.07	−0.01	0.05	0.06
9	0.05	0.08	−0.07	−0.06	−0.01	0.05	0.06
10	0.14	0.07	−0.07	−0.06	−0.01	0.05	0.06
$\sum$ 1^{st} *to* 5^{th} *lag*	−0.31				−0.32		

Notes. Panel A reports unconditional moment statistics from S&P 500 real returns and three-month U.S. Treasury real rates as well as percentiles of the same moments from model simulations of the levered consumption model. Panel B reports the same quantities from a restricted model in which disasters do not realize. Parameters are from Table VII. Additional parameters are $\gamma = 5, \delta = 99\%$, and $\phi = 2.6$.

Table A.XII: Autocorrelations: Co-Integration Model

Lag	Data		Model				
	Estimate	SE	2.5%	5%	50%	95%	97.5%
			Panel A: Consumption				
1	0.10	0.12	−0.05	−0.04	0.02	0.08	0.10
2	0.10	0.10	−0.05	−0.05	0.01	0.07	0.08
3	−0.10	0.08	−0.06	−0.05	0.01	0.07	0.09
4	−0.01	0.13	−0.05	−0.04	0.00	0.07	0.08
5	−0.12	0.09	−0.06	−0.05	0.00	0.06	0.08
6	0.09	0.08	−0.06	−0.05	0.00	0.06	0.07
7	−0.07	0.07	−0.06	−0.05	0.00	0.06	0.07
8	−0.17	0.07	−0.06	−0.05	0.00	0.05	0.07
9	−0.07	0.08	−0.06	−0.05	0.00	0.05	0.06
10	0.09	0.08	−0.06	−0.05	0.00	0.05	0.06
$\sum$ *1st to 5th lag*	−0.03				0.04		
			Panel B: Dividends				
1	0.30	0.12	−0.11	−0.10	−0.04	0.01	0.02
2	−0.06	0.10	−0.10	−0.09	−0.04	0.01	0.02
3	−0.20	0.07	−0.10	−0.09	−0.04	0.02	0.02
4	−0.19	0.13	−0.10	−0.09	−0.04	0.02	0.03
5	−0.18	0.08	−0.09	−0.08	−0.03	0.02	0.04
6	−0.01	0.09	−0.09	−0.08	−0.03	0.03	0.04
7	0.05	0.07	−0.09	−0.08	−0.02	0.03	0.03
8	0.07	0.08	−0.09	−0.08	−0.02	0.03	0.04
9	0.07	0.06	−0.08	−0.07	−0.02	0.03	0.04
10	0.10	0.06	−0.09	−0.07	−0.02	0.03	0.04
$\sum$ *1st to 5th lag*	−0.31				−0.19		
			Panel C: Excess Returns				
1	0.02	0.09	−0.17	−0.15	−0.09	−0.02	0.00
2	−0.19	0.06	−0.13	−0.13	−0.06	0.01	0.02
3	0.08	0.10	−0.12	−0.11	−0.05	0.02	0.03
4	−0.04	0.13	−0.10	−0.09	−0.03	0.03	0.04
5	−0.18	0.07	−0.09	−0.08	−0.02	0.04	0.05
6	0.08	0.08	−0.08	−0.07	−0.02	0.04	0.05
7	0.14	0.07	−0.08	−0.07	−0.01	0.04	0.05
8	−0.04	0.09	−0.07	−0.06	−0.01	0.05	0.06
9	0.05	0.08	−0.07	−0.06	−0.01	0.04	0.06
10	0.14	0.07	−0.07	−0.06	−0.01	0.05	0.06
$\sum$ *1st to 5th lag*	−0.31				−0.25		

Notes. Panel A reports unconditional autocorrelations of consumption and dividend growth and excess returns of the S&P 500 real returns on the three-month U.S. Treasury real rate as well as percentiles of the same moments from model simulations of the co-integration model. Panel B reports the same quantities from a restricted model in which disasters do not realize. Parameters are from Table VII. Additional parameters are $\gamma = 5, \delta = 99\%, \phi = 0.90, \eta = 0.30$, and $\kappa = 4$.

Table A.XIII: Asset Pricing Moments: The CRRA Case

	Data		Model				
	1871-2015	1946-2015	2.5%	5%	50%	95%	97.5%
Panel A: Unrestricted Model							
Average risk-free rate	2.17	1.26	6.67	6.81	7.50	8.19	8.28
Standard deviation of risk-free rate	4.51	3.14	4.00	4.15	5.23	6.45	6.74
Average excess return	6.00	6.82	1.27	1.38	2.05	2.74	2.90
Standard deviation of excess return	18.5	16.5	11.4	11.6	12.9	14.4	14.6
Average dividend yield	4.13	3.33	4.90	4.91	4.95	4.99	5.00
Standard deviation of dividend yield	1.57	1.42	0.26	0.27	0.33	0.38	0.40
Autocorrelation of dividend yield	78.4	90.0	62.1	63.5	71.2	78.1	79.3
Panel B: Restricted Model (No Realized Disasters)							
Average risk-free rate	2.17	1.26	6.60	6.76	7.51	8.12	8.20
Standard deviation of risk-free rate	4.51	3.14	4.13	4.24	5.17	6.46	6.77
Average excess return	6.00	6.82	2.49	2.62	3.30	4.08	4.25
Standard deviation of excess return	18.5	16.5	10.1	10.3	11.2	12.3	12.6
Average dividend yield	4.13	3.33	4.89	4.90	4.95	4.99	4.99
Standard deviation of dividend yield	1.57	1.42	0.27	0.27	0.32	0.39	0.40
Autocorrelation of dividend yield	78.4	90.0	62.4	63.6	71.1	78.0	79.3

Notes. Panel A reports unconditional moment statistics from S&P 500 real returns and three-month U.S. Treasury real rates as well as percentiles of the same moments from model simulations of the CRRA model. Panel B reports the same quantities from a restricted model in which disasters do not realize. Parameters are from Table VII. Additional parameters are $\gamma = 5, \delta = 99\%$, and $\phi = 2.6$.

Table A.XIV: Option-Implied Volatility Skew

	Data	Levered Consumption		Co-Integration	
		π_t^-	π_{t-1}^-	π_t^-	π_{t-1}^-
Mean (%)	3.71	3.57	3.56	8.12	8.06
Standard Deviation (%)	0.83	0.29	0.28	1.26	1.26

Notes. This table reports the sample mean and standard deviation of the option-implied volatility skew, defined as the difference in implied volatility of options with moneyness equal to -0.05 and 0.05 respectively (where moneyness is (strike − price) / price). Data represents the daily sample of S&P 500 index put options between 1/4/1996 and 12/29/2017. The model-implied measures of the option-implied volatility skew are obtained by plugging in our yearly estimates of $\hat{\pi}^-$, either from the current or the previous year, on the sample 1996-2017. The levered consumption and co-integration models and their calibration are described in Section IV.

References

Baron, Matthew, and Wei Xiong, 2017, Credit expansion and neglected crash risk, *Quarterly Journal of Economics* 132, 713–764.

Baron, Matthew, Wei Xiong, and Zhijiang Ye, 2023, Measuring Time-Varying Disaster Risk: An Empirical Analysis of Dark Matter in Asset Prices.

Barro, Robert J., and Gordon Y. Liao, 2021, Rare disaster probability and options pricing, *Journal of Financial Economics* 139, 750–769.

Barro, Robert J., and José F. Ursúa, 2008, Macroeconomic Crises since 1870, *Brookings Papers on Economic Activity* 2008, 255–350.

Berkman, Henk, Ben Jacobsen, and John B. Lee, 2011, Time-Varying Rare Disaster Risk and Stock Returns, *Journal of Financial Economics* 101, 313–332.

Bloom, Nicholas, 2009, The Impact of Uncertainty Shocks, *Econometrica* 77, 623–685.

Drechsler, Itamar, and Amir Yaron, 2011, What's vol got to do with it, *Review of Financial Studies* 24, 1–45.

Inklaar, R., H. de Jong, J. Bolt, and J. van Zanden, 2018, Rebasing 'Maddison': new income comparisons and the shape of long-run economic development, Maddison Project Working paper 10.

Jordà, Òscar, Moritz Schularick, and Alan M. Taylor, 2017, Macrofinancial history and the new business cycle facts, *NBER Macroeconomics Annual* 31, 213–263.

Manela, Asaf, and Alan Moreira, 2017, News implied volatility and disaster concerns, *Journal of Financial Economics* 123, 137–162.

Mayer, Thierry, and Soledad Zignago, 2011, Notes on CEPII's Distances Measures: The GeoDist Database, CEPII Working Paper N°2011-25, December 2011.

Mehra, Rajnish, and Edward C. Prescott, 1985, The equity premium: A puzzle, *Journal of monetary Economics* 15, 145–161.

Mian, Atif R., Amir Sufi, and Emil Verner, 2017, Household debt and business cycles worldwide, *Quarterly Journal of Economics* 1755–1817.

Mueller, Hannes, and Christopher Rauh, 2022, The Hard Problem of Prediction for Conflict Prevention, *Journal of the European Economic Association* 20, 2440–2467.

Officer, Robert R., 1973, The Variability of the Market Factor of the New York Stock Exchange, *Journal of Business* 46, 434.

Reinhart, Carmen M., and Kenneth S. Rogoff, 2009, *This Time is Different : A Panoramic View of Eight Centuries of* (Princeton University Press, Princeton, NJ).

Sarkees, Meredith Reid, and Frank Wayman, 2010, *Resort to War: 1816 - 2007* (CQ Press, Washington DC).

Schwert, G. William, 1989, Why Does Stock Market Volatility Change Over Time?, *Journal of Finance* 44, 1115–1153.

Tibshirani, Robert, 1996, Regression Shrinkage and Selection Via the Lasso, *Journal of the Royal Statistical Society: Series B (Methodological)* 58, 267–288.

Zou, Hui, and Trevor Hastie, 2005, Regularization and variable selection via the elastic net, *Journal of the Royal Statistical Society. Series B: Statistical Methodology* 67, 301–320.