Modeling Credit Contagion via the Updating of Fragile Beliefs*

Luca Benzoni† Pierre Collin-Dufresne‡ Robert S. Goldstein§ Jean Helwege¶

This Version: December 26, 2011

Abstract

We propose a tractable equilibrium model for pricing defaultable bonds that are subject to contagion risk. Contagion arises because agents with 'fragile beliefs' are uncertain about both the underlying state of the economy and the posterior probabilities associated with these states. As such, agents adopt a robust decision rule for updating that leads them to over-weight the posterior probabilities of 'bad' states. We estimate the model using panel data on sovereign Euro-zone CDS spreads during the recent crisis, and find that it captures levels and dynamics of spreads better than traditional affine models with the same number of observable and latent state variables.

*We thank Scott Brave, Richard Cantor, Sanjiv Das, Darrell Duffie, Lorenzo Garlappi, Alejandro Justiniano, David Lando, Eric Wan, Tan Wang, Fan Yu, and seminar participants at Moody's Advisory Research Committee, the Bank of International Settlements, Duke University, Carnegie Mellon University, the FDIC, HEC Montréal, London Business School, Groupe HEC, the Wharton School, the University of California at Berkeley, the University of Illinois at Urbana-Champaign, and the University of Texas at Dallas for helpful comments. Andrea Ajello, Olena Chyruk, Andy Fedak, Paymon Khorrami, Kuan Lee, Harvey Stephenson, Guang Yang, and Ludovico Zaraga provided excellent research assistance.

†Federal Reserve Bank of Chicago, lbenzoni@frbchi.org

‡Carson Family Professor of Finance, Columbia University, and NBER, pc2415@columbia.edu

§C. Arthur Williams Professor of Insurance, University of Minnesota, and NBER, golds144@umn.edu

¶J. Henry Fellers Professor of Business Administration, University of South Carolina, Jean.helwege@moore.sc.edu

Modeling Credit Contagion via the Updating of Fragile Beliefs

Abstract

We propose a tractable equilibrium model for pricing defaultable bonds that are subject to contagion risk. Contagion arises because agents with 'fragile beliefs' are uncertain about both the underlying state of the economy and the posterior probabilities associated with these states. As such, agents adopt a robust decision rule for updating that leads them to over-weight the posterior probabilities of 'bad' states. We estimate the model using panel data on sovereign Euro-zone CDS spreads during the recent crisis, and find that it captures levels and dynamics of spreads better than traditional affine models with the same number of observable and latent state variables.

1 Introduction

During the recent (and ongoing) Euro-zone crisis, the risk of contagion has often been cited as one of the major drivers of sovereign credit spreads. Indeed, typing the words "contagion and Euro" into a Google search returns over 750,000 results, many of which refer to articles from the financial press that relate changes in sovereign spreads of European nations to the risk or 'fear' of contagion. This dialogue raises many important questions, including:

- What is contagion risk, and what are its economic sources?
- Is there a risk-premium associated with contagion risk, and if so, what is its impact on sovereign spreads?
- To what extent is the co-movement in sovereign spreads driven by contagion risk and its risk-premium?

In this paper we propose a tractable equilibrium model in which contagion risk significantly impacts the level and the dynamics of sovereign credit spreads. Contagion arises because agents are uncertain about both the underlying state of the economy and their posterior probabilities associated with these states (they have 'fragile beliefs'). Following Hansen and Sargent (2007, 2010), we investigate agents that adopt a robust decision rule for updating that leads them to over-weight the posterior probabilities of 'bad' states and under-weight those of 'good' states. Together, these two ingredients (hidden states and fragile beliefs) can explain large sovereign spreads even if expected losses due to default are relatively small. Furthermore, the model can generate significant correlation in spreads even if common movements in macroeconomic fundamentals are relatively modest.

To be more specific, we assume there is a hidden state of nature which, if known, would impact the expected aggregate consumption growth as well as each country's default probability. Agents update the likelihood of each state based on all available information, but they are uncertain about the true data-generating process, and therefore are uncertain about the updated posterior probabilities assigned to each state. Following Hansen and Sargent (2007, 2010), we assume agents adopt robust decision rules to mitigate this 'model risk.' In particular, we assume agents use different risk-sensitivity operators to account for i) uncertainty in the model specification conditional on the state, and ii) uncertainty regarding the correct posterior distribution of the state itself.

The first component, preference for robustness regarding model parameters, has been more extensively studied in the literature. For example, it is well understood that, in a representative agent framework, the decision rule of a log-utility investor who uses an entropy penalty is observationally equivalent to that of an agent with recursive utility of the Epstein-Zin-Kreps-Porteus type (see, e.g., Barillas, Hansen and Sargent (2010)). However, the second component, which can be interpreted as a preference for robustness towards mis-specification in their beliefs, is less understood. As we show, one benefit of the 'fragile beliefs' specification we adopt is that it is very tractable, even for

fairly complex models. Indeed, when an agent has fragile beliefs, she values long-lived securities by first estimating their value as if she knew the 'true' state, and then she takes a weighted average of these values. This is very similar to the approach followed in the traditional time-separable Bayesian setup (e.g., Detemple (1995), Genotte (1985), Veronesi (1998)), except that the weights used by the fragile beliefs agent are not equal to the posterior probabilities associated with each state. Instead, she distorts the probabilities to reflect her uncertainty about the estimated posterior probabilities in an endogenous way, placing more weight on the models/states with lower utility. This minor departure from the classical time-separable Bayesian setup has significant impact on equilibrium prices. Indeed, this updating of beliefs will generate correlations in credit spreads that are significantly higher than if spreads were functions of the macroeconomic conditions only. Furthermore, since agents put higher weights on the states with lower utility, the model also generates significantly higher spreads (and credit risk-premia) than a traditional model based on time-separable preferences, for example.

One of our theoretical contributions is to derive sufficient conditions for which the prices of long-lived securities are equal to a weighted average of their conditional prices. We find that if these conditions do not hold, then the 'model averaging pricing rule' is not in general arbitrage-free, implying that these prices are not consistent with a no-trade equilibrium. These conditions are clearly related to the time-consistency of the preferences of agents with fragile beliefs (see Section 6.5 of Hansen and Sargent (2007)).

Another contribution of this paper is that we obtain closed-form solutions for bond prices even though, to capture contagion, the default intensity process falls outside the popular "doubly stochastic" (or Cox process) framework. Indeed, in a doubly stochastic setting, individual default events are inherently precluded from impacting the intensities of the surviving entities. In contrast, in our framework, agents update their beliefs by observing sovereign credit events. If there is a default, agents increase the probability they assign to the 'bad' hidden state. This updating raises the perceived default intensity (and in turn, credit spreads) of the other countries.

We then estimate our model using panel data on sovereign CDS spreads from February 2004 to September 2010 for 11 Euro-zone countries. We use a two-stage procedure. First, we follow the literature on sovereign risk[1] to identify a list of variables that have been shown to predict a country's ability or propensity to repay its debt. For each Euro-zone country, we use a dynamic principal component framework (e.g., Stock and Watson (1989, 1991)) to summarize the information in these variables into a single country-specific Macroeconomic Conditions Index (MCI). Since the data are observed at mixed frequencies, we rely on the filtering method of Aruoba, Diebold, and Scotti (2009). Assuming a multivariate AR(1) process, this approach provides us with both an estimate of the time series of these underlying variables and a parameter vector that captures their dynamics. We augment the set of explanatory variables to include the Chicago Board Option

[1]See, for example, Duffie et al. (2003), Edwards (1984), Hilscher and Nosbusch (2010), Longstaff et al. (2010), Min (1998), Pan and Singleton (2008).

Exchange (CBOE) VIX index as a measure of global economic uncertainty (e.g., Longstaff, Pan, Pedersen, and Singleton (2010), Pan and Singleton (2008)). Conditional default intensities are then specified to be linear functions of the state vector that includes country-specific MCIs and the VIX index.

In the second estimation step, we use sovereign CDS spreads panel data and time series of default events to identify the rest of the model parameters and the time series of the filtered posterior probability of the hidden state. We cast the model in a state-space framework and estimate it by quasi maximum likelihood in combination with the Kalman filter. Since both state and measurement equations in the system contain non-linearities, we rely on a square-root unscented filter (e.g., Wan and van der Merwe (2001), Christoffersen, Jacobs, Karoui, and Mimouni (2009)).

In the early part of the sample period, 2004-2007, we estimate the probability of the good state to be nearly one. This changes at the end of 2007, when the posterior probability of the bad economic state increases significantly and its fluctuations become more pronounced as the sovereign crisis unfolds. Consistent with Hansen and Sargent (2007, 2010), we find that the agent displays a preference for robust beliefs, in that she slants the risk-adjusted probability of the hidden state towards the model associated with the lowest continuation utility. That is, she attaches a higher probability of being in the bad state under the risk-neutral measure than the physical measure. Consequently, the level of risk-adjusted default intensities is higher than those computed under the physical probability measure, i.e., our model generates positive jump-to-default (JTD) risk-premia.

Overall, the model fits CDS spreads data well across Euro zone countries, both before and during the crisis. To better gauge its performance, we compare the pricing errors of our model to those of a (linear) affine specification with a state vector that includes country-specific MCIs and the VIX (as in our model), and a single latent factor, which we estimate with principal component analysis. In any arbitrage-free affine framework, sovereign credit spreads are a linear function of the state vector, with coefficients determined by no-arbitrage restrictions. Therefore, in sample, unrestricted ordinary least squares (OLS) regressions provide an upper bound on the goodness of fit such an affine model could achieve. We find that our model significantly outperforms the affine benchmark estimated via OLS regressions, with a 23-85% reduction in mean absolute pricing errors, and a typical drop in maximum errors by a factor of two. This strongly suggests that there are important nonlinearities in the behavior of credit spreads that elude affine specifications and are better captured by our model, which has nonlinearities in both the mapping of spreads onto the state variables and state variable dynamics.

Related Literature. Our paper builds on and combines two important strands of literature: event risk and Bayesian updating of beliefs. Conditions for which jump-to-default is not priced have been investigated by Jarrow, Lando and Yu (2005). However, recent empirical findings question this doubly-stochastic assumption. For example, Das et al. (2006, 2007) report that the observed clustering of defaults in actual data is inconsistent with this assumption. Duffie et al. (2009) use a

fragility-based model similar to ours to identify a hidden state variable consistent with a contagion-like response. Note that the focus of these papers is on estimating the empirical default probability, whereas our focus is on pricing. Jorion and Zhang (2007) find contagious effects at the industry level (see also related work by Jorion and Zhang (2009), and Lando and Nielsen (2010)).

Other papers investigating event risk include Jarrow and Yu (2001), who also provide a model where the default of one firm affects the intensity of another. However, their model remains tractable only for a "small" number N of firms exposed to contagion-risk (e.g., Jarrow and Yu (2001) investigate only $N = 2$). In contrast, our model remains tractable regardless of the number of entities that share in the contagious response.[2] Models of credit risk embedded within a macroeconomic setting include David (2008), Chen, Collin-Dufresne and Goldstein (2009), Chen (2010), and Bhamra, Kuehn and Strebulaev (2011).

Our approach shares many common features with those in the learning and contagion literature (e.g., David (1997), Detemple (1986), Feldman (1989), Veronesi (1999, 2000)).[3] As in these papers, the representative agent in our economy learns about a hidden state from observing aggregate consumption and other "diffusive" signals. However, in our model the agent also learns from observing the default history of an entities (firms or countries), i.e., information is revealed through both diffusion processes and jump processes. Further, we identify a time-consistent model of a representative agent that has fragile beliefs (Hansen and Sargent (2010), Hansen (2007)).[4] Time consistency allows us to price securities with long-dated cash flows in a tractable manner. This framework naturally generates a flight-to-quality (i.e., a drop in risk free rates) caused by an unexpected default, consistent with observation.

Our information-based mechanism for contagion is similar to that proposed by King and Wadhwani (1990) and Kodres and Pritsker (2002), who investigate contagion across international financial markets. There is also a large empirical literature that studies contagion in equity markets (e.g., Lang and Stulz (1992)) and in international finance (e.g., Bae, Karolyi and Stulz (2003)). Theocharides (2007) investigates contagion in the corporate bond market and finds empirical support for information-based transmission of crises.

The rest of the paper is as follows. In Section 2, we propose an intensity-based model of sovereign risk with a hidden state and show how beliefs are updated from observing default events and other signals. In Section 3 we investigate the pricing implications of this model by incorporating it in a general equilibrium framework where the representative agent has fragile beliefs. We then estimate the model in Section 4 using six years of sovereign CDS prices. We conclude in Section 5. In the Appendix, we identify necessary conditions for fragile beliefs preferences to generate time consistent price processes.

[2]Collin-Dufresne, Goldstein and Hugonnier (2004) simplify the bond pricing formula of Duffie, Schroeder and Skiadas (1996). Note, however, that the formula itself does not identify a tractable framework for pricing contagion risk. Other models of contagion include Davis and Lo (2001), Schönbucher and Schubert (2001), and Giesecke (2004).

[3]See also related work by Aït-Sahalia, Cacho-Diaz, Laeven (2010), Benzoni, Collin-Dufresne and Goldstein (2011).

[4]See also related work on belief-dependent utilities by, e.g., Veronesi (2004).

2 Updating Beliefs by Observing Default Processes

Consider an economy in which the true state of nature $\tilde{S}$ is unknown and can be in any one of $s \in (1, M)$ states. At date-t, investors do not know what state the economy is in, but form a prior $\pi_{s}(t) \equiv \text{Prob}(\tilde{S} = s|\mathcal{F}_{t})$, where $\mathcal{F}_{t}$ is the investors' information set at date-t. In this economy there are n defaultable entities (firms, countries) indexed by $i \in (1, n)$ with random default times τ_{i} driven by point processes characterized by default intensities. In particular, conditional upon being in state-s, the probability of default over the next interval dt is expressed via

$$\begin{aligned} \Pr\left[d\mathbf{1}_{\{\tau_{i}<t\}} = 1 \middle| \tilde{S} = s, \mathcal{F}_{t} \right] &\equiv \mathrm{E}\left[d\mathbf{1}_{\{\tau_{i}<t\}} \middle| \tilde{S} = s, \mathcal{F}_{t} \right] \\ &= \lambda_{is}(t^{-})\, \mathbf{1}_{\{\tau_{i}>t\}}\, dt. \end{aligned} \tag{1}$$

That is, we can interpret $\lambda_{is}(t^{-})$ as the date-t default intensity for country-i conditional upon being in state-s. Conditioning on both the state-s and the paths $\lambda_{is}(t^{-})|_{t=0}^{T}$ for some distant future date-T, we specify the default events across countries (or firms) to be independent. In technical terms, we are assuming a doubly-stochastic, or Cox-process *conditional* upon being in a particular state-s (e.g., Lando (1998)). We emphasize, however, because agents do not know the correct state-s, our model falls outside of the Cox-process framework, as will be made clear below.

Since investors do not know the actual state of nature, their estimate of the actual default intensity $\overline{\lambda}_{i}^{P}(t^{-})$ is defined implicitly through

$$\begin{aligned} \overline{\lambda}_{i}^{P}(t^{-})\, \mathbf{1}_{\{\tau_{i}>t\}}\, dt &\equiv \mathrm{E}\left[d\mathbf{1}_{\{\tau_{i}<t\}} \middle| \mathcal{F}_{t} \right] \\ &= \sum_{s=1}^{M} \pi_{s}(t)\, \mathrm{E}\left[d\mathbf{1}_{\{\tau_{i}<t\}} \middle| \tilde{S} = s, \mathcal{F}_{t} \right] \\ &= \sum_{s=1}^{M} \pi_{s}(t) \lambda_{is}(t^{-})\, \mathbf{1}_{\{\tau_{i}>t\}}\, dt. \end{aligned} \tag{2}$$

Thus, conditional on investors' information, the default intensity of country-i is equal to a weighted average of the conditional default intensities:

$$\overline{\lambda}_{i}^{P}(t^{-}) = \sum_{s=1}^{M} \pi_{s}(t) \lambda_{is}(t^{-}). \tag{3}$$

We assume that investors continuously update their estimates of the $\{\pi_{s}(t)\}$ process conditional upon whether or not they observe a default event during the interval dt. A direct application of Theorem 19.6 page 332 in Liptser and Shiryaev (2001) (see also their example 1, p. 333) gives the updating equation for $\pi_{s}(t)$:

$$\frac{d\pi_{s}(t)}{\pi_{s}(t^{-})} = \sum_{i=1}^{N} \left(\frac{\lambda_{is}(t^{-})}{\overline{\lambda}_{i}^{P}(t^{-})} - 1 \right) dM_{i}(t), \tag{4}$$

where we have defined the martingale process $dM_{i}(t)$ via:

$$dM_{i}(t) \equiv \left(d\mathbf{1}_{\{\tau_{i}\leq t\}} - \overline{\lambda}_{i}^{P}(t) \mathbf{1}_{\{\tau_{i}>t\}} dt \right). \tag{5}$$

This process has many intuitive properties. First, if the prior $\pi_s(t) = 1$ for some state-s (and thus $\pi_{s'}(t) = 0$ for all other s'), then there is no updating. That is, in an economy where the agents know for sure the intensity of the countries, then there is no learning to be done. Second, when no default is observed over an interval dt, then investors revise downward the 'high-default' states of nature (i.e., those s with $\lambda_{is}(t^-) > \overline{\lambda}_i^P(t^-)$), and in turn revise upward the 'low-default' states of nature (i.e., those s with $\lambda_{is}(t^-) < \overline{\lambda}_i^P(t^-)$). Conversely, when a default is observed over an interval dt, investors revise upward those high-default states of nature, and in turn revise downward those low-default states of nature. Third, note that $\pi_s(t) \equiv \mathrm{E}\left[\tilde{S} = s \middle| \omega_t\right]$ is a P-martingale in that $\mathrm{E}_t\left[d\pi_s(t)\right] = 0$, as can be seen from equations (2), (4) and (5).

2.1 Updating also from Continuous Information

In addition to observing country default processes, investors also observe continuous signals that provide information about the state. Specifically, we assume that investors observe $K+1$ signals with dynamics:

$$d\Omega_k(t) = \mu_{k,s}\, dt + \sigma_k\, d\mathcal{Z}_k(t) \qquad k \in (0, K), \tag{6}$$

where where the $\mu_{k,s}$ are constants,[5] and $\{d\mathcal{Z}_k(t)\}$ are independent Brownian Motions adapted to the full information filtration $\mathcal{G}_t$, which contains in particular the information on the 'true' state. In particular, we have $\mathrm{E}\left[d\mathcal{Z}_k(t)|\mathcal{G}_t\right] = 0$. Using the 'innovation' approach to filtering, we can rewrite the signal dynamics as:

$$d\Omega_k(t) = \bar{\mu}_k(t)\, dt + \sigma_k\, dZ_k(t), \tag{7}$$

where we have defined the drift of the signal process based on the investors' information filtration:

$$\bar{\mu}_k(t) \equiv \frac{1}{dt}\mathrm{E}\left[d\Omega_k(t)|\mathcal{F}_t\right] = \sum_s \pi_s(t)\, \mu_{k,s}. \tag{8}$$

We have also defined $dZ_k(t)$ via

$$dZ_k(t) = d\mathcal{Z}_k(t) + \left(\frac{\mu_{k,s} - \bar{\mu}_k(t)}{\sigma_k}\right) dt. \tag{9}$$

Note that $Z_k(t)$ is a Brownian motion in the information filtration of the investor since it has quadratic variation t from (9) and satisfies $\mathrm{E}\left[dZ_k(t)|\mathcal{F}_t\right] = 0$.

It is well-known (see, e.g., Liptser and Shiryaev (2001), David (1997), and Veronesi (1999)) that the updating equation for the posterior probability of the state from this continuous information is given by

$$\frac{d\pi_s(t)}{\pi_s(t)} \quad = \quad \sum_{k=0}^{K} \left(\frac{\mu_{k,s} - \bar{\mu}_k(t)}{\sigma_k}\right) dZ_k(t). \tag{10}$$

[5] These could be stochastic processes as long as they are $\mathcal{F}_t$ predictable.

Given that the agent observes both continuous signals $d\Omega_{k}(t)$ and defaults $d\mathbf{1}_{\tau_i>t}$, and that by definition the two are orthogonal signals (since one is a pure diffusion and the other a pure jump process, e.g., Protter (2001)), it follows from equations (4) and (10) that the updating equation is:

$$\frac{d\pi_{s}(t)}{\pi_{s}(t^{-})} = \sum_{i=1}^{N}\left(\frac{\lambda_{is}(t^{-})}{\overline{\lambda}_{i}^{P}(t^{-})} - 1\right) dM_{i}(t) + \sum_{k=0}^{K}\left(\frac{\mu_{k,s} - \bar{\mu}_{k}(t)}{\sigma_{k}}\right) dZ_{k}(t). \tag{11}$$

2.2 Model for Intensity Conditional on the State

We specify the default intensity of entity $i \in (1, N)$ conditional on being in state s as:

$$\lambda_{is}(t) = \alpha_{is} + \beta_{is}' X(t), \tag{12}$$

where $X(t)$ is a state vector which contains both country-specific variables as well as a common variables. We specify $X(t)$ to follow a multi-dimensional Gaussian affine process:

$$dX(t) = [\psi - \kappa X(t)]\, dt + \Sigma\, dW(t), \tag{13}$$

where without loss of generality we specify the volatility matrix Σ to be lower diagonal and we assume that $W(t)$ is a vector Brownian motion process (independent of $\mathcal{Z}(t)$) both in the general $\mathcal{G}_{t}$ as well as the investor specific filtration $\mathcal{F}_{t}$.

Up to this point, the state variable dynamics have been specified under the historical measure. In the following section we address the issue of pricing defaultable securities in the presence of contagion risk when the representative agent has fragile beliefs.

3 General Equilibrium with Fragile Beliefs

3.1 Information Structure

It is useful to define more formally the information structure of the model. All uncertainty in our model is summarized by a filtered reference probability space $(\Omega, \mathcal{G}, \{\mathcal{G}_{t}\}, P)$ where $\mathcal{G}_{t}$ is the natural filtration generated by the 'fundamental' shocks in the economy $(S, N, \mathcal{Z}, W)$. Specifically, S is a multinomial $\mathcal{G}_{0}$-measurable random variable whose realization selects the 'true' state s; N is a vector of default counting processes for each country $N_{i}(t)\,,\, i = 1, \ldots, n$, which each have $\mathcal{G}_{t}$-measurable default intensity $\lambda_{is}(t)$; $\mathcal{Z}$ and W are both (respectively $K+1$ and d-dimensional) vectors of independent Brownian motions. Investors in our economy however do not observe these 'fundamental' shocks. Instead, their filtration $\mathcal{F}_{t}$ is generated by observing the history of defaults (N), the vector of continuous signals (Ω_{k} for $k = 0, \ldots, J$), and the state vector X. They formulate a prior over the probabilities associated with the realization of the multinomial random variable. As we have shown above, in the information filtration $\mathcal{F}_{t}$, the default counting vector N has intensity $(\overline{\lambda}_{i}^{P}(t^{-})$ for $i = 1, \ldots, n\,)$. Also, W is a $\mathcal{F}_{t}$-adapted Brownian motion. However, $\mathcal{Z}(t)$ is not. Instead, the process $Z(t)$ defined in (9) above is a standard $\mathcal{F}_{t}$-adapted Brownian motion. Further, it is clear that X_{t} is both a $\mathcal{G}_{t}$ and $\mathcal{F}_{t}$ adapted Markov process and that $(\pi_{s}(t), N(t), \Omega(t)))$ have

jointly $\mathcal{F}_t$ adapted Markov dynamics. Lastly, it is clear that once investors know the realization of the state S then all the uncertainty unravels, so that $\mathcal{G}_t$ has same information as $\mathcal{F}_t \bigcup \{S\}$.

3.2 Endowment Process

We assume that the aggregate endowment (which in equilibrium will be consumed by the representative agent) has the following dynamics:[6]

$$d \log y \quad = \quad \mu_{0,s}\, dt + \sigma_0\, d\mathcal{Z}_0. \tag{14}$$

As such, the 'innovations' representation of log-endowment in the information filtration of the representative agent is:

$$d \log y = \overline{\mu}_0^P(t)\, dt + \sigma_0\, dZ_0, \tag{15}$$

where the $\mathcal{F}_t$-Brownian motion Z_0 and $\overline{\mu}_0^P(t) = \sum_s \pi_s(t)\, \mu_{0,s}$ are defined above.

3.3 Preferences

We assume that the representative agent displays fragile beliefs as described in Hansen and Sargent (2007). More specifically, we assume that the agent will value consumption streams using a two-step approach. First, he values each stream conditional on knowing the true model/state. Second, he takes an average of the model specific continuation values. Hansen and Sargent (2007) assume that at both stages the agent worries about the possible mis-specification of his model and therefore uses a robust decision rule. Their novel insight is that agents may use a different 'risk-sensitive operator' to deal with mis-specification of posterior beliefs associated with the model/state than to deal with model specification conditional on the state. Since we are interested in uncertainty about the updating process and what consequences this may have for asset prices, we assume that in this two-step procedure agents use standard time-separable log-utility function in the first step, and use a robust decision rule to 'model-average' the continuation values. More specifically, we assume:

- Conditional upon being in state $s \in (1, M)$, the agent has logarithmic-preferences. That is, the agent ranks consumption lotteries in state-s according to the (state contingent) index $V(\{C(\cdot)\}|\mathcal{F}_0, s)$, which satisfies:

$$V(\{C(\cdot)\}|\mathcal{F}_0, s) \quad = \quad \mathrm{E}\left[\int_0^\infty \beta\, dt\, e^{-\beta t}\, \log C(t) \middle| \mathcal{F}_0, s\right]. \tag{16}$$

- To rank consumption streams unconditionally, the agent displays fragile beliefs. In particular, the agent weights the conditional utility indices $V(\{C(\cdot)\}|\mathcal{F}_0, s)$ by using an entropy penalty characterized by a preference for robustness parameter ζ via:

$$V(\{C(\cdot)\}|\mathcal{F}_0) = \min_{\{\xi_s(0)\}>0} \left\{ \sum_{s=1}^{M} \pi_s(0) \left(\xi_s(0) V(\{C(\cdot)\}|\mathcal{F}_0, s) + \zeta\, \xi_s(0) \log \xi_s(0) \right) \right\}, \tag{17}$$

[6] We assume that the logarithmic aggregate endowment equals Ω_0 for notational convenience, so it is already in the agent's information set.

subject to the constraint

$$1 \;=\; \sum_s \pi_s(0)\xi_s(0). \tag{18}$$

Solving the constrained minimization, Hansen-Sargent (2010) show:

$$\xi_s(0) \;=\; \frac{e^{-(\frac{1}{\zeta})V(\{C(\cdot)\}|\mathcal{F}_0,s)}}{\sum_{s'}\pi_{s'}(0)\,e^{-(\frac{1}{\zeta})V(\{C(\cdot)\}|\mathcal{F}_0,s)}}. \tag{19}$$

Plugging this back into equation (17), preferences simplify to

$$V(\{C(\cdot)\}|\mathcal{F}_0) = -\zeta\log\left[\sum_{s=1}^{M}\pi_s(0)\;e^{-(\frac{1}{\zeta})V(\{C(\cdot)\}|\mathcal{F}_0,s)}\right]. \tag{20}$$

It is worth noting that if the agent chooses to consume the endowment stream (which will ultimately be the equilibrium in this exchange economy), we find

$$\begin{aligned} V(\{y(\cdot)\}|\mathcal{F}_0,s) &= \mathrm{E}\left[\int_0^\infty \beta\,dt\,e^{-\beta t}\;\log y(t)\middle|\,\mathcal{F}_0,s\right] \\ &= \int_0^\infty \beta\,dt\,e^{-\beta t}\;\Big[\log y(0)+\mu_{0,s}t\Big] \\ &= \log y(0)+\frac{\mu_{0,s}}{\beta}. \end{aligned} \tag{21}$$

Under this scenario, the fragility parameters

$$\xi_s(t) \;=\; \frac{e^{-(\frac{\mu_{0,s}}{\beta\zeta})}}{\sum_{s'}\pi_{s'}(t)\,e^{-(\frac{\mu_{0,s'}}{\beta\zeta})}} \tag{22}$$

are independent of the state variables $X(t)$, and hence change over time only through their dependence on the probabilities $\{\pi_{s'}(t)\}$. This feature is important for the model to be time-consistent, as we demonstrate below.

Given these preferences, and assuming complete markets, the representative agent chooses her consumption stream to maximize her utility given by equations (20) and (16) subject to the budget constraint

$$0 \;=\; \int_0^\infty dt\int d\omega_t\;A(\omega_t|\omega_0)\left[y(\omega_t)-C(\omega_t)\right], \tag{23}$$

where the $A(\omega_t|\omega_0)$ are the Arrow-Debreu prices (which the representative agent takes as exogenously specified). The agent's first order conditions with respect to consumption across all states of nature $C(\omega_t)$ imply:

$$\theta A(\omega_t|\omega_0) \;=\; \sum_s \pi_s(0)\,\xi_s(0)\beta e^{-\beta t}\pi(\omega_t|\omega_0,s)\frac{1}{C(\omega_t)}, \tag{24}$$

where the Lagrange multiplier θ can be determined by taking the limit $\omega_t \Rightarrow \omega_0$:

$$\theta \;=\; \frac{\beta}{C(\omega_0)}. \tag{25}$$

Combining these last two equations, we find that the optimal consumption bundle satisfies

$$A(\omega_t|\omega_0) \quad = \quad \sum_s \pi_s(0)\,\xi_s(0) e^{-\beta t}\pi(\omega_t|\omega_0,s)\frac{C(\omega_0)}{C(\omega_t)}. \tag{26}$$

In this representative agent endowment economy, consumption is exogenously given, and therefore this first order condition determines equilibrium state prices. As we show next, these state prices lead to a natural 'algorithm' for valuing long dated claims: First, we value claims as if the true state s were known and the representative agent has standard log-utility. Second, we average these different model-values by weighting them with endogenously distorted posterior probabilities of the state. Finally, we show that this pricing rule is arbitrage-free in that there exists a set of strictly positive state prices for which the two stage pricing rule described above holds at all times and states. Therefore, these state prices support a no-trade equilibrium in which a representative agent with fragile beliefs consumes the aggregate endowment given in (14). This result is related to the time consistency of fragile beliefs preferences, which is not guaranteed (see Hansen-Sargent (2007)). In fact, in Appendix A we provide an example of more general fragile beliefs preference framework where the state prices derived from the two-stage approach are not arbitrage-free. Instead, these preferences give rise to dynamic arbitrage opportunities. In that example, the representative agent is not time-consistent. At time zero, all claims (short- and long-dated), are valued such that she does not want to trade given her current and anticipated future consumption. However, at future dates, if markets reopen for trading at the prices consistent with time-zero state prices, the representative agent would like to trade, implying that the initial state prices do not support a no-trade equilibrium.

3.4 Arrow-Debreu Equilibrium

For markets to clear in this endowment economy, Arrow-Debreu prices adjust until the optimal consumption is equal to the exogenous endowment in each state. Thus, we find equilibrium state prices to be:

$$A(\omega_t|\omega_0) \quad = \quad \sum_s \pi_s(0)\,\xi_s(0)\,e^{-\beta t}\,\pi(\omega_t|\omega_0,s)\,e^{-[\log y(\omega_t)-\log y(\omega_0)]}. \tag{27}$$

Combining equations (14) and (27), we can express the Arrow-Debreu prices as

$$A(\omega_t|\omega_0) = \sum_s \pi_s^Q(t)\,\mathrm{E}\left[\left(\frac{\Lambda^s(t)}{\Lambda^s(0)}\right)\,\mathbf{1}_{\{\widetilde{\omega_t}=\omega_t\}}|\mathcal{F}_0,s\right]. \tag{28}$$

Here we have defined $\Lambda^s(t)$ to be the 'pricing kernel' conditional on the state being s:

$$\frac{d\Lambda^s(t)}{\Lambda^s(t)} \quad = \quad -r_s\,dt - \sigma_0\,d\mathcal{Z}_0(t), \tag{29}$$

where the state-contingent spot rates $\{r_s\}$ are constants,

$$r_s \quad = \quad \beta + \mu_{0,s} - \frac{\sigma_0^2}{2}, \tag{30}$$

and the 'distorted' model-risk-adjusted probabilities are given by:

$$\pi_s^Q(t) \equiv \pi_s(t)\,\xi_s(t). \tag{31}$$

More generally, this suggests that the date-t price $V_t^{D(\omega_T)}$ of a security with contingent cash flows $D(\omega_T)$ at date-T if state-ω_T occurs is:

$$\begin{aligned} V_t^{D(\omega_T)} &= \sum_s \pi_s^Q(t)\,\mathrm{E}\left[\left(\frac{\Lambda^s(T)}{\Lambda^s(t)}\right) D(T)|\mathcal{F}_t, s\right] \\ &= \sum_s \pi_s^Q(t)\,e^{-r_s(T-t)}\,\mathrm{E}^{Q_s}\left[D(T)|\mathcal{G}_t, s\right], \end{aligned} \tag{32}$$

where we have used the fact that when we condition on the realization of S, then $\mathcal{F}_t$ and $\mathcal{G}_t$ contain the same information, and we have defined the measure Q_s equivalent to P by the Radon-Nykodim derivative $\frac{dQ_s}{dP} = \frac{e^{r_s T}\Lambda^s(T)}{\Lambda^s(0)}$. By Girsanov's theorem, it follows that $\mathcal{Z}_0^{Q_s}(t)$ defined by $d\mathcal{Z}_0^{Q_s}(t) = d\mathcal{Z}_0(t) + \sigma_0\,dt$ is a $Q_s - \mathcal{G}_t$ Brownian motion, and all other Brownian motions orthogonal to $d\mathcal{Z}_0$ are unaffected by this change of measure.

We now show how this pricing rule is indeed consistent with absence of arbitrage in that there exists a well-defined pricing kernel that supports this pricing function for all states and times.

3.5 The Pricing Kernel and Market Prices of Risk

In Appendix A, we show that the $\mathcal{F}_t$ adapted process Λ_t with the following dynamics:

$$\frac{d\Lambda(t)}{\Lambda(t)} = -r(t)\,dt - \sum_{k=0}^{K}\phi_k(t)\,dZ_k(t) - \sum_i \Gamma_i(t)\,dM_i(t), \tag{33}$$

where

$$\begin{aligned} r(t) &= \sum_s \pi_s^Q(t) r^s \\ \phi_k &= \sigma_k \mathbf{1}_{\{k=0\}} - \frac{\overline{\mu}_k^Q - \overline{\mu}_k^P}{\sigma_k} \\ \Gamma_i(t) &= \frac{\overline{\lambda}_i^Q - \overline{\lambda}_i^P}{\overline{\lambda}_i^P} \\ \overline{\lambda}_i^Q &= \sum_s \pi_s^Q \lambda_{is}(t) \\ \overline{\mu}_k^Q &= \sum_s \pi_s^Q \mu_{k,s} \\ \overline{\mu}_k^P &= \sum_s \pi_s \mu_{k,s}\,, \end{aligned} \tag{34}$$

is a valid state price deflator for our economy, in the sense that for any arbitrary $\mathcal{F}_T$-measurable payoff $D(T)$ the following holds:

$$\mathrm{E}\left[\frac{\Lambda(T)}{\Lambda(t)} D(T)|\mathcal{F}_t\right] = V_t^{D(T)}\,, \tag{35}$$

where V_t^D is defined in equation (32) above.

In other words, this (strictly positive) state price density supports our conjectured pricing rule. A direct implication of this is that in an economy where all prices are determined by the pricing kernel $\Lambda(t)$, an individual with the fragile beliefs described above who consumes the aggregate consumption will not want to trade in any securities. We have thus identified a pricing system consistent with a no-trade equilibrium for this fragile beliefs agent. We show in Appendix A that this is directly related to the time consistency of the representative agent.

The fact that $\Gamma_i(t)$ differs from zero implies that sovereign jumps-to-default are priced in this economy even though default of any country does not affect aggregate consumption. This highlights a difference with time-separable frameworks such as David (1997) and Veronesi (2000). In those settings, jumps in π_s would imply jumps in the expected growth rate of consumption. Such jumps, however, would not carry a risk-premium, since with time-separable preferences changes in the pricing kernel $U'(c) \sim C(t)^\gamma$ are generated only by contemporaneous changes in consumption, and not changes in the expected growth rate. However, in the fragile beliefs framework, these jumps are priced.[7]

We now turn to the evaluation of long-dated risk-free and risky zero coupon bonds and the solution for sovereign CDS spreads.

3.6 The Risk-Free Zero-Coupon Bond

Using equation (32) and (35) the price of the zero-coupon bond that pays $D = 1$ unit of consumption in all states of nature at date-T is:

$$\begin{aligned} P(\pi(t), T-t) &= \mathrm{E}\left[\frac{\Lambda(T)}{\Lambda(t)}\right] \\ &= \sum_s \pi_s^Q(t)\, e^{-r_s(T-t)}\, \mathrm{E}^{Q_s}\left[1|\mathcal{G}_t, s\right] \\ &= \sum_s \pi_s^Q(t)\, e^{-r_s(T-t)}. \end{aligned} \tag{36}$$

3.7 Defaultable Zero-Coupon Bond

The price of the zero-coupon bond that pays $D = \mathbf{1}_{(\tau_i > T)}$ one unit of consumption at date-T if country-i does not default by that date, and zero otherwise, is:

$$\begin{aligned} B^i(\pi(t), X(t), T-t) &= \sum_s \pi_s^Q(t)\, e^{-r_s(T-t)}\, \mathrm{E}^{Q_s}\left[\mathbf{1}_{\{\tau_i > T\}}|\mathcal{G}_t, s\right] \\ &= \mathbf{1}_{\{\tau_i > t\}} \sum_{s=1}^{M} \pi_s^Q(t)\, e^{-r_s(T-t)}\, B_s^i(X(t), T-t), \end{aligned} \tag{37}$$

[7] An alternative to having expected consumption growth carry a risk premium is to rely on long run risk; see, e.g., Benzoni, Collin-Dufresne and Goldstein (2011), Drechsler and Yaron (2011), and Eraker and Shaliastovich (2008). However, this setting no longer admits a closed-form solution.

where we have defined

$$B^i_s(X(t), T-t) \equiv \mathrm{E}^{Q_s}\left[e^{-\int_t^T du\,\lambda_{i,s}(X_u)}|\mathcal{G}_t, s\right]. \tag{38}$$

The reason that the price of the risky bond can be written as a weighted sum of terms, each of which can be expressed in a "reduced-form structure" is because, conditional on being in state-s, we are in a doubly stochastic framework.

The implication of equation (38) is that $e^{-\int_0^t du\,\lambda_{i,s}(X(u))}\,B^i_s(X(t),t,T)$ is a Q_s-martingale, implying that the solution for $B^i_s(X(t),t,T)$ satisfies the PDE (in this equation, we drop the (i,s) subscripts on $B^i_s(X(t),t,T)$ and $\lambda_{i,s}$ to improve readability)

$$0 = -\lambda(X(t))\,B + B_t + \sum_j B_j\left[\psi_j - \sum_m \kappa_{jm}X_m\right] + \frac{1}{2}\sum_{j,j'} B_{j,j'}\sum_m \Sigma_{jm}\Sigma_{j'm}. \tag{39}$$

Here, we use the notation $B_t \equiv \frac{\partial}{\partial t}B$, $B_j \equiv \frac{\partial}{\partial X_j}B$, etc.

Given that $\lambda_{i,s}(X(u))$ is linear in the state vector $X(t)$ via equation (12) and that the risk-neutral dynamics are affine, it is well known that the solution to this expectation takes the form:

$$B^i_s(X(t), T-t) = e^{M_{i,s}(T-t) - N'_{i,s}(T-t)\,X(t)}, \tag{40}$$

with "initial conditions" $\left(M_{i,s}(\tau=0)=0,\ N'_{i,s}(\tau=0)=0\right)$. Collecting terms linear and independent of $X(t)$, we find that the deterministic coefficients satisfy

$$\begin{aligned} N_{i,s}(\tau) &= (\kappa')^{-1}\left[I_n - \exp(-\kappa'\tau)\right]\,\beta_{i,s} \\ M_{i,s}(\tau) &= \int_0^\tau du\,\left[-\alpha_{is} - N'_{i,s}(u)\psi + \frac{1}{2}N'_{i,s}(u)\Sigma\Sigma'N_{i,s}(u)\right]. \end{aligned} \tag{41}$$

3.8 The Risk-Neutral Survival Probability

The risk-neutral survival probabilities are defined as:

$$\begin{aligned} S^i(\pi(t), X(t), T-t) &= \sum_s \pi^Q_s(t)\,\mathrm{E}^{Q_s}\left[\mathbf{1}_{\{\tau_i>T\}}|\omega_t, s\right] \\ &= \left[\sum_{s=1}^M \pi^Q_s(t)\,B^i_s(X(t), T-t)\right]\,\mathbf{1}_{\{\tau_i>t\}}. \end{aligned} \tag{42}$$

Note that these can be obtained from equation (37) by setting the risk-free rate components to zero.

3.9 Sovereign CDS Spreads

Here we obtain an expression for the sovereign CDS from the risky and riskless bond equations (36)-(37). The present value of the payments in the fee leg of the CDS contract is:

$$\mathrm{PvFee}(c^i) = \sum_{j=1}^n B^i(\pi_t, X(t), t_j)\,c^i\,\Delta + \sum_{j=1}^n P(\pi_t, \frac{t_j+t_{j-1}}{2})\left[S^i(\pi_t, X(t), t_{j-1}) - S^i(\pi_t, X(t), t_j)\right]\,c^i\,\frac{\Delta}{2}, \tag{43}$$

where the second component is the present value of the accrued interest upon default (assumed to occur half-way between t_{j-1} and t_j for simplicity). Payments are made at pre-specified dates $t = t_0, t_1, t_2, \ldots t_n$. $\Delta = t_j - t_{j-1}$ is the time between promised coupon payment (typically one quarter).

The present value of the contingent default payment leg is:

$$\text{PvDef} = \sum_{j=1}^{n} P(\pi_t, \frac{t_j + t_{j-1}}{2}) \left[S^i(\pi_t, X(t), t_{j-1}) - S^i(\pi_t, X(t), t_j) \right] L, \tag{44}$$

where L is the expected loss given default experienced upon a sovereign default. For example, Pan and Singleton (2008) discuss the fact that market convention is to set $L = 0.75$ for sovereign risk (as opposed to $L = 0.6$ for corporate bonds). They find that their maximum likelihood estimates are not too distinct for most countries from that market convention.

The fair credit default swap spread is the number c^i that sets $\text{PvFee}(c^i) = \text{PvDef}$.

4 Model Estimation and Empirical Results

Here we estimate the model developed in the previous section using data on European sovereign CDS from 2004 to 2010. The approach we follow is to first estimate observable macro- and financial fundamental indicators that determine the willingness to pay of each individual country in the absence of a hidden state. This corresponds to the state vector X_t in our model. Then conditional on these observable indicators we estimate from the time-series and cross-section of CDS spreads the posterior probability of the hidden state (π_s) as well as the other parameters of the model (state dependent default intensity parameters and preference parameters) using quasi maximum likelihood, and treating the posterior probability as a latent variable. We then compare the performance of our model to a standard (linear) affine model with the same number of observable and latent variables as our benchmark. For simplicity we focus on the case in which there are only two states (which corresponds to one latent variable). We first describe the data. Then we explain the methodology adopted to determine a comprehensive set of observable macro and financial indicators. Next, we develop an affine benchmark case with observable and latent state variables. We go on to present the likelihood function for our model and the econometric methodology used to estimate it. Finally, we discuss results.

4.1 Data

Figure 1 shows a panel of daily five-year sovereign CDS spreads for eleven Euro zone countries: Austria, Belgium, Finland, France, Germany, Greece, Ireland, Italy, the Netherlands, Portugal, and Spain. The data are from Markit Financial Information Services and span the period from February 12, 2004, to September 30, 2010. Markit's coverage of the five-year sovereign CDS market is fairly comprehensive over the entire sample period, with daily observations typically available for each country. Notable exceptions are Finland, Ireland, and the Netherlands, for which data are

unavailable over the periods 09/19/2005-05/01/2006 (Finland); 02/12/2004-05/18/2004 (Ireland) and 02/12/2004-06/23/2006 (Netherlands). In contrast, the coverage of sovereign CDS contracts with maturities other than five years is spotty, especially in the early part of the sample period. Thus, we do not include those maturities in our analysis.

Table 1 contains summary statistics for the CDS series. A breakdown of the sample into a pre-crisis period (from 02/2004 to 12/2007) and a crisis period (from 01/2008 and on) confirms the patterns already evident in Figure 1. That is, prior to 2008 sovereign CDS spreads are low across Euro zone countries, with little time series variation (mean spreads range from 2.06 to 10.60 bps, with a 0.72-3.34 standard deviation). This all changes starting from 2008, when the spreads of Portugal, Ireland, Italy, Greece, and Spain (the so-called PIIGS countries) soar. The rest of the Euro zone trails behind them, also showing much higher and more volatile CDS spreads compared to the pre-crisis period.

To explain these fluctuations in CDS prices within our model, we first specify a state vector X that determines the default intensities for Euro zone countries, as in equation (12). The literature has identified a large set of country-specific political and macro economic indicators that relate to a country's ability or propensity to repay its sovereign debt (e.g., Duffie et al. (2003), Edwards (1984), Hilscher and Nosbusch (2010), Longstaff et al. (2010), Min (1998), Pan and Singleton (2008)). To obtain a parsimonious specification for the state vector X, for each Euro-zone country we summarize the information in these variables into a single country-specific MCI, as detailed in Section 4.2 below. We then augment our state vector X to include the Chicago Board Option Exchange VIX index as a proxy for global economic uncertainty (e.g., Longstaff, Pan, Pedersen, and Singleton (2010), Pan and Singleton (2008)).

4.2 Macroeconomic Conditions Indices for Euro Zone Countries

Following Stock and Watson (1989, 1991), we consider a dynamic factor framework to estimate an indicator of the latent macroeconomic and political state of each Euro zone country, as well as the Euro zone area. We explicitly incorporate macroeconomic variables and political risk indicators observed at various frequencies (monthly, quarterly, and yearly). To deal with these mixed frequencies, we follow the approach of Aruoba, Diebold, and Scotti (2009) (see also Harvey (1990), Section 6.3). In particular, for each country i we estimate the following model at the daily frequency using maximum likelihood in combination with the Kalman filter:

$$\begin{aligned} y_{i,t} &= \beta_i\, MCI_{i,t} + \Gamma_i\, W_{i,t} + \varepsilon_{i,t},\ \varepsilon_i \sim N(0, Q_i) \\ MCI_{i,t+1} &= \rho_i\, MCI_{i,t} + \eta_{i,t+1}, \qquad \eta_i \sim N(0, R_i)\,. \end{aligned} \tag{45}$$

The latent macroeconomic conditions indicator MCI_i is an AR(1) process with Gaussian i.i.d. errors. For identification, we normalize $R_i = 1$. The vector y_i contains country-specific variables. Similar to Aruoba, Diebold, and Scotti (2009), the vector W_i includes lagged values of y_i. The sample period goes from January 3, 2001, to September 30, 2010.

We first estimate the model on aggregate data for the whole Euro zone and we follow the existing literature on sovereign credit risk to guide the choice of variables to include in y_{EU}. We experiment with various specifications and exclude some of the variables that have an insignificant factor loading β_{EU} and do little to help identify the latent Euro zone indicator, MCI_{EU}. Similar to Aruoba, Diebold, and Scotti (2009), we also consider an extension that allows the error term η in the MCI transition equation (45) to be auto-correlated. We do not find statistical support for this specification and rule it out. We then go on to estimate the same specification using data for each Euro zone country and thus obtain the associated MCI_i indicators.

Table 2 summarizes the results. The first column lists the variables y_i^j, $j = 1, \ldots, 10$, that constitute the vector y_i. To construct these measures, we use the series described in Online Appendix along with the corresponding data sources. We treat all elements of y_i as stock variables with the exception of GDP per capita, which is a flow variable. To handle temporal aggregation of GDP per capita, we augment the state vector to include a 'cumulator' variable as in Harvey (1990), Section 6.3.3. Moreover, prior to estimation, we detrend GDP per capita.

The other columns of Table 2 show the sign and statistical significance of the loading β_i^j on the MCI_i index for the economic/political risk variable y_i^j of country i. The EU column has the results for the Euro zone area. Real economic activity is an important determinant of the fluctuations in the MCI_{EU} indicator. Both real GDP growth and GDP per capita load on MCI_{EU} with a negative and significant coefficient. This implies that *higher* MCI_{EU} values are associated with *worse* economic conditions in Europe. This is also evident from the positive and significant loading of the unemployment (UE) rate on the index. Consistent with this interpretation, the exports to GDP ratio has a negative and significant β. Inflation also loads on the MCI_{EU} indicator with a negative and significant coefficient, a feature easily explained by the properties of the sample period used for estimation. The Euro zone economy performed well prior to 2008, then poorly during the financial crisis, and then started to rebound in 2009. Inflation followed a similar pattern in our sample – higher consumer price growth coincided with better economic conditions. The ratios of Government surplus and debt to GDP have plausible loadings on MCI_{EU}, suggesting that stronger public finances go together with a better economic situation. The coefficients, however, are insignificant. Since we observe these variables at a yearly frequency, this is not entirely surprising. We retain them in the preferred specification due to their economic importance. Political stability is also associated with lower MCI_{EU} values. This is intuitive, although the coefficient is insignificant. Finally, previous studies have stressed the importance of liquidity measures (e.g, M3/GDP and Reserves/GDP) to explain sovereign credit ratings (e.g., Jaramillo (2010)). We include these variables in the model, but find them to be insignificant. Also, the results show that these measures of financial intermediation are positively related to the index. This could be specific to our sample period, e.g., central banks were pouring liquidity in the economy at the peak of the crisis.

The remaining columns of Table 2 contain results for individual countries. For the most part, the evidence is similar to our findings for the Euro zone. A notable exception concerns variables

that reflect the state of a country's public finances. For instance, the ratios of Government surplus and debt to GDP have a significant loading on the MCIs of Greece and Ireland.

Figure 2 plots the MCIs for the Euro zone and its individual countries. As mentioned previously, higher values are associated with a deterioration in economic conditions. It is evident from the plots that the indices are persistent. We find the AR(1) coefficient ρ to range from a low of 0.9931, with a 0.0035 standard error, for Portugal, to a high of 0.9994 with a 0.0005 standard error, for Ireland (with daily scaling). Consequently, the unconditional standard deviation of the processes is also high, as seen in the wide fluctuations of the MCI series.

The indices share a similar pattern through the first part of the financial crisis, when they increase rapidly. In 2009 they start to improve across the Euro zone, largely because of higher growth. Ireland's recovery is much slower. Greece stands out as the only country that shows no sign of recovery as of September 2010.

4.3 Linear Regression Benchmarks

Our model explains sovereign spreads with a combination of macroeconomic and financial variables, summarized in the state vector X, and a latent variable π that measures the posterior probability of the hidden state of the economy. In this setting, defaultable bond prices are the sum of nonlinear functions of X, weighted by the probability π (equation (37)). Thus, sovereign spreads are *non-linear* transformations of the state variables.

A natural competitor for our model is a *linear* (affine) specification that includes the same state vector X augmented with a single latent variable that captures a common component in credit spreads, possibly related to the notion of contagion. In any linear model of defaultable bonds, the spread between yields on risky and riskfree bonds is a linear function of the state vector. Here we explore the performance of such a model via linear OLS regressions that relate the CDS spreads for each Euro-zone country to the vector of state variables X. In the simplest case, the only explanatory variable is the county-specific MCI. We then augment these regressions to include aggregate political and macroeconomic information for the Euro zone region as well as measures of global economic uncertainty. Finally, we include a proxy for a latent variable that helps explain the covariation in CDS spreads across Euro zone countries that is not spanned by macroeconomic and financial variables.

The most general version of these regressions is given by:

$$\begin{aligned} \text{CDS}_i &= b_{0,i} + b_{1,i}\,\text{MCI}_i + b_{2,i}\,\text{MCI}_{\text{EU}_i^{\perp}} + b_{3,i}\log\,\text{VIX} + b_{4,i}\,\text{BB-BBB spread} + b_{5,i}\,\text{PC}^1_{\text{CDS}_{-i}} + \varepsilon_i \quad (46) \\ \Delta\text{CDS}_i &= b_{1,i}\,\Delta\text{MCI}_i + b_{2,i}\,\Delta\text{MCI}_{\text{EU}_i^{\perp}} + b_{3,i}\,\Delta\log\,\text{VIX} + b_{4,i}\,\Delta\text{BB-BBB spread} + b_{5,i}\,\Delta\text{PC}^1_{\text{CDS}_{-i}} + \varepsilon_i\,. \quad (47) \end{aligned}$$

Equation (46) relates the levels of CDS spreads to the levels of the explanatory variables, consistent with our bond pricing model. A potential concern with this specification has to do with the persistence of the variables included in the regressions. For instance, Granger and Newbold (1974)

warn that autocorrelation in the regression residuals could produce spuriously high R^2 statistics. Further, it is well known that ignoring the autocorrelation in the error term makes significance tests on the coefficients invalid. As a partial remedy to these concerns, we still rely on consistent OLS estimates of the model, but compute coefficients' t-ratios using standard errors that are robust to autocorrelation and heteroskedasticity. Moreover, in equation (47) we also estimate the same specification in differences. This alternative approach is common in the empirical credit risk literature, but is also subject to criticism. For instance, Maeshiro and Vali (1988) document a loss of estimation efficiency caused by the adoption of a differenced model when the disturbances of the original (levels) linear regression are autocorrelated. Doshi, Ericsson, Jacobs, and Turnbull (2011) discuss these issues in more detail and conclude that a model expressed in levels might be preferred when fitting daily CDS data, as we do in our application.

When estimating regressions (46)-(47), our measure for the dependent CDS_i variable is the daily 5-year credit default swap basis-point spread for country i. MCI_i is the macroeconomic conditions index for country i, while $\text{MCI}_{\text{EU}_i^{\perp}}$ is the residual of an OLS regression of the daily Euro zone MCI_{EU} against the daily country-i MCI_i and a constant. The log VIX variable is the logarithmic daily S&P 500 percentage implied volatility index published by the CBOE. The BB-BBB spread variable is the difference between Bank of America Merrill Lynch corporate bond effective percentage yield indices, sampled at the daily frequency. To compute the $\text{PC}^1_{\text{CDS}_{-i}}$ variable, for each Euro zone country $j \neq i$ we first regress CDS_j against the MCI_j and the logarithmic VIX. $\text{PC}^1_{\text{CDS}_{-i}}$ is the first principal component extracted from the panel of the regression residuals (we exclude Finland, Ireland, and the Netherlands due to data availability). This approach allows us to extract common information in CDS spreads across Euro zone countries that is not captured by global financial uncertainty and local macroeconomic conditions. When estimating the regressions in differences, we measure ΔCDS_i with the series of daily overlapping changes in the 5-year CDS spreads for country i relative to the prior month, $\Delta\text{CDS}_i(t) = \text{CDS}_i(t) - \text{CDS}_i(t-21)$. The sample period goes from 02/12/2004 to 09/30/2010, with the exception of Finland, Ireland, and the Netherlands, for which data are unavailable over the periods 09/19/2005-05/01/2006 (Finland); 02/12/2004-05/18/2004 (Ireland) and 02/12/2004-06/23/2006 (Netherlands).

We report OLS estimates for France, Germany, Greece, and Italy in Table 3 (regressions in levels include a constant, which we omit from the table). Here, we focus on these four countries as they provide a diverse representation of the Euro zone; a complete set of results is in the Online Appendix. In each panel, column (1.1) has the results for the simplest model, with the country-specific MCI being the only independent variable. The associated coefficient is positive and statistically significant, consistent with the interpretation that a deterioration in the country's economic conditions leads to higher CDS spreads (coefficient t-ratios in brackets are based on Newey-West heteroskedasticity and autocorrelation robust standard errors). The explanatory power of the regression is typically high, with R^2 coefficients ranging from 11% for Italy to 72% for Greece.[8]

[8] A notable exception is the coefficient for Portugal's MCI (Panel J of Table 2 in the Online Appendix); while

We then add $\text{MCI}_{\text{EU}_i^{\perp}}$ in column (1.2). The associated regression coefficient is typically positive and significant, but the improvement in fit is generally small. This suggests that economic conditions in the rest of the Euro zone have some impact, albeit moderate, on local CDS spreads. A notable exceptions is France, for which the coefficient is negative and insignificant.

Consistent with Longstaff et al. (2011), global economic uncertainty, measured by the logarithmic VIX, helps explain the variation in the spreads (column 1.3). This is evident in the error diagnostics in the two bottom rows of each panel, where we report mean and max of the absolute residuals. Compared to the results in column (1.1), the regression with the log-VIX alone produces similar mean absolute errors. For strong Euro zone economies like Germany and France, the mean absolute error in column (1.3) is lower than the one for model (1.1), suggesting that global economic uncertainty plays a significant role in explaining credit risk for these countries. The reverse applies to troubled countries like Greece, for which country-specific macroeconomic conditions do a better job at explaining CDS data. When we combine the log-VIX with the country-specific MCI in column (1.4), both variables typically have positive and significant coefficients. Nonetheless, collinearity in these two measures can cloud the interpretation of this conclusion. Augmenting the regression to include another measure of global economic uncertainty, the BB-BBB spread, leads to a further moderate increase in fit (see, e.g., the lower mean absolute error in column 1.5). The results for Greece in columns (1.4)-(1.5) stand out; local economic conditions drive most of the variation in its CDS spreads; proxies for global financial uncertainty are insignificant.

The regressions in column (1.6) contain only three explanatory variables: the country-specific MCI, the logarithmic VIX, and our proxy for a latent factor, $\text{PC}^1_{\text{CDS}_{-i}}$. This model gives the best overall fit. In particular, the mean and maximum absolute errors in the two bottom rows of each panel are much lower in model (1.6) compared to other specifications. This evidence suggests that regional information latent in Euro-zone spreads helps explain country-specific fluctuations in sovereign credit risk.

Columns (2.1)-(2.6) report estimates for the same regressions, estimated in changes (equation (47)). Consistent with Longstaff et al. (2011), in this case the explanatory power of changes in local and regional macroeconomic conditions is very limited. Changes in global financial indicators improve the fit slightly. Also in this case, regressions that include our proxy for a latent factor, $\Delta\text{PC}^1_{\text{CDS}_{-i}}$, provide the best fit.

In sum, this analysis serves two purposes. First, it illustrates the ability of a parsimonious set of variables, country-specific MCIs in combination with a global financial indicator such as the VIX index, to explain sovereign credit risk in the Euro zone. The explanatory power of macroeconomic and financial indicators is more evident in regression for spread *levels*, consistent with recent work by Doshi, Ericsson, Jacobs, and Turnbull (2011) that focuses on CDS contracts on corporate bonds. Second, the results provide a benchmark for the performance of our CDS pricing model. In an arbitrage-free (linear) affine framework, sovereign credit spreads would be a linear function of the

positive, it is insignificant, and the regression's R^2 is nearly zero.

state vector, where the loadings of credit spreads on the state variables would be determined by no-arbitrage restrictions. Therefore, in sample, unrestricted OLS regressions give an upper bound on the goodness of fit such an affine model could achieve (given that it uses the same state variable). We construct this benchmark to also include information that is common across Euro zone countries and independent of macroeconomic conditions and financial uncertainty; by construction our proxy $\text{PC}^1_{\text{CDS}_{-i}}$ should capture the most important latent factor. Therefore, if our model performs better than this affine benchmark, then this strongly suggests that there are important non-linearities in the behavior of credit spreads better captured by our model. In Section 4.5 below, we compare the errors produced by our non-linear CDS pricing models to the OLS residuals of the preferred regression in column (1.6).

4.4 Model Specification and Estimation

Based on the results in the previous section, we specify our state vector X to include the N=11 MCIs as well as the logarithmic VIX index, $X = [\, MCI_1, \ldots, MCI_{11}, \log VIX \,]'$. We assume that $M = 2$ and the hidden economic state s can be either 'good' or 'bad.' That is, the agent views the probability of a default event to be higher in the bad economic state than in the good one, $\lambda_{iG} < \lambda_{iB}$. We assume that the default intensity λ_{iB} for country i, $i = 1, \ldots, N$, depends only on the country-specific macroeconomic index as well as global uncertainty. That is, we restrict the elements of the β_{iB} vector in equation (12) to be zero, except for the two coefficients that load on MCI_i and the logarithmic VIX index. Further, we fix some of the β_{iG} coefficients to zero in order to guarantee that intensities remain positive.

The latent variable π, which measures the probability that the economy is in the good state, completes the set of model state variables. Thus, we define an augmented state vector $\overline{X} = [X', \pi]'$, with π dynamics given in equation (11). In particular, we assume that the agent updates her posterior using information on default events $d\mathbf{1}_{\{\tau_i \leq t\}}$ of the Euro zone countries, $i = 1, \ldots, 11$, as well as a single continuous signal $d\Omega_k$, $k = 1$, with dynamics given in equation (7).

Model estimation proceeds in two stages. In the *first step*, we specify and estimate the dynamics of the X vector. We map the discrete-time AR(1) process for the MCIs, given in equation (45), into the continuous-time dynamics in equation (13). We assume that the MCIs are independent across countries, with mean-reversion coefficients $\kappa_{i,i} = -\log(\rho_i)$, zero long-run means, i.e., $\psi_i = 0$, and diffusion terms normalized to unity, $\Sigma_{i,i} = 1$. The independence assumption across the MCIs greatly simplifies estimation, as we can use the ρ_i coefficients computed separately for each MCI process in Section 4.2.

Prior to estimation, we demean the logarithmic VIX series and therefore restrict $\psi_{VIX} = 0$. Moreover, we allow the VIX process to depend on lagged realization of the MCIs. We explore this linkage in linear regressions that have the daily log-VIX on the left-hand side and include lagged log-VIX and MCI realizations among the explanatory variables. We find a strong and persistent auto-regressive component in the VIX. Interestingly, including the MCIs of troubled Euro zone

economies (the PIIGS countries) increases the explanatory power of the regressions. This suggests that the turmoil in Europe has had an impact on global economic uncertainty. Interpreting the point estimates and statistical significance of these coefficients is however difficult, as there is collinearity among the MCIs. In contrast, we do not find the MCIs of Euro zone economies outside of the PIIGS circle to improve the explanatory power of the regressions, and therefore we exclude them from our specification. In this framework, the lagged log-VIX remains significant but the persistence of the autoregressive component declines.

We also explore the possibility that shocks to the VIX correlate with shocks to the MCIs. For each MCI_i, we correlate the residuals of the AR(1) process in equation (45) with the residuals of our linear regression model for the log-VIX. We find such correlations to be insignificant and fix them at zero in the rest of the analysis. With this additional restriction, the log-VIX equation in the X dynamics (13) takes the form

$$d\log VIX(t) = -\left[\kappa_{VIX}\log VIX(t) + \sum_{i\in PIIGS}\kappa_{VIX,i}MCI_i(t)\right]dt + \Sigma_{VIX}dW_{VIX}(t)\,. \tag{48}$$

We fix κ_{VIX}, $\kappa_{VIX,i}$, $i \in PIIGS$, and Σ_{VIX} at the OLS estimates of equation (48). Finally, we fix $\mu_G = 0.018$, $\mu_B = 0.005$, and $\sigma_0 = 0.03$ in the endowment dynamics (14). These parameters also determine the value of the state-contingent spot rate (equation (30)), which has to be constant to ensure that preferences are time-consistent, as we discuss in the Appendix. The robustness parameter ζ and the discount rate β enter equation (22) as a product; thus, for identification we set $\beta = 0.01$.

The *second estimation step* relies on the panel of sovereign CDS spreads and default events to identify the remaining model parameters and the time series of the filtered posterior probability of the hidden economic state. We cast the model in a state-space framework. The set of measurement equations includes the sovereign CDS spreads for each Euro zone country, which we assume to be determined by the model up to a Gaussian *i.i.d.* pricing error,

$$\widetilde{CDS}_{i,t} = CDS_i(\overline{X}_t) + \upsilon_{i,t},\quad \upsilon_{i,t}\sim N(0,V_i),\quad i=1,\ldots,11\,. \tag{49}$$

Moreover, we assume that the MCIs and the log-VIX are measured without error. Equations (11) and (13) define the state dynamics for $\overline{X} = [X',\pi]'$ and complete the state-space representation.

There are two sources of non-linearity in the system. First, the CDS pricing formula in the measurement equation (49); see equation (37) and the discussion in Section (3.9). Second, the probability of the hidden state π has non-linear dynamics (11). To accommodate these features, we discretize the model and estimate it by quasi maximum likelihood in combination with the square-root unscented Kalman filter (UKF); see, e.g., Wan and van der Merwe (2001) and Christoffersen et al. (2009).

We denote by Y the vector of CDS spreads, MCI indices and log-VIX. In addition to the elements of Y, we also observe sovereign default events. Thus, the transition density of the observable

variables is given by

$$P(\,Y_t, N_t \,|\, Y_{t-1}, N_{t-1}; \Theta) = P(\,N_t \,|\, Y_{t-1}, N_{t-1}; \Theta) \times P(\,Y_t |\, N_t, Y_{t-1}, N_{t-1}; \Theta)\,, \tag{50}$$

where Θ is the vector of model coefficients. The first term on the right-hand side of equation (50) is the default probability for the Euro zone countries. In our sample, we do not observe defaults, i.e., for all dates t and countries $i = 1, \ldots, 11$, $d\mathbf{1}_{\{\tau_i < t\}} = 0$, and $P(N_t|Y_{t-1}; \Theta) = \prod_{i=1}^{11} \sum_{s=1}^{M} \pi_s(t) \exp(-\lambda_{is}(t^-)\, dt)$, where $\lambda_{is}(t^-)$ is given in equation (12). Conditional on time t default information and time $t-1$ values of Y, we approximate the distribution of $Y(t)$ on the right-hand side of equation (50) with a multivariate Gaussian density, and compute its conditional mean and variance using the UKF. We then maximize the likelihood function,

$$P(\,Y_t, N_t, t = 0, \ldots, T \,|\, \Theta\,) = P(\,Y_0, N_0 \,|\, \Theta\,) \times \prod_{t=1}^{T} P(\,Y_t, N_t \,|\, Y_{t-1}, N_{t-1}; \Theta\,)\,. \tag{51}$$

4.5 Empirical Results

Time series of estimated probability of the hidden state. Figure 3 shows default intensities estimated over the period from February 12, 2004 to September 30, 2010. During the pre-crisis period, sovereign defaults of member countries are very unlikely, as evident from the low default intensity estimates. Consistent with this interpretation, our filtered posterior probability of the hidden state, π_{Good}, is almost one during that period (Figure 4). That is, agents are nearly sure that the Euro zone economy is in the good state.

This situation starts to change at the end of 2007, when we estimate a significant decrease in the posterior probability of the good economic state. Since then, π_{Good} estimates start to fluctuate in connection with credit spreads. This suggests that much of the co-variation in sovereign spreads is due to the 'contagion' variable (in our model the posterior probability of the unobservable hidden state), rather than to observable common and country-specific macroeconomic and financial covariates (the state variable X in equation (12)). Movements in the posterior probability of the states (or credit spreads, for that matter) do not appear to be random and unexplainable, however. Indeed, on days of major turning points in π_{Good} we find ex-post a lot of contemporaneous news reports (discussed in the Online Appendix) which seem qualitatively consistent with the directional change in our estimated posterior probability of the good state. This is reassuring since our π_{Good} variable is a latent variable inverted from prices, and thus not directly estimated from such news.[9]

Fragile belief preference parameters. Consistent with Hansen and Sargent (2010), we find that the agent displays model uncertainty aversion and slants the risk-adjusted probability of the hidden state towards the model associated with the lowest continuation utility. We estimate the robustness parameter ζ at 1.79, with a standard error of 0.08. This determines a risk adjustment

[9] Of course, one would expect CDS spreads to be related to such news, and our posterior probability is clearly linked to spreads as is evident from the figure.

ξ in the posterior probability that the agent assigns to being in the bad economy (equations (22) and (31)). It is evident from Figure 4 that under the risk-neutral measure the agent attaches a higher probability of being in the bad economy. Consequently, the levels of the risk-adjusted default intensities, $\overline{\lambda}_i^Q$, are higher than those computed under the physical probability measure, $\overline{\lambda}_i$ (see Figure 3 and magnified versions of the same plots in the Online Appendix).

Estimates of default risk premia. The right-hand side axes in Figures 3 show the ratios of the two intensities, $\overline{\lambda}_i^Q/\overline{\lambda}_i$. Across countries and throughout the sample period, the ratio ranges from one to two. Prior to the crisis, the ratios have a mild downward sloping pattern and are lower than in the post-2007 period, especially for strong European economies like Germany and France. Starting from the end of 2007 they increase significantly and fluctuate between 1.5 and 2. The increase in $\overline{\lambda}_i^Q/\overline{\lambda}_i$ coincides with the drop in our estimate of the probability that the economy is in the good state (Figure 4). Since the end of 2007, the posterior π declines from nearly one to a low of approximately 0.4 in June 2010. This greater uncertainty results in a higher risk premium on sovereign bonds.

These results are in line with empirical papers (Berndt et al. (2005), Driessen (2005)) that estimate the jump to default risk-premium, as measured by the ratio of risk-neutral to historical default intensities, and find it to be of the order two to five (i.e., short term credit spreads should be two to five times higher than historical default rates with constant recovery rates). In contrast to these models, however, we provide a theoretical explanation for such a JTD risk-premium. First, the updating mechanism described above, makes default event not conditionally diversifiable. Second, because the representative agent displays fragile beliefs, any small shift in the probability of the bad states is magnified when it comes to pricing. The combination of Bayesian updating of hidden states and fragile beliefs representative agent generates a sizable and time-varying JTD risk-premium.

Decomposition of credit spreads. To better gauge the importance of model uncertainty premia in the presence of fragile beliefs, we compute the component of CDS spreads associated with uncertainty aversion. We measure it by the difference between (1) CDS prices predicted by the model in the presence of fragile beliefs and (2) CDS prices computed under the same coefficients, except for turning off uncertainty aversion ($\zeta \Rightarrow \infty$). Figure 5 shows that these measures fluctuate with a pattern similar to that of our π_{Good} estimate in Figure 4 and exhibit a great deal of co-movement across countries. Prior to the crisis, the cross-country average of these measures is approximately 19% of the average level in CDS spreads across the Euro zone. This estimate increases to 36% during the period from January 2008, consistent with the agent demanding higher compensation for model uncertainty.

Model fit of CDS spreads. Figures 6 compares model-implied CDS spreads with actual data (magnified versions of these displays are in the Online Appendix). The plots illustrate the ability of the model to fit the time series of CDS spreads across regimes. In particular, it captures the low

level and variability of pre-crisis data, and it does well at matching the wild fluctuations in spreads during the financial crisis. Moreover, a comparison across countries shows that the model can fit the cross section of Euro zone spreads.

Comparison with a (linear) affine benchmark. Table 4 lends additional support to these conclusions. The first two rows compare the mean absolute CDS pricing errors produced by our contagion model to those produced by a linear model that includes the country-specific MCIs, the log-VIX, and a single latent factor, as explained in Section 4.3. The last two rows show a similar comparison based on the maximum absolute CDS pricing errors. The improvement over the linear specification is evident across countries. For Greece, the mean absolute error is almost one third of the error associated with the linear regressions. For other countries the improvement is also significant, ranging from 23-85% and with an average improvement of 67% across the Euro zone. The maximum pricing deviation improves considerably as well, with an average 47% reduction across countries.

5 Conclusions

We investigate a general equilibrium framework that captures contagion risk in defaultable bonds. The model has two important ingredients: a hidden state of nature which impacts expected consumption growth and default probabilities, and a representative agent with fragile beliefs (Hansen and Sargent (2010)). Even though capturing contagion implies that our intensity-based model falls outside of the "doubly stochastic" framework, bond prices remain tractable, in turn facilitating empirical investigation. In particular, the model can capture large and highly correlated credit spreads even when default probabilities and correlations in macroeconomic fundamentals are low. Finally, we also identify conditions for which the marginal utility of the agent with fragile beliefs generate time-consistent state prices.

We apply the model to a panel of sovereign CDS spreads for Euro zone countries from February 12, 2004 to September 30, 2010. Default intensities depend on (1) indices that summarize country-specific macroeconomic conditions; (2) the VIX, as a measure of global financial uncertainty; and (3) a latent variable that captures contagion risk. Estimation via the unscented Kalman filter shows that agents update their posterior probability of contagion as the Euro-zone sovereign crisis unfolds. Prior to the crisis, our estimate of the probability that the economy is in the good state is nearly one. Starting from December 2007 this estimate decreases significantly, with fluctuations that match the flow of news closely. Default intensities increase dramatically over the same period with a great deal of comovement across countries, driven by the agent's assessment that the economy is more likely to be in the bad state. We find the agent to be averse to model uncertainty, as under the risk-neutral measure she slants the probability of the hidden state towards the model with higher default intensities. This model uncertainty premium pushes risk-neutral default intensities up during the crisis and accounts for a large portion of CDS spreads. The model fits CDS spreads

data well across Euro zone countries before and through the crisis, and it significantly outperforms affine specifications that include the same observable variables and a single latent factor. That is, our model captures important nonlinearities between observed yields and state variables that affine models cannot.

Our sample period ends on September 30, 2010. Since then, the European sovereign crisis has continued to unfold. A default by Greece on its sovereign debt has become progressively more likely, and concerns about the solvency of other Euro zone countries have escalated as well. Nonetheless, there seems to have been some decoupling of the fate of Greece from the rest of the Euro zone, as measured by a decline in correlations between CDS spreads for Greece and other countries. On October 27, 2011, Euro zone leaders agreed upon a new package of measure aimed at containing the crisis. The deal included a 'voluntary' restructuring of Greek debt held by the private sector with a 50% haircut on the bonds' value. Yet, ISDA ruled that such 'voluntary' renegotiation, no matter how much arm-twisting by European Governments was involved, did not constitute a default event and, therefore, would not trigger CDS protection.

These recent events highlight a limitation of our analysis, which assumes that at most two hidden states describe the underlying economy. First, the presence of additional hidden economic states could help capture scenarios in which the default of one or more member countries is mainly idiosyncratic in nature and therefore has limited contagious effect on other sovereign entities. Second, it has become apparent during the recent crisis that the CDS market could reflect technical legal risks that are quite separate from the risks driving cash-bond default risk. In principle, this uncertainty constitute an additional hidden state as well.

Identification of these states would of course require a data sample that contains sufficient information about these events. This is not the case of our sample period, during which concerns about a Greek default were predominant, as evident from our empirical analysis, but no sovereign defaults were observed. Also, macroeconomic data, which are an input for our model, become available with a lag, which makes it impossible to work with financial data available at the time of writing.

While we acknowledge these limitations, we feel that our model and its empirical exploration provide an interesting and fruitful framework for investigating contagion risk in credit markets.

A Appendix

A.1 Proof that Robust and Fragile Preferences are not Time-Consistent

Here we demonstrate that equilibrium state prices in an economy in which the representative agent has preferences for both robustness and fragility are not time-consistent. Define $A(\omega_{_T}|\omega_{_t})$ as the date-t price of an Arrow-Debreu (A/D) security that pays \$1 at date-T iff state $\omega_{_T}$ occurs. We examine the necessary conditions for the following formula to hold:

$$A(\uparrow\uparrow|0) \stackrel{?}{=} A(\uparrow\uparrow \mid \uparrow)\, A(\uparrow|0). \tag{A.1}$$

Intuitively, the left-hand side (LHS) is the date-0 cost of a buy and hold strategy that pays \$1 at date-2 if ($\omega_{_2}$ =$\uparrow\uparrow$). In contrast, the right-hand side (RHS) can be thought of as the cost of a dynamic trading strategy which purchases, at date-0, $A(\uparrow\uparrow \mid \uparrow)$ shares of the A/D security that pays \$1 at date-1 if ($\omega_{_1}$ =$\uparrow$). Then, if $\uparrow$ occurs, the strategy pays off $A(\uparrow\uparrow \mid \uparrow)$, which is just enough at date-1 in state-$\uparrow$ to purchase one share of the A/D security that pays \$1 if ($\omega_{_2}$ =$\uparrow\uparrow$) occurs. Thus, both sides of equation (A.1) imply two different methods of obtaining \$1 at date-2 iff the ($\uparrow\uparrow$) state occurs, and hence, by absence of arbitrage, the present value of these two portfolios should be the same. We show that, in an economy in which state prices are derived from the marginal utility of an agent with fragile beliefs, in general this relation does not hold. Consequently, in general the preferences are not time-consistent. Then we prove that, for our specification, the fragile beliefs economy is actually arbitrage-free.

A.1.1 The Economy

Consider an infinite-period, discrete-time economy in which the exogenously specified dividend, which in equilibrium equals consumption, follows the binomial process:

$$\log C(t+dt) = \begin{cases} \log C(t) + \sigma\sqrt{dt} & \text{if } \uparrow \text{ occurs} \\ \log C(t) - \sigma\sqrt{dt} & \text{if } \downarrow \text{ occurs}\,. \end{cases} \tag{A.2}$$

We emphasize that, even though we use the notation dt, we are investigating a discrete-time economy, and only in some instances specialize to the continuous time limit.

The agent is uncertain which of $s \in (1, S)$ states the economy is in, but she has priors $\pi(\tilde{S} = s|\omega_{_t}) \equiv \pi_{_s}(t)$. Conditional upon being in state-s, the probability of an up state is

$$\pi(\uparrow|s, \omega_{_t}) \;\; = \;\; \frac{1}{2} + \frac{\mu_{_{0,s}}}{2\sigma}\sqrt{dt}. \tag{A.3}$$

Note that the probability of an up (or down) move is time invariant in that $\pi(\uparrow|s, \omega_{_t}) = \pi(\uparrow|s)\ \forall\omega_{_t}$.

Standard Bayesian updating implies that, conditional upon observing an $\uparrow$ event at date-$(t+1)$,

the probability that the economy is in state-s updates to:

$$\begin{aligned}\pi_{s}(\uparrow) &\equiv \pi(s|\omega_{t}\cup\uparrow)\\ &= \frac{\pi(\uparrow|s)\,\pi(s|\omega_{t})}{\pi(\uparrow|\omega_{t})}\\ &= \frac{\pi(\uparrow|s)\,\pi(s|\omega_{t})}{\sum_{s'}\pi(\uparrow|s')\,\pi(s'|\omega_{t})}\\ &\equiv \frac{\pi(\uparrow|s)\,\pi_{s}(t)}{\sum_{s'}\pi(\uparrow|s')\,\pi_{s'}(t)}.\end{aligned} \tag{A.4}$$

A.1.2 Preference for Robustness

We first assume that the economy is known to be in state-s. Generalizing the conditional log-preferences of equation (16), here we specify the agent's conditional utility as having preference for robustness:

$$V(t|s) = (1-e^{-\beta dt})\log C(t) + e^{-\beta\,dt}\min_{\xi>0,\mathrm{E}[\xi]=1}\mathrm{E}_{t}\left[\xi V(t+dt|s) + \zeta_{1}\xi\log\xi\,|s\right]. \tag{A.5}$$

The Lagrangian for this constrained minimization is:

$$\begin{aligned}\mathcal{L} &= \pi(\uparrow|s)\left[\xi(\uparrow|s)V(\uparrow|s) + \zeta_{1}\xi(\uparrow|s)\log\xi(\uparrow|s)\right]\\ &\quad+\pi(\downarrow|s)\left[\xi(\downarrow|s)V(\downarrow|s) + \zeta_{1}\xi(\downarrow|s)\log\xi(\downarrow|s)\right] + \lambda\left[1-\pi(\uparrow|s)\xi(\uparrow|s) - \pi(\downarrow|s)\xi(\downarrow|s)\right].\end{aligned}$$

The first order condition gives:

$$\frac{\partial\mathcal{L}}{\partial\xi(\uparrow|s)}: 0 = \pi(\uparrow|s)\;\left[V(\uparrow|s) + \zeta_{1}\log\xi(\uparrow|s) + \zeta_{1} - \lambda\right], \tag{A.6}$$

implying that

$$\xi(\uparrow|s) = \exp\left[\left(\frac{\lambda-\zeta_{1}}{\zeta_{1}}\right) - \frac{V(\uparrow|s)}{\zeta_{1}}\right]. \tag{A.7}$$

To identify λ, we plug back into constraint that $\mathrm{E}[\xi]=1$ to find

$$\begin{aligned}\xi(\uparrow|s) &= \frac{e^{-V(\uparrow|s)/\zeta_{1}}}{\mathrm{E}[e^{-V(t+dt|s)/\zeta_{1}}]}\\ &= \frac{e^{-V(\uparrow|s)/\zeta_{1}}}{\pi(\uparrow|s)\,e^{-V(\uparrow|s)/\zeta_{1}} + \pi(\downarrow|s)\,e^{-V(\downarrow|s)/\zeta_{1}}},\end{aligned} \tag{A.8}$$

with an analogous equation for $\xi(\downarrow|s)$.

Plugging this back into the original equation yields:

$$V(t|s) = (1-e^{-\beta dt})\log C(t) - \zeta_{1}\,e^{-\beta\,dt}\,\log\left[\pi(\uparrow|s)\,e^{-V(\uparrow|s)/\zeta_{1}} + \pi(\downarrow|s)\,e^{-V(\downarrow|s)/\zeta_{1}}\right]. \tag{A.9}$$

Recursively, we also find[10]

$$V(\uparrow|s) = (1-e^{-\beta dt})\log C(\uparrow) - \zeta_{1}\,e^{-\beta\,dt}\,\log\left[\pi(\uparrow|s)\,e^{-V(\uparrow\uparrow|s)/\zeta_{1}} + \pi(\downarrow|s)\,e^{-V(\uparrow\downarrow|s)/\zeta_{1}}\right]. \tag{A.10}$$

[10] This is a slight abuse of notation: $V(\uparrow|s)$ should really be written as $V\left(\omega_{t+1}=(\omega_{t}\cup\uparrow)|s\right)$.

In what follows, two points are important. First:

$$\frac{\partial V(t|s)}{\partial V(\uparrow|s)} \quad = \quad e^{-\beta\, dt}\, \pi(\uparrow|s)\, \xi(\uparrow|s). \tag{A.11}$$

In particular, note that the RHS is independent of date-t. The interpretation of this result is that, conditional upon being in a state-s, things work as they would in a Black-Scholes binomial tree, with, e.g., $\pi(\uparrow\uparrow|s,0) = [\pi(\uparrow|s,0)]^2$ and $A(\uparrow\uparrow|s,0) = [A(\uparrow|s,0)]^2$.

Second, the solution to equation (A.9) is

$$V(t|s) \quad = \quad \log C(t) + B_{s}. \tag{A.12}$$

Indeed, plugging this proposed solution into equation (A.9), and then using equation (A.2), we find the functional form of equation (A.12) to be self-consistent (that is, the $\log C_t$ term cancels), and further identify the functional form of the constants B_s:

$$B_{s} \quad = \quad -\zeta_1 \left(\frac{e^{-\beta\, dt}}{1-e^{-\beta\, dt}}\right) \log\left(\pi(\uparrow|s)\, e^{-\sigma\sqrt{dt}} + \pi(\downarrow|s)\, e^{\sigma\sqrt{dt}}\right). \tag{A.13}$$

Using equation (A.3) and taking the continuous-time limit, we get

$$B_{s}^{(dt\to 0)} \quad = \quad \frac{\mu_{0,s}}{\beta} - \frac{\sigma^2}{2\beta\zeta_1}. \tag{A.14}$$

This can be seen as generalizing the results of the log-utility framework inherent in equation (21) to the case in which the robustness parameter ζ_1 is finite.

A.1.3 Fragile Beliefs

The agent must still decide how to weight the different states. We assume she does this by maximizing the following objective:

$$V(t) \quad = \quad \sum_{s=1}^{m} \pi_{s}(t) \left[\xi_{s}(t) V(t|s) + \zeta_2 \xi_{s}(t) \log \xi_{s}(t)\right], \tag{A.15}$$

subject to the constraint

$$0 \quad = \quad \lambda\left(1 - \sum_{s=1}^{m} \pi_{s}(t)\xi_{s}(t)\right). \tag{A.16}$$

Setting up the Lagrangian, the first order condition gives

$$\frac{\partial}{\partial \xi_{s}} : 0 \quad = \quad \pi_{s}(t)\left[V(t|s) + \zeta_2 \log \xi_{s}(t) + \zeta_2 - \lambda\right]. \tag{A.17}$$

Applying the expectation constraint to identify λ, we obtain

$$\xi_{s}(t) = \frac{e^{-V(t|s)/\zeta_2}}{\sum_{s'} \pi_{s'}(t)\, e^{-V(t|s')/\zeta_2}}. \tag{A.18}$$

Using equation (A.12), we can rewrite this as

$$\xi_{s}(t) = \frac{e^{-B_{s}/\zeta_2}}{\sum_{s'} \pi_{s'}(t)\, e^{-B_{s'}/\zeta_2}}. \tag{A.19}$$

It is worth noting that $\xi_s(t)$ depends on date-t only through $\{\pi_{s'}(t)\}$. Therefore, with slight abuse of notation, we have

$$\xi_s(\uparrow) = \frac{e^{-B_s/\zeta_2}}{\sum_{s'} \pi_{s'}(\uparrow)\, e^{-B_{s'}/\zeta_2}}\,. \tag{A.20}$$

Also, we emphasize that equation (A.4) yields

$$\begin{aligned}
\pi_s(\uparrow)\,\xi_s(\uparrow) &= \frac{\left(\frac{\pi(\uparrow|s)\;\pi_s(t)}{\pi(\uparrow|\omega_t)}\right)\; e^{-B_s/\zeta_2}}{\sum_{s'}\left(\frac{\pi(\uparrow|s')\;\pi_{s'}(t)}{\pi(\uparrow|\omega_t)}\right)\; e^{-B_{s'}/\zeta_2}} \\
&= \frac{\pi(\uparrow|s)\;\pi_s(t)\, e^{-B_s/\zeta_2}}{\sum_{s'}\pi(\uparrow|s')\;\pi_{s'}(t)\, e^{-B_{s'}/\zeta_2}} \\
&= \frac{\pi(\uparrow|s)\;\pi_s(t)\,\xi_s(t)}{\sum_{s'}\pi(\uparrow|s')\;\pi_{s'}(t)\,\xi_{s'}(t)}.
\end{aligned} \tag{A.21}$$

The fact that both the numerator and denominator are linear in $\pi_s^Q(t) \equiv \pi_s(t)\xi_s(t)$ is important in what follows and implies that fragility is well-specified.

Getting back to the issue at hand and plugging equation (A.18) into equation (A.15), we find

$$V(t) = -\zeta_2 \log\left[\sum_s \pi_s(t)\, e^{-V(t|s)/\zeta_2}\right], \tag{A.22}$$

where, as explained above,

$$V(t|S) = \left(1 - e^{-\beta dt}\right)\log C(t) - \zeta_1 e^{-\beta\, dt}\,\log\left[\pi(\uparrow|s)e^{-V(\uparrow|s)/\zeta_1} + \pi(\downarrow|s)e^{-V(\downarrow|s)/\zeta_1}\right], \tag{A.23}$$

and

$$\frac{\partial V(t)}{\partial V(t|s)} = \pi_s(t)\,\xi_s(t). \tag{A.24}$$

A.2 Arrow-Debreu Prices

To identify the Arrow-Debreu prices, we consider the agent starting at the optimal controls, and then modifying those controls by purchasing ϵ shares of the Arrow-Debreu security that pays $\$1(\omega_1 = \uparrow)$. As such, her current consumption drops by $\epsilon A(\uparrow|0)$, and her date-1 consumption in the $\uparrow$ state increases by ϵ, with all other consumption in time and event space held constant:

$$[C(t), C(\uparrow)] \Rightarrow [C(t) - \epsilon A(\uparrow|0), C(\uparrow) + \epsilon]\,. \tag{A.25}$$

Such an infinitesimal change has no effect on optimal utility:

$$\begin{aligned}
0 &= \delta V(t) \\
&= \sum_s \pi_s(t)\,\xi_s(t)\,\delta V(t|s) \qquad\qquad (A.26) \\
&= \sum_s \pi_s(t)\,\xi_s(t)\,\left[\left(1 - e^{-\beta dt}\right)\left(\frac{1}{C_t}\right)[-\epsilon A(\uparrow|0)] + e^{-\beta dt}\pi(\uparrow|s)\xi(\uparrow|s)\left(1 - e^{-\beta dt}\right)\left(\frac{1}{C_\uparrow}\right)(\epsilon)\right],
\end{aligned}$$

where equation (A.26) comes from equation (A.24). The solution to this is

$$\begin{aligned} A(\uparrow|0) &= \left(\frac{C_t}{C_\uparrow}\right) e^{-\beta\,dt} \sum_s \pi_s(t)\,\xi_s(t)\,\pi(\uparrow|s)\,\xi(\uparrow|s) \\ &= e^{-\sigma\sqrt{dt}} e^{-\beta\,dt} \sum_s \pi_s(t)\,\xi_s(t)\,\pi(\uparrow|s)\,\xi(\uparrow|s). \end{aligned} \tag{A.27}$$

It is important to note that the only t-dependence on the right-hand side is through the $\pi_s(t)\,\xi_s(t)$. In particular, if at date-1 an $\uparrow$-state occurs, the price the A/D security that pays $\$1(\omega_2 = \uparrow\uparrow)$ (with some abuse of notation) is:

$$\begin{aligned} A(\uparrow\uparrow|\uparrow) &= e^{-\sigma\sqrt{dt}} e^{-\beta\,dt} \sum_s \pi_s(\omega_t \cup \uparrow)\,\xi_s(\omega_t \cup \uparrow)\,\pi(\uparrow|s)\xi(\uparrow|s) \\ &\equiv e^{-\sigma\sqrt{dt}} e^{-\beta\,dt} \sum_s \pi_s(\uparrow)\,\xi_s(\uparrow)\,\pi(\uparrow|s)\,\xi(\uparrow|s) \\ &= e^{-\sigma\sqrt{dt}} e^{-\beta\,dt} \sum_s \left(\frac{\pi(\uparrow|s)\;\pi_s(t)\,\xi_s(t)}{\sum_{s'}\pi(\uparrow|s')\;\pi_{s'}(t)\,\xi_{s'}(t)}\right) \pi(\uparrow|s)\,\xi(\uparrow|s) \\ &= e^{-\sigma\sqrt{dt}} e^{-\beta\,dt} \frac{\sum_s \pi_s(t)\,\xi_s(t)\pi^2(\uparrow|s)\,\xi(\uparrow|s)}{\sum_{s'}\pi_{s'}(t)\,\xi_{s'}(t)\,\pi(\uparrow|s')}, \end{aligned} \tag{A.28}$$

where we have used equation (A.21) in the second-to-last line.

Finally, consider the two-period infinitesimal variation:

$$[C(t), C(\uparrow\uparrow)] \Rightarrow [C(t) - \epsilon A(\uparrow\uparrow|0), C(\uparrow\uparrow) + \epsilon]\ . \tag{A.29}$$

We find

$$A(\uparrow\uparrow|0) = e^{-2\sigma\sqrt{dt}} e^{-2\beta\,dt} \sum_s \pi_s(t)\,\xi_s(t)\,\pi^2(\uparrow|s)\xi^2(\uparrow|s). \tag{A.30}$$

A.3 Time Consistency

The LHS of equation (A.1) is equation (A.30). Combining equations (A.27) and (A.28), we find the RHS of equation (A.1) is

$$A(\uparrow\uparrow|\uparrow)A(\uparrow|0) = \left[e^{-\sigma\sqrt{dt}} e^{-\beta\,dt} \frac{\sum_s \pi_s(t)\,\xi_s(t)\pi^2(\uparrow|s)\,\xi(\uparrow|s)}{\sum_{s'}\pi_{s'}(t)\,\xi_{s'}(t)\,\pi(\uparrow|s')}\right]\left[e^{-\sigma\sqrt{dt}} e^{-\beta\,dt} \sum_s \pi_s(t)\,\xi_s(t)\,\pi(\uparrow|s)\,\xi(\uparrow|s)\right].$$

Dividing both sides by $e^{-2\sigma\sqrt{dt}} e^{-2\beta\,dt}$ and multiplying both sides by $\sum_{s'}\pi_{s'}(t)\,\xi_{s'}(t)\,\pi(\uparrow|s')$, we obtain

$$\begin{aligned} \text{“}LHS\text{”} &= \left[\sum_s \pi_s(t)\,\xi_s(t)\,\pi^2(\uparrow|s)\,\xi^2(\uparrow|s)\right]\left[\sum_s \pi_s(t)\,\xi_s(t)\,\pi(\uparrow|s)\right] \\ \text{“}RHS\text{”} &= \left[\sum_s \pi_s(t)\,\xi_s(t)\,\pi^2(\uparrow|s)\,\xi(\uparrow|s)\right]\left[\sum_s \pi_s(t)\,\xi_s(t)\,\pi(\uparrow|s)\,\xi(\uparrow|s)\right]. \end{aligned} \tag{A.31}$$

Now, since the priors $\{\pi_s(t)\}$ are completely arbitrary (except that they sum to unity), it follows that fragility combined with robustness typically generates an economy with time-inconsistent state prices. However, there are at least two cases where the LHS equals the RHS:

- There is no hidden state. That is, $\pi_{s}(t) = 1$ for one value of s and zero for all others.
- $\xi(\uparrow|s) = \xi(\downarrow|s) = 1$, which implies that preference for robustness has been 'turned off' (i.e., $\zeta_1 = \infty$).

The first case rediscovers the well-known result that recursive preferences are time-consistent. The second case implies that a combination of fragility and conditional time-separable preferences may possibly be time-consistent. Note that all we have identified here are necessary conditions. In the next section, we show that the model we investigate in the text is in fact time-consistent.

A.4 Derivation of the Pricing Kernel in our Economy (Necessary Conditions)

Here we identify the stochastic discount factor $\Lambda(\omega_{T})$ that determines the price of a generic asset $V^{D(\omega_{T})}(\omega_{t})$ with state-contingent cash flows $D(\omega_{T})$. The pricing kernel is defined via:

$$\Lambda(\omega_{t})\, V^{D(\omega_{T})}(\omega_{t}) \;=\; \int d\omega_{T} \pi(\omega_{T}|\omega_{t})\, \Lambda(\omega_{T})\, D(\omega_{T}). \tag{A.32}$$

It is convenient to notionally specify pricing kernel dynamics as:

$$\frac{d\Lambda(t)}{\Lambda(t)} \;=\; -r(t)\, dt - \sum_{k=0}^{K} \phi_{k}(t)\, dZ_{k}(t) + \sum_{i=1}^{n} \Gamma_{i}(t)\, \left[d\mathbf{1}_{(\tau_{i}<t)} - \overline{\lambda}_{i}^{P}(t)\, dt \right]. \tag{A.33}$$

We want to identify the risk-free rate $r(t)$, the market prices of Brownian motion risk $\{\phi_{k}(t)\}$, and jump risk $\Gamma_{i}(t)$. To this end, note that $\frac{\Lambda(t+dt)}{\Lambda(t)} = 1 + \frac{d\Lambda}{\Lambda}$; this implies that we can express the price of the asset with cash flows $D(\omega_{t+dt})$ paid out at time-$(t+dt)$ as:

$$\begin{aligned} V^{D(\omega_{t+dt})}(\omega_{t}) \;=\; & \int d\omega_{t+dt} \pi(\omega_{t+dt}|\omega_{t}) \\ & \times \left[1 - r(t)\, dt - \phi_{0}(t)\, dZ_{0}(t) + \sum_{i=1}^{n} \Gamma_{i}(t)\, \left[d\mathbf{1}_{(\tau_{i}<t)} - \overline{\lambda}_{i}^{P}(t)\, dt \right] \right] D(\omega_{t+dt}). \end{aligned} \tag{A.34}$$

From equations (29) and (32), the security price has the expression

$$V^{D(\omega_{t+dt})}(\omega_{t}) \;=\; \sum_{s} \pi_{s}^{Q}(t) \int d\omega_{t+dt} \pi(\omega_{t+dt}|\omega_{t}, s) \left[1 - r_{s}\, dt - \sigma_{0}\, d\mathcal{Z}_{0}(t) \right] D(\omega_{t+dt}). \tag{A.35}$$

We now use these two different expressions to identify $r(t)$, $\{\phi_{k}(t)\}$, and $\{\Gamma_{i}(t)\}$. In particular, we consider four securities:

1. Consider a risk-free security that pays $D(\omega_{t+dt}) = 1$ in all states of nature. Comparing equations (A.34) and (A.35), we obtain:

$$1 - r(t)\, dt \;=\; \sum_{s} \pi_{s}^{Q}(t) \left[1 - r_{s}\, dt \right], \tag{A.36}$$

which implies that the risk-free rate $r(t)$ satisfies

$$r(t) \;=\; \sum_{s} \pi_{s}^{Q}(t)\, r_{s}. \tag{A.37}$$

2. Consider a security that pays $D(\omega_{t+dt}) \equiv dZ_0 = d\mathcal{Z}_0 + \left(\frac{\mu_{0,s} - \overline{\mu}_0^P(t)}{\sigma_0}\right) dt$, where we have used equation (9). Comparing equations (A.34) and (A.35), we see that the price of risk on dZ_0 is

$$\begin{aligned} -\phi_0\, dt &= -\sigma_0\, dt + \sum_s \pi_s^Q(t) \left(\frac{\mu_{0,s} - \overline{\mu}_0^P(t)}{\sigma_0}\right) dt \\ &\equiv -\sigma_0\, dt + \left(\frac{\overline{\mu}^Q(t) - \overline{\mu}_0^P(t)}{\sigma_0}\right) dt, \end{aligned} \tag{A.38}$$

where we have defined $\overline{\mu}^Q(t) = \sum_s \pi_s^Q(t)\mu_{0,s}$. This expression yields:

$$\phi_0 = \sigma_0 - \left(\frac{\overline{\mu}^Q(t) - \overline{\mu}_0^P(t)}{\sigma_0}\right). \tag{A.39}$$

3. Consider a security that pays, for some $k \in (1, K)$, $D(\omega_{t+dt}) \equiv dZ_k = d\mathcal{Z}_k + \left(\frac{\mu_{k,s} - \overline{\mu}_k^P(t)}{\sigma_k}\right) dt$, where we have used equation (9). Comparing equations (A.34) and (A.35), we observe that the price of risk on dZ_k is

$$\begin{aligned} -\phi_k\, dt &= \sum_s \pi_s^Q(t) \left(\frac{\mu_{k,s} - \overline{\mu}_k^P(t)}{\sigma_k}\right) dt \\ &\equiv \left(\frac{\overline{\mu}_k^Q(t) - \overline{\mu}_0^P(t)}{\sigma_k}\right) dt, \end{aligned} \tag{A.40}$$

where we have defined $\overline{\mu}_k^Q(t) = \sum_s \pi_s^Q(t)\mu_{k,s}$. Simplifying, we find

$$\phi_k = -\left(\frac{\overline{\mu}_k^Q(t) - \overline{\mu}_k^P(t)}{\sigma_k}\right). \tag{A.41}$$

4. Finally, consider a security that pays $D(\omega_{t+dt}) = d\mathbf{1}_{\{\tau_i < t\}}$. Comparing equations (A.34) and (A.35), we obtain:

$$\overline{\lambda}_i^P(t)\,(1 + \Gamma_i(X_t))\, dt = \sum_s \pi_s^Q(t)\, \lambda_{i,s}(X_t)\, dt. \tag{A.42}$$

Defining the risk-neutral intensity via

$$\overline{\lambda}_i^Q(X_t) = \sum_s \pi_s^Q(t)\, \lambda_{i,s}(X_t), \tag{A.43}$$

we get

$$\Gamma_i(X_t) = \frac{\overline{\lambda}_i^Q(X_t) - \overline{\lambda}_i^P(X_t)}{\overline{\lambda}_i^P(X_t)}. \tag{A.44}$$

A.5 Proof that our Fragile Beliefs Economy is Arbitrage-Free

Our candidate pricing kernel has following dynamics:

$$\frac{d\Lambda(t)}{\Lambda(t)} = -r(t)\,dt - \sum_{k=0}^{K} \phi_k(t)\,dZ_k(t) - \sum_i \frac{\overline{\lambda}_i^Q - \overline{\lambda}_i^P}{\overline{\lambda}_i^P}\,dM_i(t), \tag{A.45}$$

where $dM_i(t) \equiv \big(dN_i(t) - \overline{\lambda}_i(t)\big)$. Note that the pricing kernel is defined with respect to the filtration of the agent. That is, the $\{Z_k(t)\}$ are $\mathcal{F}_t$-Brownian motions and $N_i(t)$ has $\mathcal{F}_t$-intensity $\overline{\lambda}_i(t)$). Further, the candidate market price of Brownian risk is given by:

$$\phi_k = \sigma_k \mathbf{1}_{\{k=0\}} - \frac{\overline{\mu}_k^Q - \overline{\mu}_k^P}{\sigma_k} \tag{A.46}$$

$$\overline{\mu}_k^Q = \sum_s \pi_s^Q \mu_{k,s} \tag{A.47}$$

$$\overline{\mu}_k^P = \sum_s \pi_s \mu_{k,s}. \tag{A.48}$$

Now, according to our calculations,

$$r(t) = \sum_s \pi_s^Q(t) r^s \tag{A.49}$$

$$\overline{\lambda}_i^P = \sum_s \pi_s(t) \lambda_{is}(t) \tag{A.50}$$

$$\overline{\lambda}_i^Q = \sum_s \pi_s^Q(t) \lambda_{is}(t) \tag{A.51}$$

$$\pi_s^Q(t) = \frac{\chi_s \pi_s(t)}{\sum_s \chi_s \pi_s(t)} \tag{A.52}$$

for constants $\chi_s = e^{-\frac{\mu_{0,s}}{\beta\zeta}}$ and a $\pi_s(t)$ process with dynamics

$$d\pi_s(t) = \pi_s(t) \sum_{k=0}^{K} \frac{\mu_{k,s} - \overline{\mu}_k^P(t)}{\sigma_k} dZ_k(t) + \sum_{i=1}^{N} \alpha_{is}(t^-)\,dM_i(t). \tag{A.53}$$

A.5.1 Pricing the Risk-Free Bond

We first consider the pricing of the risk-free bond. We want to show that the calculated value of a zero-coupon bond using the marginal utility of an agent with fragile beliefs is equal to the calculated value when the pricing kernel of equation (A.45) is used. In other words, we want to show that:

$$\mathrm{E}\left[\frac{\Lambda(T)}{\Lambda(t)} \,|\, \mathcal{F}_t\right] = \sum_s \pi_s^Q(t)\,\mathrm{E}\left[\frac{\Lambda^s(T)}{\Lambda^s(t)}\,\middle|\, s, \mathcal{F}_t\right]. \tag{A.54}$$

Alternatively, since:

$$\begin{aligned}\mathrm{E}\left[\frac{\Lambda(T)}{\Lambda(t)} \,|\, \mathcal{F}_t\right] &= \mathrm{E}_t^Q\left[e^{-\int_t^T r(u)\,du}\right] \\ &= \mathrm{E}_t^Q\left[e^{-\int_t^T \left(\sum_s \pi_s^Q(u) r_s\right) du}\right],\end{aligned} \tag{A.55}$$

and since

$$\mathrm{E}\left[\left.\frac{\Lambda^s(T)}{\Lambda^s(t)}\right| s, \mathcal{F}_t\right] \;=\; E^{Q_s}\left[\left. e^{-r_s(T-t)}\right| \mathcal{F}_t, s\right] \;=\; e^{-r_s(T-t)}, \tag{A.56}$$

we need to show that:

$$\mathrm{E}_t^Q[e^{-\int_t^T \sum_s \pi_s^Q(u) r_s \, du} | \mathcal{F}_t] \;=\; \sum_s \pi_s^Q(t)\, e^{-r_s(T-t)}. \tag{A.57}$$

To prove this, it is sufficient to show that $M(t)$ defined as:

$$M(t) \;=\; e^{-\int_0^t \sum_s \pi_s^Q(u) r_s \, du} \sum_s \pi_s^Q(t)\, e^{-r_s(T-t)} \tag{A.58}$$

is a $\{Q, \mathcal{F}\}$-martingale.

Applying Itô's lemma to $M(t)$ we find:

$$dM(t) \;=\; e^{-\int_0^t \sum_s \pi_s^Q(u)\, r_s \, du} \sum_s \pi_s^Q(t)\, e^{-r_s(T-t)} \left\{ -\sum_{s'} \pi_{s'}^Q \, r^{s'} \, dt + r_{\cdot} \, dt + \frac{d\pi_s^Q(t)}{\pi_s^Q(t)} \right\}. \tag{A.59}$$

Therefore, a sufficient condition for $\mathrm{E}^Q[dM(t) \,|\, \mathcal{F}_t] = 0$ is that

$$\mathrm{E}^Q\left[\left. \frac{d\pi_s^Q(t)}{\pi_s^Q(t)} \right| \mathcal{F}_t\right] \;=\; \left(\sum_{s'} \pi_{s'}^Q r^{s'} - r_{\cdot} \right) dt. \tag{A.60}$$

Since in our equilibrium we have

$$r_{\cdot} \;=\; constant + \mu_{0,s}, \tag{A.61}$$

we see that necessary and sufficient condition for $\mathrm{E}^Q[dM(t) \,|\, \mathcal{F}_t] = 0$ is

$$\mathrm{E}^Q\left[\left. \frac{d\pi_s^Q(t)}{\pi_s^Q(t)} \right| \mathcal{F}_t\right] \;=\; \left(\sum_{s'} \pi_{s'}^Q \mu_{0,s'} - \mu_{0,s} \right) dt. \tag{A.62}$$

This result also shows that we cannot freely parameterize r_s independent of the updating equation. That is, even with 'robustness' turned off and the conditional preferences modeled as time-separable, in order for the fragile beliefs utility to be time-consistent, additional restrictions on the model are required.

To determine the appropriate market price of risk for our economy, here we derive the dynamics of π^Q and see what the restriction (A.62) implies for ϕ.

A.5.2 Dynamics of π^Q

Recall $\pi_s^Q(t) = \frac{\chi_s \pi_s(t)}{\sum_s \chi_s \pi_s(t)}$, where the $\chi_s = e^{-\frac{\mu_{0,s}}{\beta\zeta}}$ are constants, as can be seen from equation (22). Also recall that

$$d\pi_{\cdot}(t) = \pi_{\cdot}(t) \sum_{k=0}^{K} \frac{\mu_{k,s} - \overline{\mu}_k^P(t)}{\sigma_k} dZ_k(t) + \sum_{i=1}^{N} \alpha_{i\cdot}(t^-)\, dM_i(t). \tag{A.63}$$

For simplicity, here we set $N = 1$ and drop the i-subscripts in the following. We then give the general result below, which follows trivially. We find

$$\frac{d\pi_s^Q(t)}{\pi_s^Q(t)} = \frac{d\pi_s^c(t)}{\pi_s(t)} - \frac{\sum_s \chi_s d\pi_s^c(t)}{\sum_s \chi_s \pi_s(t)} + \frac{(\sum_s \chi_s d\pi_s^c(t))^2}{(\sum_s \chi_s \pi_s(t))^2} - \frac{d\pi_s^c(t)}{\pi_s(t)} \frac{\sum_s \chi_s d\pi_s(t)}{\sum_s \chi_s \pi_s(t)} + \frac{1}{\pi_s^Q} \left(\frac{\chi_s \pi_s \frac{\lambda_s}{\overline{\lambda}^P}}{\sum_s \chi_s \pi_s \frac{\lambda_s}{\overline{\lambda}^P}} - 1 \right) dN_t, \tag{A.64}$$

where the superscript 'c' denotes the continuous part. Note that the jump component simplifies to $\left(\frac{\lambda_s}{\overline{\lambda}^Q} - 1\right) dN_t$.

Now we determine the dynamics of $\sum_s \chi_s \pi_s(t)$:

$$\frac{\sum_s \chi_s d\pi_s(t)}{\sum_s \chi_s \pi_s(t)} = \sum_{k=0}^{K} \frac{(\overline{\mu}_k^Q - \overline{\mu}_k^P)}{\sigma} dZ_k(t)^P + \left(\frac{\overline{\lambda}^Q}{\overline{\lambda}^P} - 1 \right) (dN_t - \overline{\lambda}^P dt). \tag{A.65}$$

Combining equations (A.64) and (A.65) and simplifying terms yields:

$$\frac{d\pi_s^Q(t)}{\pi_s^Q(t)} = \sum_{k=0}^{K} \frac{\mu_{k,s} - \overline{\mu}_k^Q}{\sigma_k} \left(dZ_k(t)^P - \frac{\overline{\mu}_k^Q - \overline{\mu}_k^P}{\sigma_k} dt \right) + \frac{\lambda_s - \overline{\lambda}^Q}{\overline{\lambda}^Q} (dN_t - \overline{\lambda}^Q dt). \tag{A.66}$$

Note that, interestingly, the jump component of π^Q is a Q-martingale. However, under Q the drift of π^Q is in general not equal to zero. In fact, since the market price of Brownian risk is ϕ_k we have the following:

$$\mathrm{E}^Q \left[\frac{d\pi_s^Q(t)}{\pi_s^Q(t)} \middle| \mathcal{F}_t \right] = -\sum_{k=0}^{K} \frac{\mu_{k,s} - \overline{\mu}_k^Q}{\sigma_k} \left(\phi_k + \frac{\overline{\mu}_k^Q - \overline{\mu}_k^P}{\sigma_k} \right) dt. \tag{A.67}$$

Recall that we want to find the market price of risk ϕ such that equation (A.62) holds, i.e.,

$$\mathrm{E}^Q \left[\frac{d\pi_s^Q(t)}{\pi_s^Q(t)} \middle| \mathcal{F}_t \right] = \left(\overline{\mu}_0^Q - \mu_{0,s} \right) dt. \tag{A.68}$$

Combining this expression with equation (A.67), we obtain the restriction:

$$\sigma_k \mathbf{1}_{\{k=0\}} = \phi_k + \frac{\overline{\mu}_k^Q - \overline{\mu}_k^P}{\sigma_k}, \quad \forall k = 0, 1, \ldots, K. \tag{A.69}$$

In turn, the Q-dynamics of π^Q are given by

$$\frac{d\pi_s^Q(t)}{\pi_s^Q(t)} = (\overline{\mu}_0^Q - \mu_{0,s})\, dt + \sum_{k=0}^{K} \frac{\mu_{k,s} - \overline{\mu}_k^Q}{\sigma_k} dZ_k^Q(t) + \frac{\lambda_s - \overline{\lambda}^Q}{\overline{\lambda}^Q} (dN_t - \overline{\lambda}^Q\, dt). \tag{A.70}$$

A.5.3 Pricing an Arbitrary Contingent Claim

The last section proved that the price of the risk-free bond is consistent with the market price dynamics implied by our pricing kernel given in equation (A.45) above. More generally, here we show that the price of *any* arbitrary contingent claim priced off of the fragile beliefs agent's marginal utility is consistent with the prices given by the economy with this pricing kernel. In

other words, the fragile beliefs economy is arbitrage-free, and there exists a no-trade equilibrium with a representative agent, endowed with the fragile beliefs preferences we have specified, who consumes the aggregate consumption.

Consider a claim to an arbitrary $\mathcal{F}_T$ measurable payoff X_T. We want to show that the value of the payoff in the economy where prices are determined by the pricing kernel of equation (A.45) agrees with the price of the claim evaluated off of the gradient of the fragile beliefs agent. In other words, we want to show that:

$$\mathrm{E}\left[\frac{\Lambda(T)}{\Lambda(t)}X_T\,\middle|\,\mathcal{F}_t\right]=\sum_s \pi_s^Q(t)\mathrm{E}\left[\frac{\Lambda^s(T)}{\Lambda^s(t)}X_T\,\middle|\,\mathcal{F}_t,s\right]. \tag{A.71}$$

Alternatively, since

$$\begin{aligned}\mathrm{E}\left[\frac{\Lambda(T)}{\Lambda(t)}X_T\,\middle|\,\mathcal{F}_t\right] &= \mathrm{E}^Q[e^{-\int_t^T r(u)du}X_T\,|\,\mathcal{F}_t]\\ &= \mathrm{E}^Q[e^{-\int_t^T \sum_s \pi_s^Q(u)r_s\,du}X_T\,|\,\mathcal{F}_t],\end{aligned} \tag{A.72}$$

and

$$\mathrm{E}\left[\frac{\Lambda^s(T)}{\Lambda^s(t)}X_T\,\middle|\,\mathcal{F}_t,s\right]=E^{Q_s}[e^{-r_s(T-t)}X_T|\mathcal{F}_t,s], \tag{A.73}$$

we need to show that

$$\mathrm{E}^Q[e^{-\int_t^T \sum_s \pi_s^Q(u)r_s\,du}X_T\,|\,\mathcal{F}_t]=\sum_s \pi_s^Q(t)E^{Q_s}[e^{-r_s(T-t)}X_T|\mathcal{F}_t,s]. \tag{A.74}$$

It suffices to verify that $M(t)$ defined as

$$M(t) = e^{-\int_0^t \sum_s \pi_s^Q(u)r_s\,du}\sum_s \pi_s^Q(t)\,E^{Q_s}[e^{-r_s(T-t)}\,X_T|\mathcal{F}_t,s] \tag{A.75}$$

is a $\{Q,\mathcal{F}\}$-martingale. Note that from the definition of the Q_s measure, using the martingale representation theorem, there exist $\{\mathcal{F}_t\}$ adapted processes $\psi_{\cdot,\cdot}(t)$ so that

$$\begin{aligned}H_t &:= \mathrm{E}^{Q_s}[X_T|\mathcal{F}_t,s]\\ &= \mathrm{E}^{Q_s}[X_T|\mathcal{F}_0,s]+\int_0^t\left\{\sum_{k=0}^K \psi_{k,s}(u)\,d\mathcal{Z}_k^{Q_s}(u)+\psi_{W,s}(u)\,dW(u)+\psi_{N,s}(dN(u)-\lambda_s dt)\right\},\end{aligned} \tag{A.76}$$

where for simplicity we assume the jump process (N) is one-dimensional (the extension to multi-dimensional $i=1,\ldots,n$ is straightforward). Applying Itô's lemma to $M(t)$ we find:

$$\begin{aligned}dM(t) &= e^{-\int_0^t \sum_s \pi_s^Q(u)r_s\,du}\sum_s \pi_s^Q(t)\,e^{-r_s(T-t)}\times\\ &\left\{-H_t\left(\sum_{s'}\pi_{s'}^Q r^{s'}\,dt+r_s\,dt+\frac{d\pi_s^Q(t)}{\pi_s^Q(t)}\right)+\sum_k \psi_{k,s}(t)\,d\mathcal{Z}_k^{Q_s}(t)+\psi_{W,s}(t)\,dW(t)\right.\\ &\left.+\sum_k \psi_{k,s}(t)\frac{\mu_{k,s}-\overline{\mu}_k^Q}{\sigma_k}\,dt+\frac{\psi_{N,s}(u)\lambda_s}{\overline{\lambda}^Q(t)}(dN(t)-\overline{\lambda}^Q(t))\right\}.\end{aligned} \tag{A.77}$$

Now, since by definition of the risk-neutral measure we have

$$dZ_k^Q(t) \quad = \quad dZ_k(t) + \left(\sigma_k \mathbf{1}_{\{k=0\}} - \frac{\overline{\mu}_k^Q - \overline{\mu}_k^P}{\sigma_k} \right) dt. \tag{A.78}$$

Further, by definition of the Q_s measure we have:

$$d\mathcal{Z}_k^{Q_s}(t) \quad = \quad d\mathcal{Z}_k + \sigma_k \mathbf{1}_{\{k=0\}} \, dt. \tag{A.79}$$

Finally, we have defined

$$dZ_k(t) \quad = \quad d\mathcal{Z}_k(t) + \frac{\mu_{k,s} - \overline{\mu}_k^P}{\sigma_k} \, dt. \tag{A.80}$$

Combining, we obtain:

$$d\mathcal{Z}_k^{Q_s}(t) + \frac{\mu_{k,s} - \overline{\mu}_k^Q}{\sigma_k} \, dt \quad = \quad dZ_k^Q(t). \tag{A.81}$$

Thus

$$\mathrm{E}_t^Q \left[\psi_{k,s}(t) d\mathcal{Z}_k^{Q_s}(t) + \psi_{k,s}(t) \frac{\mu_{k,s} - \overline{\mu}_k^Q}{\sigma_k} dt \,|\, \mathcal{F}_t \right] \quad = \quad 0. \tag{A.82}$$

Previously, we have proved that

$$\mathrm{E}_t^Q \left[\left. \frac{d\pi_s^Q(t)}{\pi_s^Q(t)} \right| \mathcal{F}_t \right] \quad = \quad (\sum_{s'} \pi_{s'}^Q r^{s'} - r_{s}) \, dt. \tag{A.83}$$

Therefore, since $W(t)$ is a $\mathcal{F}_t$-Brownian motion and $N(t)$ is has $\mathcal{F}_t$ jump intensity $\overline{\lambda}(t)$, we have indeed proved that

$$\mathrm{E}_t^Q[dM(t)|\mathcal{F}_t] = 0. \tag{A.84}$$

This implies that we have determined a valid pricing kernel for our economy.

References

[1] Y. Aït-Sahalia, J. Cacho-Diaz, and R. Laeven. Modeling financial contagion using mutually exciting jump processes. Working Paper, Princeton University, 2010.

[2] S. B. Aruoba, F. X. Diebold, and C. Scotti. Real-time measurement of business conditions. *Journal of Business and Economic Statistics*, 27:417–427, 2009.

[3] K.-H. Bae, A. Karolyi, and R. Stulz. A new approach to measuring financial contagion. *Review of Financial Studies*, 16:717–763, 2003.

[4] J. Bai, P. Collin-Dufresne, R. Goldstein, and J. Helwege. How large can jump-to-default risk be? Unpublished working paper, University of Minnesota, 2011.

[5] L. Benzoni, P. Collin-Dufresne, and R. Goldstein. Explaining asset pricing puzzles associated with the 1987 market crash. *Journal of Financial Economics*, 101:552–573, 2011.

[6] A. Berndt, R. Douglas, D. Duffie, M. Ferguson, and D. Schranz. Measuring default risk premia from default swap rates and EDFs. Working Paper, Stanford University, 2005.

[7] H. Bhamra, L.-A. Kuehn, and I. A. Strebulaev. The aggregate dynamics of capital structrure and macroeconomic risk. Forthcoming, Review of Financial Studies, 2011.

[8] H. Chen. Macroeconomic conditions and the puzzles of credit spreads and capital structure. *Journal of Finance*, 65:2171–2212, 2010.

[9] L. Chen, P. Collin-Dufresne, and R. Goldstein. On the relation between the credit spread puzzle and the equity premium puzzle. *Review of Financial Studies*, 22:3367–3409, 2009.

[10] P. F. Christoffersen, K. Jacobs, L. Karoui, and K. Mimouni. Nonlinear filtering in affine term structure models: Evidence from the term structure of swap rates. Unpublished working paper. University of Toronto and University of Houston, 2009.

[11] P. Collin-Dufresne, R. Goldstein, and J. Hugonnier. A general formula for the valuation of defaultable securities. *Econometrica*, 72:1377–1407, 2004.

[12] S. Das, D. Duffie, N. Kapadia, and L. Saita. Common failings: How corporate defaults are correlated. *Journal of Finance*, 62:93–117, 2007.

[13] S. Das, L. Freed, G. Geng, and N. Kapadia. Correlated default risk. *Journal of Fixed Income*, 16:1–26, 2006.

[14] A. David. Fluctuating confidence in stock markets: Implications for returns and volatility. *Journal of Financial and Quantitative Analysis*, 32:427–462, 1997.

[15] A. David. Inflation uncertainty, asset valuations, and the credit spreads puzzle. *Review of Financial Studies*, 21:2487–2534, 2008.

[16] A. David and P. Veronesi. Option prices with uncertain fundamentals: Theory and evidence on the dynamics of implied volatilities. Working Paper, University of Chicago, 2002.

[17] M. Davis and V. Lo. Infectious defaults. *Quantitative Finance*, 1:382–387, 2001.

[18] J. Detemple. Asset pricing in a production economy with incomplete information. *Journal of Finance*, 41:383–393, 1986.

[19] H. Doshi, J. Ericsson, K. Jacobs, and S. M. Turnbull. On pricing credit default swaps with observable covariates. Working Paper, McGill University and University of Houston, 2011.

[20] I. Drechsler and A. Yaron. What's vol got to do with it. *Review of Financial Studies*, 24:1–45, 2011.

[21] J. Driessen. Is default event risk priced in corporate bonds? *Review of Financial Studies*, 18:165–195, 2005.

[22] D. Duffie, L. Pedersen, and K. J. Singleton. Modeling sovereign yield spreads: A case study of Russian debt. *Journal of Finance*, 58:119–159, 2003.

[23] D. Duffie, M. Schroder, and C. Skiadas. Recursive valuation of defaultable claims and the timing of resolution of uncertainty. *Annals of Applied Probability*, 6:1075–1090, 1996.

[24] S. Edwards. LDC foreign borrowing and default risk: An empirical investigation, 1976-80. *American Economic Review*, 74:726–734, 1984.

[25] B. Eraker and I. Shaliastovich. An equilibrium guide to designing affine pricing models. *Mathematical Finance*, 18:519–543, 2008.

[26] D. Feldman. The term structure of interest rates in a partially observable economy. *Journal of Finance*, 44:789–812, 1989.

[27] K. Giesecke. Correlated default with incomplete information. *Journal of Banking and Finance*, 28:1521–1545, 2004.

[28] J. Gomes and L. Schmid. Equilibrium credit spreads and the macroeconomy. Working Paper, Duke University, 2010.

[29] C. W. J. Granger and P. Newbold. Spurious regressions in econometrics. *Journal of Econometrics*, 2:111–120, 1974.

[30] L. P. Hansen. Beliefs, doubts and learning: Valuing macroeconomic risk. *American Economic Review, Papers and Proceedings*, 97:1–30, 2007.

[31] A. C. Harvey. *Forecasting, Structural Time Series Models and the Kalman Filter*. Cambridge University Press, 1990.

[32] J. Hilscher and Y. Nosbusch. Determinants of sovereign risk: Macroeconomic fundamentals and the pricing of sovereign debt. *Review of Finance*, 14:235–262, 2010.

[33] L. Jaramillo. Determinants of investment grade status in emerging markets. Working Paper, International Monetary Fund, 2010.

[34] R. Jarrow, D. Lando, and F. Yu. Default risk and diversification: Theory and application. *Mathematical Finance*, 15:1–26, 2005.

[35] R. Jarrow and F. Yu. Counterparty risk and the pricing of defaultable securities. *Journal of Finance*, 56:1765–99, 2001.

[36] P. Jorion and G. Zhang. Good and bad credit contagion: Evidence from credit default swaps. *Journal of Financial Economics*, 84:860–883, 2007.

[37] P. Jorion and G. Zhang. Credit contagion and counterparty risk. *Journal of Finance*, 64:2053–2087, 2009.

[38] M. A. King and S. Wadhwani. Transmission of volatility between stock markets. *Review of Financial Studies*, 3:5–33, 1990.

[39] L. E. Kodres and M. Pritsker. A rational expections model of financial contagion. *Journal of Finance*, 57:769–799, 2002.

[40] D. Lando. On cox processes and credit risky securities. *Review of Derivatives Research*, 2:99–120, 1998.

[41] D. Lando and M. S. Nielsen. Correlation in corporate defaults: Contagion or conditional independence? *Journal of Fixed Income*, 19:355–372, 2010.

[42] L. Lang and R. Stulz. Contagion and competitive intra-industry effects of bankruptcy announcements. *Journal of Financial Economics*, 32:45–60, 1992.

[43] R. Liptser and A. Shiryaev. *Statistics of Random Processes*, volume I, II. Springer-Verlag, New York, NY, 2001.

[44] F. Longstaff, J. Pan, L. Pedersen, and K. J. Singleton. How sovereign is sovereign credit risk? *American Economic Journal: Macroeconomics*, 3:75–103, 2010.

[45] A. Maeshiro and S. Vali. Pitfalls in the estimation of a differenced model. *JBES*, 6:511–515, 1988.

[46] H. Min. Determinants of emerging market bond spread: Do economic fundamentals matter? Working Paper, World Bank, 1998.

[47] J. Pan and K. J. Singleton. Default and recovery implicit in the term structure of sovereign cds spreads. *Journal of Finance*, 63:2345–2384, 2008.

[48] P. J. Schönbucher and D. Schubert. Copula-dependent default risk in intensity models. 2001. Working Paper, Bonn University.

[49] G. Theocharides. Contagion: Evidence from the bond market. Working Paper, SKK Business School, Korea, 2007.

[50] P. Veronesi. Stock market overreactions to bad news in good times: a rational expectations equilibrium model. *Review of Financial Studies*, 12:975–1007, 1999.

[51] P. Veronesi. How does information quality affect stock returns? *Journal of Finance*, 55:807–837, 2000.

[52] P. Veronesi. Belief dependent utilities, aversion to state uncertainty, and asset prices. Working Paper, University of Chicago, 2004.

[53] E. A. Wan and R. van der Merwe. The unscented kalman filter. In S. Haykin, editor, *Kalman Filtering and Neural Networks*, chapter 7, pages 221–280. Wiley Publishing, 2001.

[54] M. W. Watson and J. H. Stock. New indexes of coincident and leading economic indicators. In O. J. Blanchard and S. Fischer, editors, *Macroeconomics Annual*, volume 4, pages 351–409. M.I.T. Press, 1989.

[55] M. W. Watson and J. H. Stock. A probability model of the coincident economic indicators. In K. Lahiri and G. Moore, editors, *Leading Economic Indicators: New Approaches and Forecasting Records*, chapter 4, pages 63–85. Cambridge University Press, 1991.

Figures and Tables

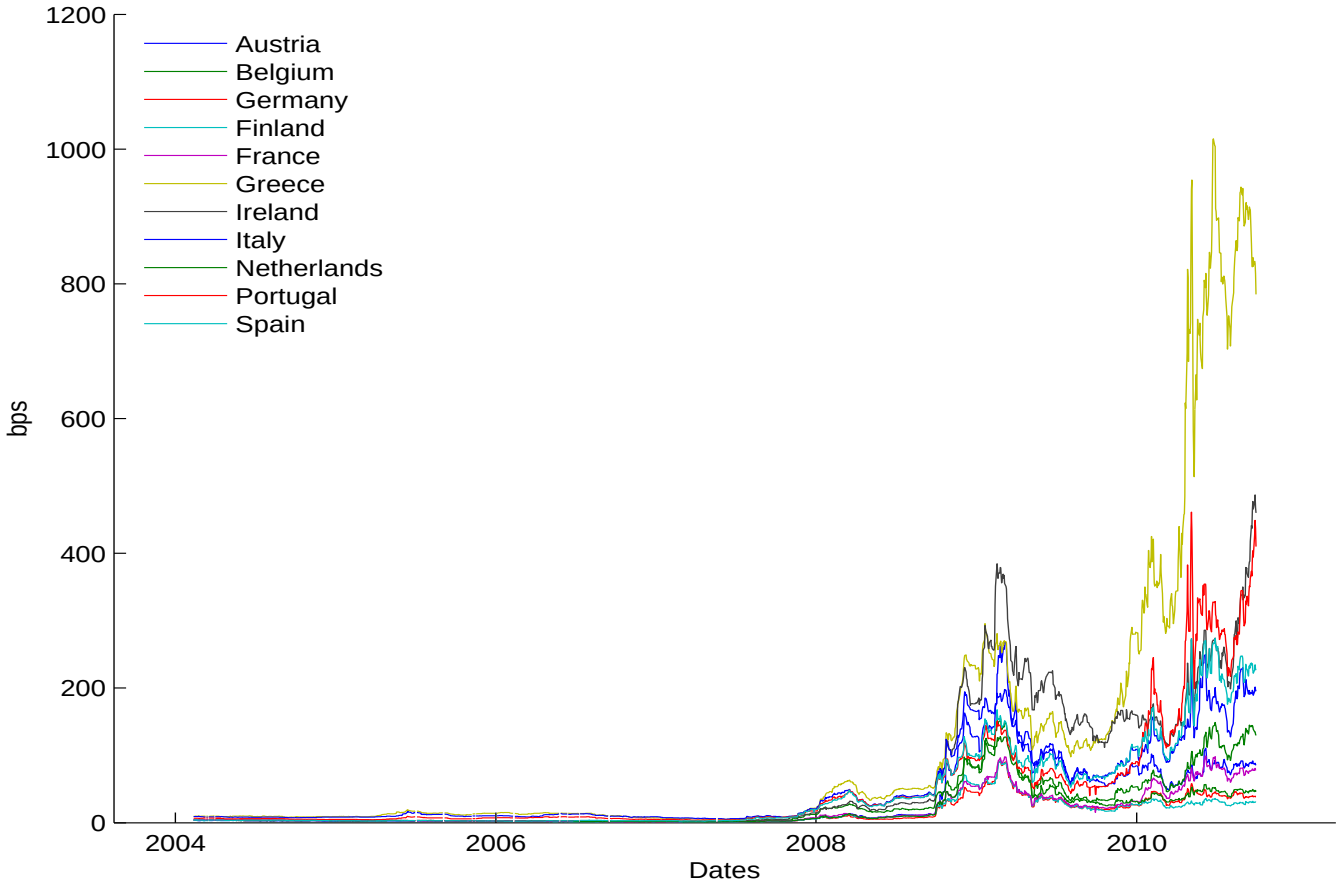

Figure 1: **Sovereign CDS Spreads.** The plots show 5-year sovereign CDS spreads for Euro zone countries. The sample period is 02/12/2004-09/30/2010. Source: Markit Financial Information Services.

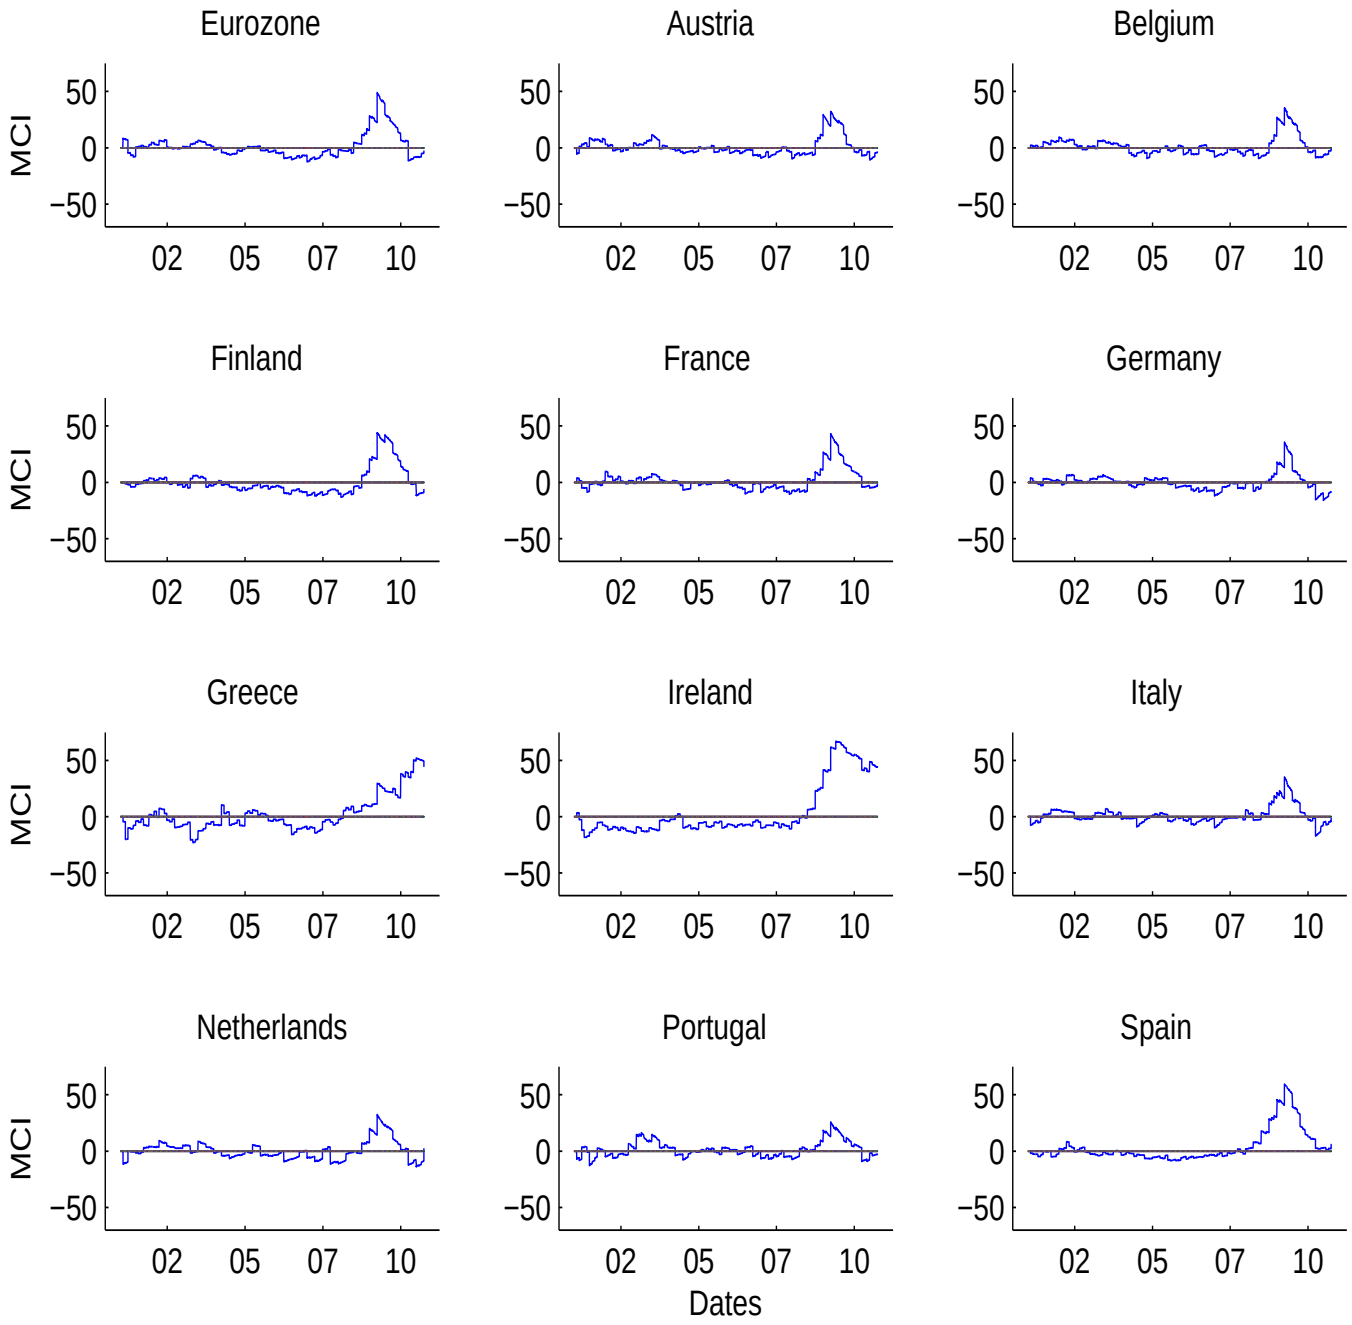

Figure 2: **Macroeconomic Conditions Indices.** The plots show the macroeconomic conditions indices for Euro zone countries. The sample period is 01/03/2001-09/30/2010.

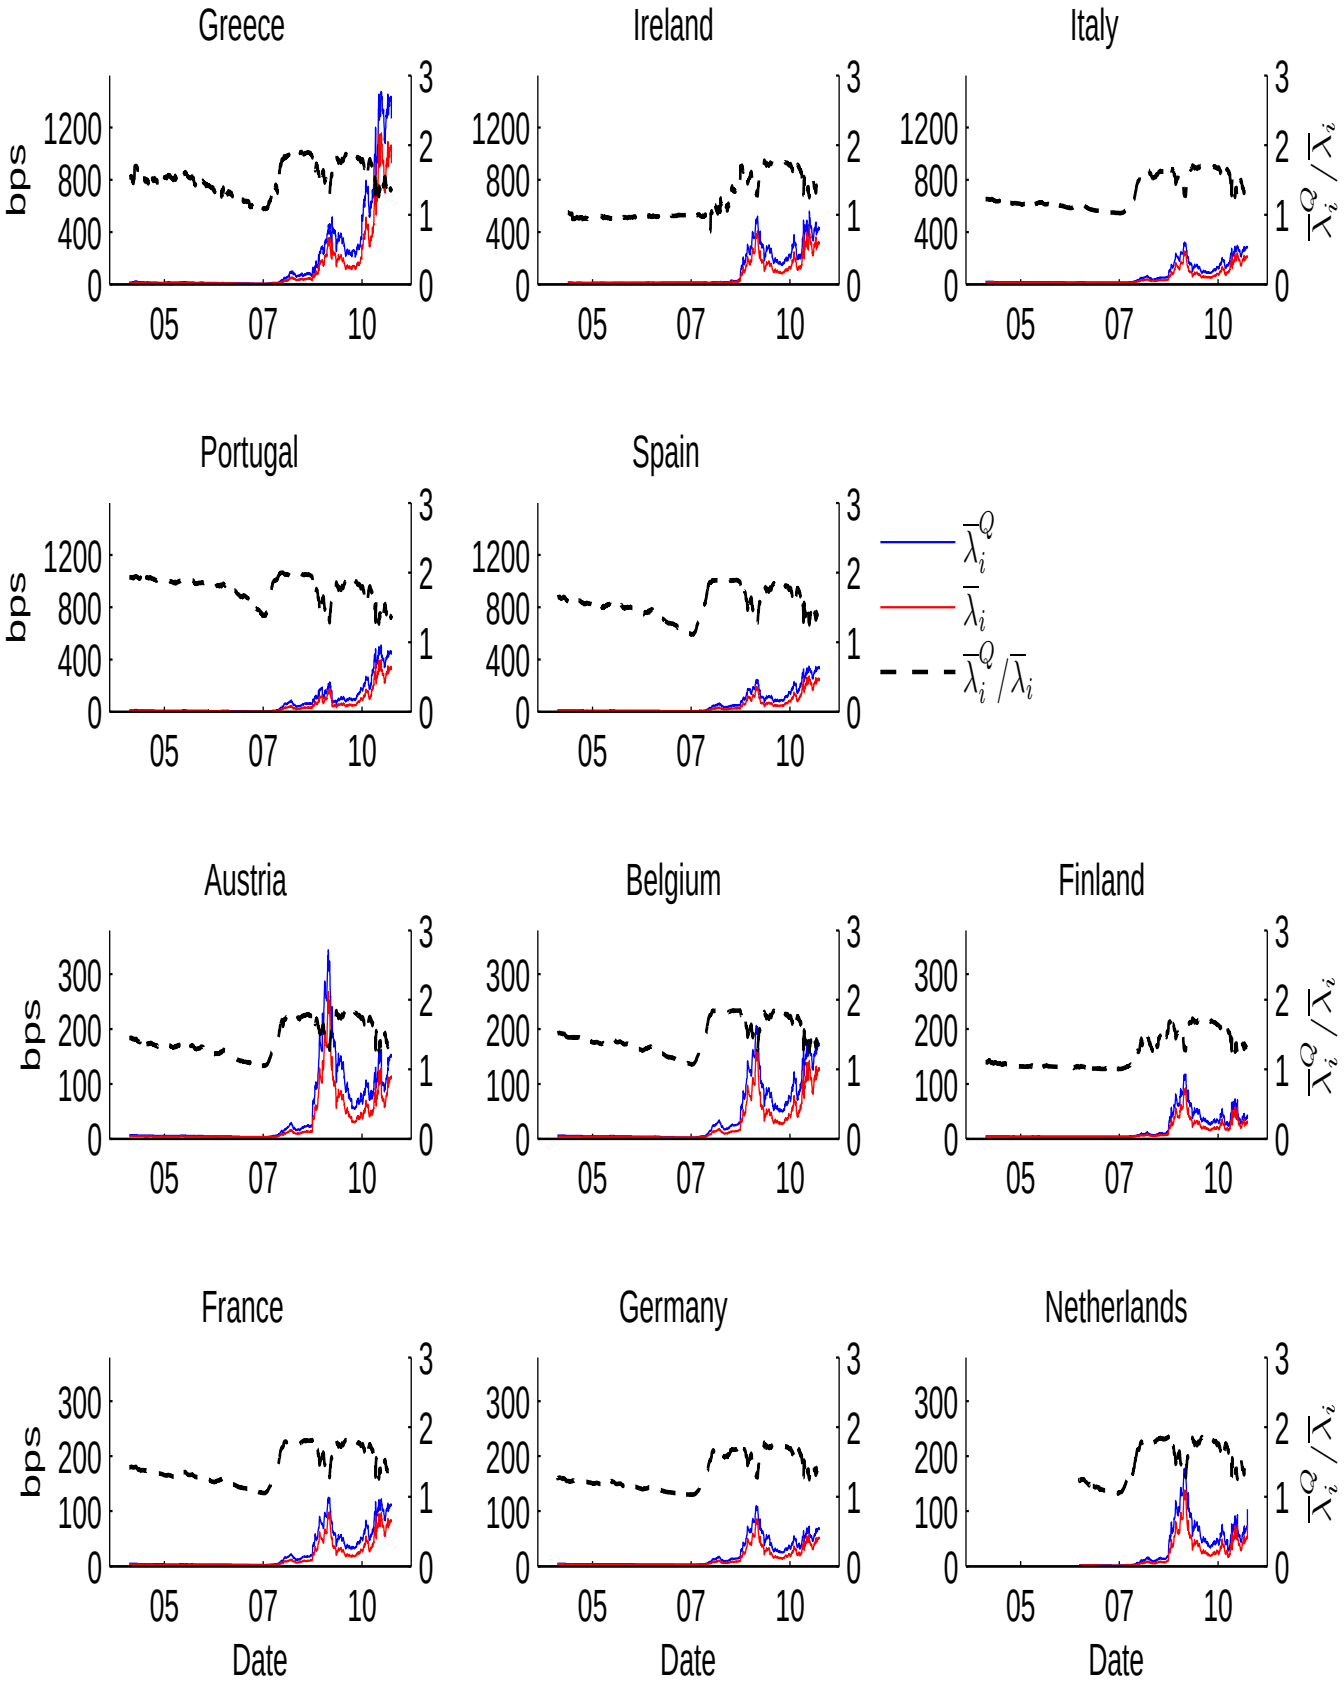

Figure 3: **Default Intensities.** The plots show the P- and Q-measure estimates for the default intensities, $\overline{\lambda}_i$ and $\overline{\lambda}_i^Q$, for Euro zone countries. The left-hand side axis shows default intensities measured in basis points. The right-hand side axis shows the $\overline{\lambda}_i^Q / \overline{\lambda}_i$ ratio. We estimate the model via quasi maximum likelihood in combination with the unscented Kalman filter. The sample period is 02/12/2004-09/30/2010.

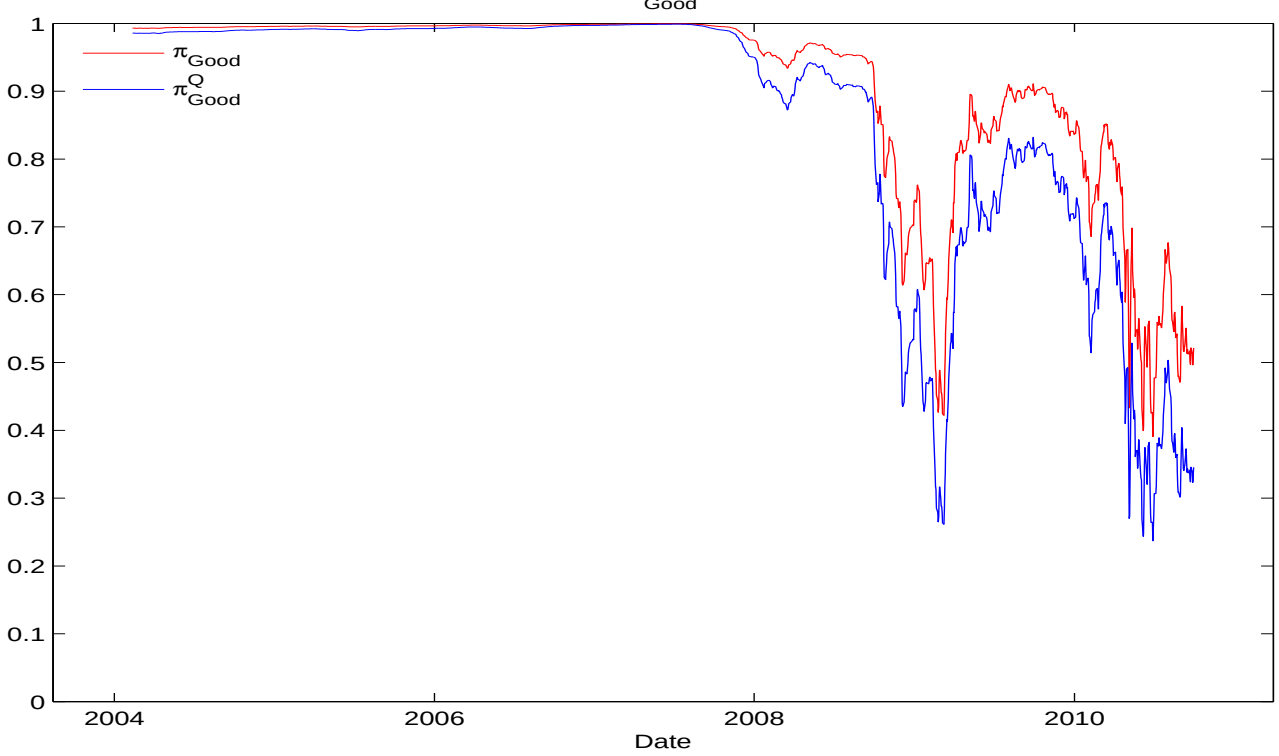

Figure 4: **Hidden State Probability.** The red line shows the unscented Kalman filter estimate of the probability that the economy is in the 'good' state, π_{Good}. The blue line shows the risk-adjusted measure that the economy is in the 'good' state, π^{Q}_{Good}. The sample period is 02/12/2004-09/30/2010.

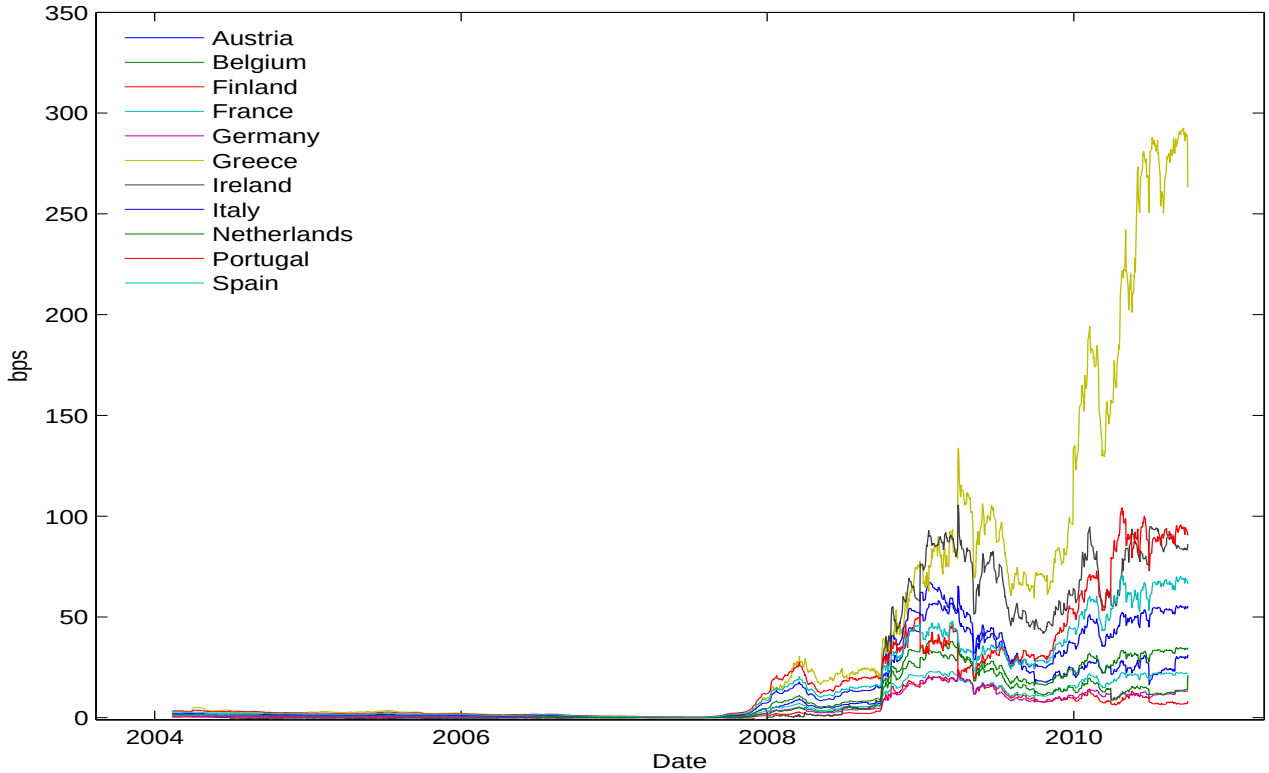

Figure 5: **Fragility Risk Premium**. The plots show the component of CDS spreads that the model attributes to the agents' aversion to model uncertainty, measured by the difference between (1) CDS prices predicted by the model in the presence of fragile beliefs and (2) CDS prices computed under the same coefficients, except that the parameter of uncertainty aversion $\zeta \Rightarrow \infty$. The sample period is 02/12/2004-09/30/2010.

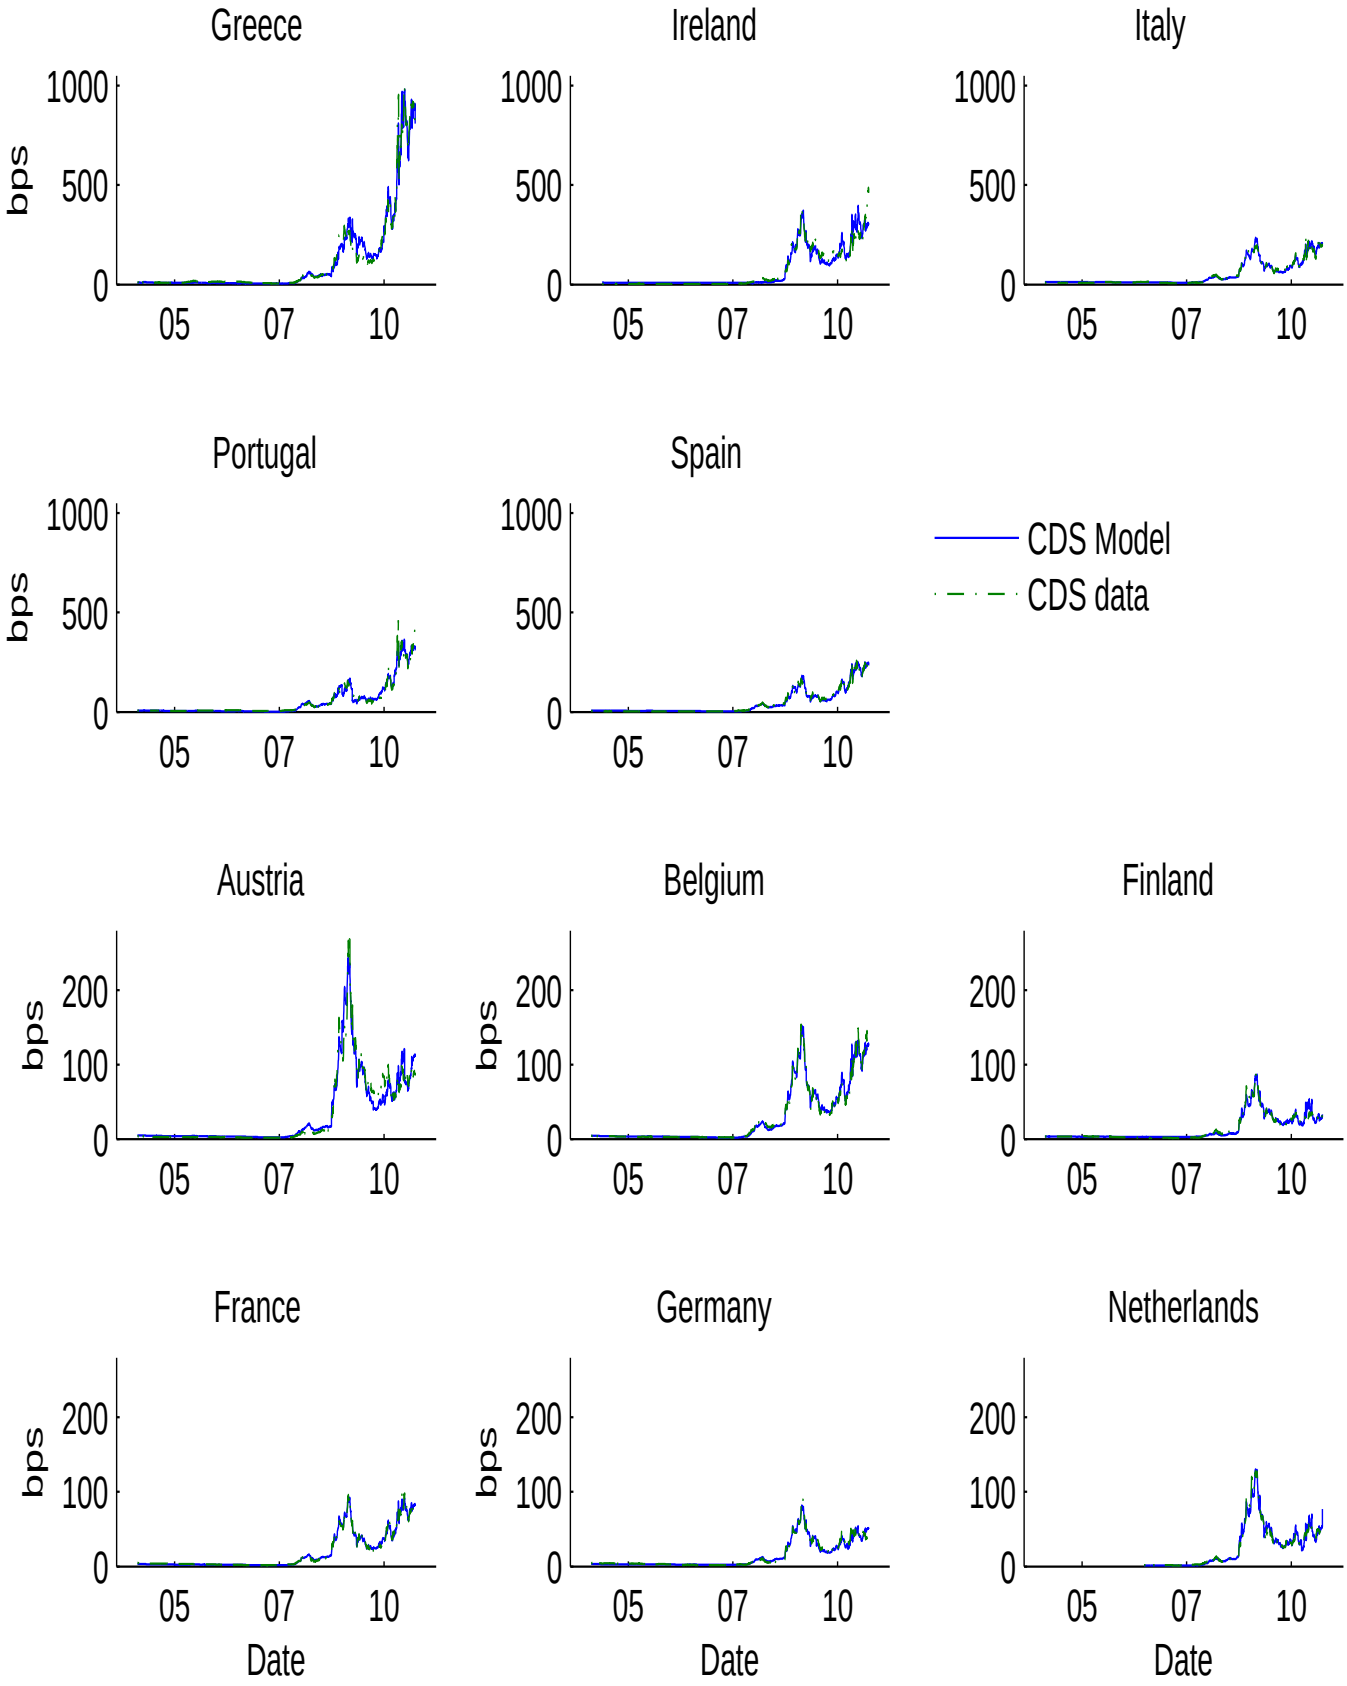

Figure 6: **Sovereign CDS spreads: Model vs. Data.** The plots contrast model-implied 5-year sovereign CDS spreads with actual data for Euro zone countries. We estimate the model via quasi maximum likelihood in combination with the unscented Kalman filter. The sample period is 02/12/2004-09/30/2010.

Table 1: **Sovereign CDS Spreads: Summary Statistics.** The table shows summary statistics for the daily 5-year sovereign CDS spreads across Euro-zone countries. The data are from Markit Financial Information Services and span the period from 02/12/2004 to 09/30/2010, with the exception of Finland, Ireland, and the Netherlands, for which data are unavailable over the periods 09/19/2005-05/01/2006 (Finland); 02/12/2004-05/18/2004 (Ireland) and 02/12/2004-06/23/2006 (Netherlands).

	Mean	Max	Std. Dev.	Skewness	Kurtosis
Full sample period: 02/12/2004-09/30/2010					
Austria	31.47	268.88	47.94	1.88	6.85
Belgium	27.24	155.53	37.71	1.62	4.66
Finland	14.32	92.23	17.50	1.74	6.08
France	18.09	99.25	24.38	1.53	4.16
Germany	13.92	91.38	17.56	1.63	5.32
Greece	114.68	1015.24	210.20	2.63	9.32
Ireland	68.64	486.69	100.10	1.47	4.40
Italy	48.47	249.73	58.60	1.38	3.70
Netherlands	26.10	127.83	28.85	1.42	4.79
Portugal	50.79	461.22	81.43	2.46	8.84
Spain	44.07	273.62	63.03	1.71	5.25
Pre-crisis sample period: 02/12/2004-12/31/2007					
Austria	2.54	6.08	0.72	1.81	8.91
Belgium	3.19	13.20	1.49	4.33	25.17
Finland	2.57	5.75	1.00	0.17	2.35
France	2.81	7.54	1.02	1.25	6.49
Germany	2.91	6.50	1.08	0.58	2.98
Greece	10.60	21.70	3.34	0.50	2.97
Ireland	3.22	14.69	1.92	3.92	20.07
Italy	9.57	21.21	2.47	1.60	7.13
Netherlands	2.06	6.34	1.03	2.45	9.17
Portugal	6.54	19.35	2.14	2.43	12.51
Spain	3.85	19.70	2.50	4.17	22.29
Euro crisis sample period: 01/01/2008-09/30/2010					
Austria	71.69	268.88	52.08	0.96	4.72
Belgium	60.69	155.53	38.41	0.74	2.44
Finland	28.01	92.23	17.71	1.21	4.64
France	39.34	99.25	25.36	0.46	2.07
Germany	29.24	91.38	18.23	0.68	3.71
Greece	259.50	1015.24	263.89	1.46	3.85
Ireland	152.11	486.69	101.84	0.50	3.01
Italy	102.61	249.73	56.27	0.41	2.09
Netherlands	39.26	127.83	28.24	1.18	4.22
Portugal	112.36	461.22	96.65	1.49	4.22
Spain	100.04	273.62	64.09	0.95	3.06

Table 2: **Economic/political risk indicators and the MCIs.** For each Euro zone country i, the table shows the sign and statistical significance of the loading β_i^j on the MCI_i index for the economic/political risk indicator y_i^j. The sample period is 01/03/2001-09/30/2010. The state-space model is

$$\begin{aligned} y_{i,t} &= \beta_i\, MCI_{i,t} + \Gamma_i\, W_{i,t} + \varepsilon_{i,t},\ \varepsilon_i \sim N(0, Q_i) \\ MCI_{i,t+1} &= \rho_i\, MCI_{i,t} + \eta_{i,t+1}, \qquad \eta_i \sim N(0, R_i) \end{aligned}$$

	EU		*Austria*		*Belgium*		*Finland*		*France*		*Germany*		*Greece*		*Ireland*		*Italy*		*Netherlands*		*Portugal*		*Spain*	
Pol. stability	−		+	✓	+		−		−		−		−		−	✓	+		+		−		−	✓
Inflation	−	✓	−	✓	−	✓	−	✓	−	✓	−	✓	+		−	✓	−	✓	−		−	✓	−	✓
M3/GDP	+		+		+		+		−		+		−		+		+	✓	+		+		+	
Reserves/GDP	+		−		+		+	✓	−		+		−		+		+		+		−	✓	+	
UE rate	+	✓	+	✓	+	✓	+	✓	+	✓	+	✓	+	✓	+	✓	+	✓	+		+	✓	+	✓
Real GDP growth	−	✓	−	✓	−	✓	−	✓	−	✓	−	✓	−	✓	−		−	✓	−	✓	−	✓	−	✓
Exports/GDP	−	✓	−	✓	−	✓	−	✓	−	✓	−	✓	−		+		−	✓	−		−	✓	−	✓
Real GDP/Pop.	−	✓	−	✓	−	✓	−	✓	−	✓	−	✓	−	✓	−	✓	−	✓	−	✓	−	✓	−	✓
Gov. surplus/GDP	−		−	✓	+		−		−		+		−	✓	−	✓	−	✓	−		−		−	
Debt/GDP	+		−		+	✓	+		+	✓	−		+	✓	+	✓	−		+	✓	+		+	

Legend:
$(+) \Rightarrow \beta_i^j > 0$
$(-) \Rightarrow \beta_i^j < 0$
✓$\Rightarrow \beta_i^j$'s p-val $< .1$

Table 3: **Linear Model Regressions.** For each Euro zone country, the table shows OLS regressions for the models,

$$\text{Model 1}: \quad \text{CDS}_i = b_{0,i} + b_{1,i}\,\text{MCI}_i + b_{2,i}\,\text{MCI}_{\text{EU}_i^\perp} + b_{3,i}\log\text{ VIX} + b_{4,i}\,\text{BB-BBB spread} + b_{5,i}\,\text{PC}^1_{\text{CDS}_{-i}} + \varepsilon_i$$

$$\text{Model 2}: \quad \Delta\text{CDS}_i = b_{1,i}\,\Delta\text{MCI}_i + b_{2,i}\,\Delta\text{MCI}_{\text{EU}_i^\perp} + b_{3,i}\,\Delta\log\text{ VIX} + b_{4,i}\,\Delta\text{BB-BBB spread} + b_{5,i}\,\Delta\text{PC}^1_{\text{CDS}_{-i}} + \varepsilon_i\,.$$

The dependent CDS_i variable is the daily 5-year credit default swap basis-point spread for country i. MCI_i is the macroeconomic conditions index for country i. $\text{MCI}_{\text{EU}_i^\perp}$ is the residual of an OLS regression of the daily Euro zone MCI_{EU} against the daily country-i MCI_i and a constant. The log VIX variable is the logarithmic daily S&P 500 percentage implied volatility index published by the Chicago Board Option Exchange. The BB-BBB spread variable is the difference between Bank of America Merrill Lynch corporate bond effective percentage yield indices, sampled at the daily frequency. To compute the $\text{PC}^1_{\text{CDS}_{-i}}$ variable, for each Euro zone country $j \neq i$ we regress CDS_j against the MCI_j and the logarithmic VIX. $\text{PC}^1_{\text{CDS}_{-i}}$ is the first principal component extracted from the panel of the regression residuals (we exclude Finland, Ireland, and the Netherlands due to data availability). Compared to Model 1, Model 2 has both left- and right-hand side variables in differences rather than in levels, e.g., we measure ΔCDS_i with the series of daily overlapping changes in the 5-year CDS spreads for country i relative to the prior month, $\Delta\text{CDS}_i(t) = \text{CDS}_i(t) - \text{CDS}_i(t-21)$. The sample period goes from 02/12/2004 to 09/30/2010. The results for Austria, Belgium, Finland, Ireland, the Netherlands, Portugal, and Spain are in the Online Appendix. The coefficient t-ratios in brackets are based on (Newey-West) heteroskedasticity and autocorrelation robust standard errors.

	Panel A: France											
	(1.1)	(1.2)	(1.3)	(1.4)	(1.5)	(1.6)	(2.1)	(2.2)	(2.3)	(2.4)	(2.5)	(2.6)
MCI_i	1.02	1.02		0.30	0.13	0.79	-0.30	-0.42		-0.35	-0.18	-0.06
	(6.67)	(6.84)		(1.79)	(0.79)	(12.88)	(-1.02)	(-1.39)		(-1.30)	(-0.74)	(-0.41)
$MCI_{EU_i^\perp}$		-0.80						-0.62				
		(-1.45)						(-2.29)				
$log\ VIX$			37.97	33.97	16.15	27.58			14.37	14.79	6.68	16.56
			(10.26)	(6.81)	(2.35)	(20.18)			(4.79)	(4.64)	(2.17)	(6.75)
BB-BBB spread					9.35						7.05	
					(3.22)						(3.36)	
$PC^1_{CDS_{-i}}$						0.13						0.07
						(23.65)						(7.14)
Adj. R^2	0.21	0.22	0.46	0.48	0.51	0.88	0.01	0.04	0.11	0.14	0.24	0.50
mean(abs(ε))	15.09	15.49	12.15	11.66	10.92	6.40	4.36	4.40	4.78	4.97	4.51	3.87
max(abs(ε))	86.08	82.45	63.15	66.21	67.56	33.17	46.33	43.64	49.97	42.60	38.47	35.14

Table 3, continued

Panel B: Germany

	(1.1)	(1.2)	(1.3)	(1.4)	(1.5)	(1.6)	(2.1)	(2.2)	(2.3)	(2.4)	(2.5)	(2.6)
MCI_i	0.81	0.81		0.15	0.02	0.83	-0.27	-0.42		-0.31	-0.22	-0.14
	(3.72)	(4.31)		(0.94)	(0.14)	(7.87)	(-1.07)	(-1.26)		(-1.31)	(-1.07)	(-0.93)
$MCI_{EU_i^{\perp}}$		1.81						-0.39				
		(7.85)						(-1.13)				
$log\ VIX$			29.40	28.08	7.75	22.00			10.43	10.83	4.13	12.16
			(9.21)	(8.22)	(1.66)	(15.02)			(4.15)	(4.12)	(1.50)	(5.59)
BB-BBB spread					9.99						5.87	
					(4.10)						(2.77)	
$PC^1_{CDS_{-i}}$						0.08						0.05
						(11.95)						(5.32)
Adj. R^2	0.15	0.47	0.54	0.54	0.61	0.80	0.02	0.03	0.09	0.12	0.23	0.38
mean(abs(ε))	12.19	8.60	8.70	8.54	7.53	5.76	3.44	3.39	3.68	3.83	3.46	3.39
max(abs(ε))	64.27	62.14	51.53	50.25	43.62	30.99	41.38	41.76	43.02	40.66	35.85	32.54

Panel C: Greece

	(1.1)	(1.2)	(1.3)	(1.4)	(1.5)	(1.6)	(2.1)	(2.2)	(2.3)	(2.4)	(2.5)	(2.6)
MCI_i	10.62	10.62		10.73	10.74	1.48	0.95	0.81		0.85	1.04	0.73
	(10.94)	(17.31)		(9.92)	(9.83)	(1.65)	(1.01)	(1.01)		(0.92)	(1.15)	(1.47)
$MCI_{EU_i^{\perp}}$		-5.19						-2.51				
		(-6.16)						(-1.84)				
$log\ VIX$			199.49	-8.42	3.52	170.83			62.61	61.87	38.96	215.50
			(5.82)	(-0.42)	(0.11)	(6.45)			(2.36)	(2.35)	(1.21)	(7.18)
BB-BBB spread					-5.68						20.28	
					(-0.34)						(1.93)	
$PC^1_{CDS_{-i}}$						1.61						1.22
						(11.45)						(7.33)
Adj. R^2	0.73	0.82	0.17	0.73	0.73	0.91	-0.03	0.01	0.02	0.02	0.04	0.64
mean(abs(ε))	81.49	65.67	111.00	81.66	81.81	44.12	24.78	26.04	27.13	27.46	26.64	17.37
max(abs(ε))	481.63	368.16	817.56	482.66	478.67	273.40	530.28	524.57	510.57	514.43	521.18	280.36

Table 3, continued

	Panel D: Italy											
	(1.1)	(1.2)	(1.3)	(1.4)	(1.5)	(1.6)	(2.1)	(2.2)	(2.3)	(2.4)	(2.5)	(2.6)
MCI_i	2.27	2.27		-1.00	-1.87	1.57	-0.08	-0.66		-0.33	-0.41	0.48
	(3.49)	(3.51)		(-1.55)	(-2.61)	(4.04)	(-0.13)	(-1.00)		(-0.63)	(-0.96)	(1.42)
$MCI_{EU_i^{\perp}}$		3.97						-1.36				
		(5.10)						(-2.15)				
$log\ VIX$			95.48	107.58	36.14	76.43			30.32	31.32	9.02	30.67
			(11.48)	(8.08)	(2.32)	(11.24)			(4.63)	(4.40)	(1.13)	(5.93)
BB-BBB spread					38.22						20.10	
					(5.17)						(3.97)	
$PC^1_{CDS_{-i}}$						0.29						0.16
						(15.82)						(6.27)
Adj. R^2	0.11	0.24	0.51	0.52	0.61	0.80	-0.02	0.02	0.10	0.10	0.28	0.46
mean(abs(ε))	40.43	36.07	29.05	30.10	25.87	19.79	9.69	10.20	10.70	10.82	9.67	8.68
max(abs(ε))	223.23	211.05	154.89	144.19	139.07	104.57	105.49	104.49	93.58	94.95	93.34	88.07

Table 4: **CDS Pricing Errors: Model vs. OLS Regressions.** The table compares 5-year CDS spread pricing errors from the model to those from the OLS regressions. We estimate the sovereign CDS spreads model via quasi maximum likelihood in combination with the unscented Kalman filter. For each country, the OLS linear regressions explain the sovereign CDS spreads with the country-specific MCI, the logarithmic VIX, and a proxy for a latent factor (this is the regression in column 6 of Table 3). The sample period is 02/12/2004-09/30/2010.

	Austria	Belgium	Finland	France	Germany	Greece	Ireland	Italy	Netherlands	Portugal	Spain
						mean absolute error					
Contagion Model	5.72	2.05	1.95	1.27	1.70	16.04	14.83	4.78	2.54	6.84	4.66
Linear OLS fit	13.94	11.41	4.99	6.40	5.76	44.12	19.24	19.79	8.09	38.23	24.25
						max absolute error					
Contagion Model	55.68	23.67	19.76	9.36	12.55	201.23	178.34	45.29	28.87	118.47	38.17
Linear OLS fit	109.21	55.82	39.83	33.17	30.99	273.40	183.41	104.57	40.22	231.08	102.51