Contagion in the European Sovereign Debt Crisis*

A structural network analysis

Brent Glover†
Carnegie Mellon University

Seth Richards-Shubik‡
Lehigh University,
and NBER

May 2021

Abstract

We use a network model of credit risk to measure market expectations of the potential spillovers from a sovereign default. Specifically we develop an empirical version of the Eisenberg and Noe (2001) framework for financial contagion, which emphasizes a direct mechanism that operates through balance sheets. We estimate the model with data on sovereign credit default swap spreads and the detailed structure of financial linkages among thirteen European sovereigns from 2005 to 2011. Simulating the estimated model, we find that a sovereign default typically generates small spillovers to other sovereigns based on this mechanism, but there were non-trivial effects on the credit risk of Portugal.

Keywords: financial networks; sovereign debt crisis; contagion; structural estimation; systemic risk
JEL Codes: D85, F34, F36, G01, L14

*We thank Ana Babus, Irina Balteanu, Steven N. Durlauf, Matthew Elliott, Michael Gofman, Christian Julliard, Andrew Karolyi, Robin Lumsdaine, Raoul Minetti, Christian Opp, Mark Ready, Eli Remolona, Bruce Sacerdote, Alireza Tahbaz-Salehi, Lowell J. Taylor, Martin Weidner, Mark Wright, and participants at the AFA Meetings, Tepper-LAEF Macro-Finance Conference, SED Meetings, EMG-ECB Conference on Emerging Markets Finance, University College London applied seminar, Financial and Economics Networks Conference at the University of Wisconsin, Inquire Europe Symposium, the Minnesota Junior Finance Conference, the CIRANO-Sam M. Walton College of Business workshop on Networks in Trade and Finance, and the Barcelona GSE Summer Forum workshop on Networks for helpful comments.

†Tepper School of Business, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, E-mail: `gloverb@andrew.cmu.edu`

‡Department of Economics, Lehigh University, 621 Taylor St, Bethlehem, PA 18015, E-mail: `sethrs@lehigh.edu`

1 Introduction

Over the past decade, there have been many studies of contagion in the European sovereign debt crisis.[1] They offer sophisticated analyses of comovements in sovereign bonds and other assets, using factor models, vector autoregressions, and related techniques to isolate the effects of contagion from those of economic fundamentals and other confounders. In these analyses, contagion is typically defined in terms of the transmission of shocks in observable asset prices, and how this transmission changes during known crisis periods or in response to extreme shocks.[2] Less is known, however, about the effects of specific causal mechanisms that are defined by a theoretical model of financial contagion. The purpose of this paper is to use such a model to estimate the effect of one well-defined mechanism for contagion during the European sovereign debt crisis.

Our model is an empirical version of the classic Eisenberg and Noe (2001) framework for contagion in financial networks. In that framework, the mechanism for contagion is the direct losses that occur when a member of a financial network defaults. These losses may trigger additional defaults among other members that hold claims on the defaulting party (e.g., sovereign bonds), in what could be referred to as a "direct loss" or "balance sheet" mechanism for contagion.[3] We estimate the effect of this mechanism at the aggregate level between sovereigns, using data on credit default swap (CDS) spreads for sovereign bonds and the cross-holdings of sovereign debt among thirteen European countries. While national governments typically do not hold claims on each other directly, this mechanism is relevant at the aggregate level because sovereigns take on risk from their domestic banks via explicit or implicit guarantees and other bailouts.[4] Additionally, both Allen and Gale

[1]See Constancio (2012) for an early survey and Caporin, Pelizzon, Ravazzolo, and Rigobon (2018) for a recent example.

[2]Forbes (2012) provides a survey and thoughtful discussion of the definition of contagion.

[3]Glasserman and Young (2016) refer to this mechanism as "direct loss spillovers through default" and specify their model in terms of a network of balance sheets.

[4]See, for example, Bolton and Jeanne (2011), Alter and Schüler (2012), Merler, Pisani-Ferry, et al. (2012), De Bruyckere, Gerhardt, Schepens, and Vander Vennet (2013), Acharya, Drechsler, and Schnabl (2014), Alter and Beyer (2014), Kallestrup, Lando, and Murgoci (2016), and Farhi and Tirole (2018).

(2000) and Elliott, Golub, and Jackson (2014) suggest that their network models, which have mechanisms related to ours, could be applied to a network of countries.

To accomplish our analysis, we first develop the empirical model and consider its requirements for identification. While there is an extensive theoretical literature on contagion in financial networks (see Allen and Babus, 2009; Glasserman and Young, 2016, for surveys), there are few empirical applications of the models proposed in that literature,[5] and we are not aware of any papers that structurally estimate a model based on the Eisenberg and Noe (2001) framework.[6]

Our empirical model turns out to be straightforward to estimate, fits the data well, and yields informative results. We believe it could be useful in many other applications on financial networks.

The key data for the model are on financial linkages and default probabilities. We use data from the Bank for International Settlements (BIS) and the International Monetary Fund (IMF) to construct a network of aggregate bilateral claims among the sovereigns in our sample, for each quarter from 2005 to 2011. We impute the market expectations for their risk-neutral default probabilities from the spreads on their 5-year CDS contracts. The model is estimated with these data, along with quarterly data on GDP.

We then use the estimated model to make a series of simulations that quantify the effects of the balance sheet mechanism for contagion. Each simulation assumes a default by one sovereign and computes the predicted changes in the default probabilities of the other sovereigns in the network. The results indicate that the potential for contagion from this mechanism was small on average, but there were substantial effects on certain sovereigns. In particular, this mechanism contributed to the risk of contagion from Greece to Portugal at the height of the crisis: our simulations indicate that in 2011-Q1, for example, the direct

[5] Elliott, Golub, and Jackson (2014) and Glasserman and Young (2015) use relevant empirical data to provide informative numerical illustrations, but these are not intended to be econometric analyses.

[6] Cohen-Cole, Patacchini, and Zenou (2011) and Denbee, Julliard, Li, and Yuan (2021) apply linear social interactions models to interbank networks. These papers are discussed in more detail below. Gofman (2017) estimates a network-based model of the interbank lending market to evaluate the effects of limiting bank size and interconnectedness.

losses from a Greek default would have increased the probability of a Portuguese default by 100 bps, from a base of 750 bps.

The default simulations also provide a natural measure of the systemic risk posed by each sovereign. The measure represents the expected spillover losses based on the increased default probabilities of other sovereigns in the network, given a default by the focal sovereign. These losses are normalized by the amount of external debt in the initial default, so that the measure indicates the potential for contagion per unit of a sovereign's debt. This is closely related to other measures that have been proposed in the literature, such as the "contagion index" in Glasserman and Young (2015), "node depth" in Glasserman and Young (2016), and "systemicness" in Bonaldi, Hortaçsu, and Kastl (2015). With this measure we show how the risk of contagion from the balance sheet mechanism evolved during the crisis, and we arrive at a potentially surprising result for the country with the greatest potential for contagion per unit of debt: Austria. Because much of Austria's debt was held by Italy, a financially vulnerable sovereign with substantial external debt, the model predicts relatively high spillover losses in the event of an Austrian default (although the probability of such an event was relatively small).

The mechanism in our model is one of multiple mechanisms for contagion that were likely at play during the debt crisis. Forbes (2012) discusses several mechanisms and channels that have been considered in the literature on financial contagion. As Glasserman and Young (2016) note, only some of these mechanisms directly involve networks of financial linkages, for example losses in asset values from specific counterparties (e.g., Eisenberg and Noe, 2001), and sudden contractions of liquidity in interbank lending (e.g., Allen and Gale, 2000). Others do not, such as correlations in financial or economic shocks (e.g., Ang and Longstaff, 2013), portfolio rebalancing (e.g., Kodres and Pritsker, 2002), and sudden changes in investor beliefs about credit risk (e.g., Benzoni, Collin-Dufresne, Goldstein, and Helwege, 2015). There are also important internal amplification mechanisms in a sovereign debt crisis: reductions in economic output and falling investor confidence make it harder for a sovereign to make payments and roll over debt, further exacerbating the crisis. Naturally, estimating the effects of different mechanisms may require different data and modeling approaches. More importantly, counteracting different mechanisms would require different policy interventions. For example, a central bank can purchase troubled assets from and provide liquidity to individual entities to counteract direct losses and funding runs, while more general, blanket assurances are sometimes used to attempt to change investor beliefs about credit risk (as in ECB President Lagarde's statement in March 2020 that there were "no limits to our commitment to the Euro"[7]). By focusing on one mechanism, we are able to indicate the scope for particular kinds of interventions, such as targeted purchases of bonds from vulnerable sovereigns. Thus, while the balance sheet mechanism in our model is one of multiple possible mechanisms, and may or may not account for the majority of the contagion among sovereigns during this crisis, it is well defined, reasonably identifiable, and clearly relevant for specific policy actions.

To understand the identification of our model, we draw on the microeconometric literature on social interactions. The conditions for the identification of endogenous spillover effects are clearly defined in that literature (e.g., Blume, Brock, Durlauf, and Ioannides, 2011), and our empirical model falls into a particular class that has been analyzed previously (Krauth, 2006; Soetevent and Kooreman, 2007). The basic simultaneity problem—i.e., that credit risk is jointly determined among interconnected entities—is addressed because our model explicitly solves for equilibrium payments and defaults among the sovereigns in the network. The key identifying variation to estimate the effect of the balance sheet mechanism comes from the relationship between the pairwise correlations in default probabilities and the bilateral financial linkages among sovereigns: in other words, the extent to which differential financial linkages account for the differential comovements in sovereign credit risk. Prior to estimating the model, we show that this variation is present and significant, net of both sovereign and time-period fixed effects that would absorb most of the relevant confounding

[7] `https://www.nytimes.com/2020/03/18/business/ecb-coronavirus-bond-buy.html`

factors. Having this understanding of identification is useful, to go beyond a descriptive assessment of contagion and instead examine a causal mechanism.

To be clear, we make important identifying assumptions in order to estimate the model. Our empirical approach is most closely related to three existing papers that estimate network models of spillovers in interbank markets. Cohen-Cole, Patacchini, and Zenou (2011) and Denbee, Julliard, Li, and Yuan (2021) use detailed data on interbank loans to define networks among banks and then recover spillovers in the costs of making loans or in the value of holding liquid assets, respectively. Bonaldi, Hortaçsu, and Kastl (2015) develop a model of loan repayments in a network of banks, using the Eisenberg and Noe (2001) framework, and estimate a reduced-form network of spillovers in funding costs, which is then interpreted via the model. The models in these papers and ours are fundamentally static, as are the models in theoretical literature on contagion in financial networks. Accordingly, the empirical approach is to apply a static model to repeated observations on the same set of entities, treating each period independently. This involves two key assumptions: that financial linkages established in a previous period are exogenous, and that unobserved shocks are independent over time. Hence, unobserved factors that determine both financial linkages and credit risk over time would bias our results.[8] However a preliminary regression analysis (Section 4.2) shows that the observed relationship between the pairwise correlations in default probabilities and the strength of financial linkages remains significant even when we allow for unrestricted time effects. Separately, because our model is static, it cannot address internal amplification mechanisms and other dynamic factors that have been the focus of the macroeconomic literature on sovereign default.[9] Some of these assumptions may seem stark, but they are required to make the application of a structural network model be feasible, and thereby to make progress on the estimation of a theoretically defined, causal mechanism for contagion

[8]Appendix A provides a discussion of this and other potential sources of bias.

[9]This literature, beginning with the seminal contribution of Eaton and Gersovitz (1981), typically studies the debt issuance, credit risk, and default decisions of a single sovereign, with a small, open economy. Such models would not apply to a network of large, interlinked economies, although Arellano, Bai, and Lizarazo (2017) make progress with a model of two sovereign debtors and a continuum of lenders.

in a financial network.

Several notable papers in the finance literature also draw on similar data and identifying variation as we do, in their analyses of contagion among European sovereigns. Acharya, Drechsler, and Schnabl (2014) and Kallestrup, Lando, and Murgoci (2016) both use BIS data on aggregate bilateral claims to construct weighted measures of exposure to credit risk from foreign sovereigns, and include these in regression models for sovereign or bank CDS spreads. Acharya, Drechsler, and Schnabl (2014) find a small but meaningful association between bank CDS spreads and their home country's aggregate exposure to foreign sovereign CDS spreads (a 2% effect in relative terms). This is roughly similar to our preliminary regression estimate of the association between sovereign default probabilities and their aggregate exposure to foreign sovereign default probabilities (a 0.018 to 0.026 effect in absolute terms, see Table 3). Kallestrup, Lando, and Murgoci (2016) find much larger associations with such aggregate exposures to foreign sovereigns, for both sovereign and bank CDS spreads (about 0.4 in absolute terms). However these methods do not—and are not intended to—address the simultaneity of credit risk among interconnected entities, which is one purpose of our equilibrium model. More broadly, Kalbaska and Gatkowski (2012), De Bruyckere, Gerhardt, Schepens, and Vander Vennet (2013), and Brutti and Sauré (2015) also use CDS spreads to measure sovereign credit risk and BIS data (or similar data from bank stress tests) to construct aggregate exposures to foreign sovereigns. The latter two papers use the variation in foreign exposures to analyze pairwise correlations in CDS spreads, and find that differential exposures account for some to much of the differential comovements in credit risk across countries.[10] Collectively, these existing analyses provide evidence that the key empirical relationship implied by our model—the association between aggregate financial linkages and credit risk comovements—was present during the European sovereign debt crisis.

[10]De Bruyckere, Gerhardt, Schepens, and Vander Vennet (2013) estimate that a one standard deviation increase in a bank's exposure to a foreign sovereign raises its excess correlation with that sovereign by 9%, while Brutti and Sauré (2015) estimate that a one standard deviation increase in a country's exposure to Greek debt raises the transmission rate of shocks in Greek CDS spreads to the sovereign's own CDS spreads by 43%.

The remainder of this paper proceeds as follows. Section 2 presents the theoretical framework for our network model of contagion. Section 3 develops the empirical version of the model, and discusses its identification and estimation. The data are described in Section 4, along with an assessment of our constructed measure of aggregate bilateral claims, and the preliminary regression analysis of the association between these claims and the pairwise correlations in default probabilities. Our main results on the effects of the balance sheet mechanism and systemic risk from each sovereign are presented in Section 5.

2 Theoretical Framework

Our model is based on the Eisenberg and Noe (2001) framework, as are a number of recent papers in the theoretical literature on financial contagion.[11] While there are important distinctions in the details of these models and the results they produce, the broad features are as follows. The models in these papers describe a *payment equilibrium* among a set of financial entities that hold claims on each other and also have outside assets or liabilities that are partly stochastic. Given a network of claims among the entities and realizations of their shocks, the payment equilibrium determines a vector of repayments that clears the system.[12] Default is exogenous and occurs when an entity has insufficient assets to meet all of its obligations in full. Contagion in this framework is therefore understood as defaults or other losses that occur as a consequence of incomplete repayments received from other members of the network. This is an immediate, direct mechanism for the spillovers from a default, which we refer to as the "direct loss" or "balance sheet" mechanism for contagion.

Applying this framework to our context, each country is treated as a single, aggregate financial entity, and countries are connected through their aggregate holdings of each other's sovereign debt. This relies on the close connection in credit risk between the domestic banks

[11]See, for example, Gouriéroux, Héam, and Monfort (2012), Rogers and Veraart (2013), Acemoglu, Ozdaglar, and Tahbaz-Salehi (2015), and Glasserman and Young (2015, 2016).

[12]The network of financial linkages is not endogenized in these models.

and the central government in each country, which has been well documented.[13] Losses in value or increases in risk of bank holdings of foreign sovereign debt impact the government due to the existence of explicit or implicit guarantees, and this in turn affects the credit risk of the government's own sovereign debt.

Thus the entities in our network are sovereigns $i = 1, \ldots, N$. They are observed over a number of time periods $t = 1, \ldots, T$, but each period is treated independently as the payment equilibrium is a fundamentally static solution concept.[14] In each period, sovereigns hold debt claims on each other that were established in a previous period. The face value of sovereign i's gross, aggregate claims on sovereign j, payable at date t, is denoted c_{ijt}. These bilateral claims are collected into a matrix $\mathbf{C}_t$, which defines a weighted, directed graph that constitutes the financial network in period t. Sovereigns have additional obligations to unspecified entities outside the network, so that the total external debt owed by sovereign i in period t, denoted D_{it}, is more than just the sum of the claims on i from the other sovereigns in the network (i.e., $D_{it} \geq \sum_{j \neq i} c_{jit}$). Sovereigns also have access to their country's aggregate economic output, $Y_{it} \in \mathbb{R}^+$, which is stochastic and assumed to evolve exogenously. Finally, sovereigns are exposed to an exogenous financial shock, $X_{it} \in \mathbb{R}$.

The payment equilibrium determines which sovereigns are solvent in a particular period, given their total debt (D_{it}), aggregate output (Y_{it}), financial shocks (X_{it}), and the equilibrium payments on their established claims (c_{ijt}). The total (re)payments received in equilibrium are denoted R_{it}. A sovereign is solvent if its current assets exceed its current liabilities:[15]

$$s_{it} \equiv \mathbb{1} \left\{ R_{it} + Y_{it} + X_{it} > D_{it} \right\}, \tag{1}$$

and we refer to the s_{it} as "solvency indicators." If sovereign j is solvent in period t ($s_{jt} = 1$),

[13] See, for example, Bolton and Jeanne (2011), Alter and Schüler (2012), Merler, Pisani-Ferry, et al. (2012), De Bruyckere, Gerhardt, Schepens, and Vander Vennet (2013), Acharya, Drechsler, and Schnabl (2014), Alter and Beyer (2014), Kallestrup, Lando, and Murgoci (2016), and Farhi and Tirole (2018).

[14] This is a limitation, but as we discuss in Appendix A, it does not appear to qualitatively impact our estimate of spillovers from the balance sheet mechanism.

[15] In the empirical model, we allow a nonzero default threshold, which can be positive or negative.

then sovereign i receives the full value of its claims on j; i.e., c_{ijt}. If, on the other hand, sovereign j defaults, its creditors receive a proportion of their claims. This proportion is equal across creditors, and is given by a *recovery function*, r, specified in the model.[16]

We consider two specifications of the recovery function in the empirical analysis. One uses a fixed, exogenous recovery rate, $\delta \in [0, 1)$, so that if country j defaults, its creditors receive δc_{ijt}. This is the "fixed" recovery function:

$$r^{\mathrm{F}} \equiv \begin{cases} 1 & \text{if } s = 1 \\ \delta & \text{if } s = 0 \end{cases}$$

This assumption of a fixed recovery rate is common in the credit risk literature, and the value we choose ($\delta = 0.4$) is consistent with historical recovery rates for sovereign defaults.[17] The discrete losses that occur with this specification can be motivated as a consequence of the renegotiations involved in a sovereign default (see, e.g., Benjamin and Wright, 2009; Yue, 2010). With this recovery function, the total repayments received in period t are

$$R_{it} = \sum_{j \neq i} c_{ijt} r^{\mathrm{F}}(s_{jt}) = \sum_{j \neq i} c_{ijt}[\delta + (1 - \delta)s_{jt}]. \tag{2}$$

The other specification of the recovery function sets the recovery rate equal to the ratio of current assets to current liabilities ($[R_{it}+Y_{it}+X_{it}]/D_{it}$), which follows the model in Eisenberg and Noe (2001). This is the "proportional" recovery function:

$$r^{\mathrm{P}} = \begin{cases} 1 & \text{if } s = 1 \\ (R + Y + X)/D & \text{if } s = 0 \end{cases}$$

[16] Glasserman and Young (2016) define this concept of a recovery function, which maps the current ratio (i.e., the ratio current assets to current liabilities) to a proportional repayment rate. The function is weakly increasing and typically equals 1 when the current ratio is ≥ 1.

[17] In a sample of historical sovereign debt restructurings, Sturzenegger and Zettelmeyer (2008) estimate a range of recovery rates from 30% to 75%. Pan and Singleton (2008) and Longstaff, Pan, Pedersen, and Singleton (2011) use a recovery rate of 25%, while Ang and Longstaff (2013) use 50%.

With this recovery function, the total repayments received in period t are

$$R_{it} = \sum_{j \neq i} c_{ijt} r^{\mathrm{P}}([R_{it} + Y_{it} + X_{it}]/D_{it}) = \sum_{j \neq i} c_{ijt} \max\{0, \min\{1, [R_{it} + Y_{it} + X_{it}]/D_{it}\}\}. \quad (3)$$

Finally, with the designated recovery function, a payment equilibrium can be characterized by a vector of repayments $(R_{it})_{i=1}^{N}$, or equivalently by a vector of solvency indicators $(s_{it})_{i=1}^{N}$, that solve the system of equations defined by (1), using either (2) or (3) to determine the repayments.

Depending on the values of Y_{it} and X_{it} across all countries, there may be multiple solutions to (1) when the fixed recovery function (2) is used. Similar to the models in Rogers and Veraart (2013) and Elliott, Golub, and Jackson (2014), this is a consequence of the discrete loss that occurs with a default. When there are multiple solutions (i.e., multiple equilibria), we follow these papers and select the "best-case" equilibrium in which the fewest countries default.[18] For example, suppose that given the claims, debts, and shocks among all the countries in the network, there are two solutions for countries i and j: either both default ($s_{it} = s_{jt} = 0$) or both are solvent ($s_{it} = s_{jt} = 1$), while all other countries remain solvent. This is possible if i and j are both close to the default threshold and need the repayments from each other in order to remain solvent. In such cases, we always select the equilibrium where marginal countries such as these pay each other back and remain solvent. This would be the result if there were some coordination process, as it is reasonable to presume that all countries would be weakly better off if there were fewer defaults. The best-case solution can be found with a simple iterative procedure: start with repayment amounts as though all countries were solvent; use (1) to determine which countries would, in fact, default; reduce the repayment amounts based on these defaults; use (1) to determine if any additional countries would default; repeat this process until no further countries would default.[19]

[18] As in Elliott, Golub, and Jackson (2014) the set of equilibria constitutes a finite lattice, so there is a well-defined maximum with the fewest defaults.

[19] Eisenberg and Noe (2001), Rogers and Veraart (2013), and Elliott, Golub, and Jackson (2014) use similar algorithms.

Finally, we think it is useful to describe—*informally*—how the payment equilibrium could fit into a larger process for the evolution of the financial network over time. This makes clear the assumptions about timing that are involved in our use of the data. It also helps to clarify how biases could arise if our econometric assumptions are violated, such as the exogeneity of financial linkages. (These potential biases are discussed in detail in Appendix A.) Accordingly, for these limited purposes, we can put the payment equilibrium in the context of a process that repeats over time if we suppose that each period unfolds as follows:

0. Countries are endowed with bilateral claims (c_{ijt}) and total debts (D_{it}).
1. Output (Y_{it}) and financial shocks (X_{it}) are realized.
2. Repayments (R_{it}) are jointly determined in the payment equilibrium for period t.
3. Claims and debts are established for the next period.
4. CDS contracts are traded for credit events in the next period.

To be clear, our model only pertains to the payment equilibrium in step 2. This follows the empirical approaches in Cohen-Cole, Patacchini, and Zenou (2011), and Denbee, Julliard, Li, and Yuan (2021), which similarly estimate structural models of spillovers in financial networks. As in our analysis, these papers apply static equilibrium models to repeated observations on a fixed set of entities. In order to treat each time period independently, any adjustment costs or other dynamic aspects of the decision problems are ignored, and unobserved shocks are assumed to be independent over time.[20] The network of financial linkages can then be considered exogenous if actions in a previous period define the network, as in the timeline above or in Denbee, Julliard, Li, and Yuan (2021), for example.

Any attempt to go beyond this static approach and incorporate the dynamic decision problem in step 3 would confront a substantial challenge of finding equilibrium policy functions for the entities in the network, where the state space involves an $N \times N$ matrix of financial claims. It would also require a number of additional modeling assumptions. We

[20]Bonaldi, Hortaçsu, and Kastl (2015) similarly assume that their reduced-form errors are independent over time.

have chosen instead to follow the above papers in the network literature and estimate a static model of the payment equilibrium, which provides a clear framework to evaluate one potential mechanism for contagion.

3 Empirical Approach

Our goal is to estimate an empirical version of the equilibrium solvency condition (1), which can then be used to quantify the potential spillovers from a sovereign default that arise from the balance sheet mechanism for contagion. Because defaults are not observed in our sample (2005-2011), and in general are very rare among developed sovereigns, we match equilibrium predictions from the model to observable market beliefs about the probability that each sovereign will default. Specifically, we use the observed spreads on sovereign CDS contracts to impute a sovereign's risk-neutral default probability.

To map the data to our model, we suppose that CDS spreads at the end of period $t-1$ reflect the market's assessment of the risk-neutral probability that each sovereign will be solvent in the payment equilibrium in period t. These market expectations should therefore be equal to the expected value of the solvency indicators, s_{it}, conditional on the information available at the end of period $t-1$ (when the claims payable in period t have already been established). We use p_{it} to denote these conditional expectations, taken under the risk-neutral measure $\mathcal{Q}$. Formally, we define these as:

$$p_{it} \equiv \mathbb{E}^{\mathcal{Q}} \left[s_{it} | \mathbf{C}_t, (D_{jt}, Y_{j,t-1}, X_{j,t-1})_{j=1}^{N} \right]. \tag{4}$$

These expectations can be found, given a joint distribution of output $((Y_{jt})_{j=1}^{N})$ and shocks $((X_{jt})_{j=1}^{N})$, conditional on their lagged values, by solving for the payment equilibrium (which gives the vector of solvency indicators, $(s_{jt})_{j=1}^{N})$ over this distribution.

To adapt the theoretical solvency condition in (1) to work with the available data, we need to allow for the fact that the exact amounts of claims and debts due each period, and the available tax revenues for debt payments, are not observed. Our data on bilateral claims (c) and total debts (D) consist of their stocks observed at a quarterly frequency. The measure of aggregate output (Y) is quarterly GDP, and the financial shocks (X) are unobserved. Accordingly, we introduce parameters that express the proportions of these variables that are relevant, on average, for the payment equilibrium in a single period. In addition we allow the threshold required for solvency to take some value other than zero, which could be positive or negative.[21] Thus the empirical version of the solvency condition is specified as

$$s_{it} = \mathbb{1}\left\{\gamma R_{it} - \alpha D_{it} + \beta Y_{it} + X_{it} > \pi_i + \pi_t\right\}. \tag{5}$$

The parameters γ and α express the proportions of the observed financial claims that are payable each period, and β gives the proportion of aggregate output that is available to the central government for payments on its debt obligations. The solvency threshold for sovereign i in period t is $\pi_i + \pi_t$. This threshold varies across sovereigns to capture differences in relatively fixed obligations such as social pension payments, and varies over time to reflect changes in factors like the overall availability of capital.

We then need to specify the forecasted distributions of aggregate output (Y_{it}) and financial shocks (X_{it}) conditional on their values in period $t-1$, so that we can integrate the solutions to (5) over their joint distribution and thereby compute the expectations in (4). For output, we specify the forecasted distribution as a function of its previous level ($Y_{i,t-1}$) and growth rate ($\Delta Y_{i,t-1}$). To capture common macroeconomic trends among the sovereigns in our network, we partition the previous growth rate into a common component and country-specific residuals using a principal components analysis. The common component of the growth rate in country i, denoted $\Delta Y^c_{i,t-1}$, is the first principal component (PC) for period $t-1$ multiplied by the loading for country i. The residual is $\Delta Y^r_{i,t-1} = \Delta Y_{i,t-1} - \Delta Y^c_{i,t-1}$. As

[21] The economic and legal environment of sovereign borrowing is such that there is not a clearly defined default threshold. In the case of a corporate borrower, equityholders would choose to optimally default on their obligations when the value of the equity claim goes to zero. An analogous condition does not exist in the case of a sovereign borrower.

the notation indicates, $\Delta Y^c_{i,t-1}$ varies across countries because it incorporates the loadings. This allows some countries to be more exposed to the aggregate European economy than others. The mean of the forecast for Y_{it} is then specified as a linear combination of the previous level and these two components of the growth rate: $\beta_1 Y_{i,t-1} + \beta_2 \Delta Y^c_{i,t-1} + \beta_3 \Delta Y^r_{i,t-1}$. The distribution of Y_{it} around this mean is assumed to be normal with variance σ^2_Y. Thus, the forecasted distribution of aggregate output for sovereign i in period t is

$$Y_{it} | (Y_{i,t-1}, \Delta Y^c_{i,t-1}, \Delta Y^r_{i,t-1}) \sim \mathcal{N}(\beta_1 Y_{i,t-1} + \beta_2 \Delta Y^c_{i,t-1} + \beta_3 \Delta Y^r_{i,t-1},\ \sigma^2_Y).$$

This represents the market beliefs at the end of period $t-1$.

The shock X_{it} is also specified to have a normal distribution, with mean zero and variance σ^2_X. The variance is the same across countries, but we normalize all variables in levels to be relative to the size of a country's economy. This is equivalent to setting the standard deviation of the financial shocks in each country to be proportional to the size of its economy; e.g., $\sigma_{Xi} = \sigma_X Y_{i0}$, where Y_{i0} is a baseline level of aggregate output (for which we use annual GDP in 2004). Thus, we effectively allow for larger shocks in countries with larger economies.[22] Beyond this, the output and financial shocks are assumed to be independent across countries and over time, as previously discussed.

Applying these specifications, the network-wide vector of conditional expectations in (4), which we refer to as the *risk-neutral solvency probabilities*, can be expressed as follows:

$$(p_{it})_{i=1}^N = \int \mathbb{1}\left\{\gamma R_{it} - \alpha D_{it} + \beta_0(\beta_1 Y_{i,t-1} + \beta_2 \Delta Y^c_{i,t-1} + \beta_3 \Delta Y^r_{i,t-1} + \tilde{Y}_{it}) + X_{it} > \pi_i + \pi_t\right\}_{i=1}^N \cdot \prod_{j=1}^N \frac{1}{\sigma_Y}\phi\left(\frac{\tilde{Y}_{jt}}{\sigma_Y}\right)\frac{1}{\sigma_X}\phi\left(\frac{X_{jt}}{\sigma_X}\right) \mathrm{d}\tilde{Y}_{jt}\mathrm{d}X_{jt},$$

where $\tilde{Y}_{it}$ is the deviation of Y_{it} from its conditional mean and ϕ is the standard normal density. The vector of indicator functions ($\mathbb{1}\{\dots\}_{i=1}^N$) in the integral gives the vector of

[22] This assumption also appears in the theoretical literature (e.g., Glasserman and Young, 2015).

solvency indicators ($(s_{it})_{i=1}^N$) as a function of the vectors of observables and shocks. The interdependencies across sovereigns arise because the repayments (R_{it}) depend on the solvency of other sovereigns (s_{jt}). To simplify this expression, we combine the shocks $\tilde{Y}_{it}$ and X_{it} as $\epsilon_{it} \equiv \tilde{Y}_{it} + X_{it}$ and normalize the parameters so that ϵ_{it} has unit variance (as in a standard probit model). Also, because β_0 is not separately identified from β_1, β_2, and β_3, we set $\beta_0 = 1$. Consequently, the estimates of β_1, β_2, β_3 are to be interpreted as a combination of the forecasts for future output and the availability of output for debt payments. Finally, we use a simple linear trend to capture any changes in the default threshold over time, so that π_t is specified as θt.[23] This yields the ultimate specification that we take to the data:

$$(p_{it})_{i=1}^N = \int 1\left\{\gamma R_{it} - \alpha D_{it} + \beta_1 Y_{i,t-1} + \beta_2 \Delta Y_{i,t-1}^c + \beta_3 \Delta Y_{i,t-1}^r + \epsilon_{it} > \pi_i + \theta t\right\}_{i=1}^N \cdot \prod_{i=1}^N \phi(\epsilon_{it}) \mathrm{d}\epsilon_{it} \quad (6)$$

The integral is computed via simulation. For each vector of pseudo-random draws of $(\epsilon_{it})_{i=1}^N$, we solve the system of equations defined by (5) for the vector of solvency indicators, using the specification above. If there are multiple equilibria (as is possible with the fixed recovery function, r^{F}), the solution algorithm finds the best-case equilibrium, as described in Section 2. Finally, the average of the vectors of solvency indicators across all draws provides an approximation of the vector of solvency probabilities above.

The parameters of (6) are estimated by minimizing the squared error between the empirical, risk-neutral solvency probabilities, derived from the observed CDS spreads, and the predicted solvency probabilities from the above model. The identification of the model is discussed next.

[23] The results in Section 4.2 indicate that a linear time trend fits the data reasonably well and that our conclusions would be robust to more flexible specifications. Having a fixed effect for each time period is problematic because it would greatly increase the parameter space and would raise an incidental parameter problem in our nonlinear model.

3.1 Identification

To consider identification, our empirical model can be understood within a certain class of models from the microeconometric literature on social and spatial interactions. The class consists of static equilibrium models where individual actions are nonlinear functions of the realized actions of other players (i.e., these models are based on simultaneous-move games of complete information, typically with binary actions). In our case it is the solvency outcomes in (5) that are nonlinear functions of the realized solvencies of other countries. Krauth (2006) and Soetevent and Kooreman (2007) are two primary examples that estimate models from this class and provide detailed analyses of identification. Their approaches, like ours, involve making joint predictions for the vector of equilibrium outcomes in order to address the mutual endogeneity of outcomes among interdependent agents. Also both employ selection rules when multiple equilibria are present, as do we—in our case, motivated by the theoretical literature (e.g., Rogers and Veraart, 2013; Elliott, Golub, and Jackson, 2014).

The results in Krauth (2006, Section 2.4.2, p. 251) show that our model is semi-parametrically identified under our assumption that the shocks (ϵ_{it}) are independent across countries and over time. The main difference between our model and those in Krauth (2006) and Soetevent and Kooreman (2007) is that in our case the spillovers take place on a weighted, directed graph (i.e., the network of financial linkages), while in theirs the interaction effects are uniform within groups (e.g., school classrooms where all students are equally connected). The variation in exposures introduced by using individual linkages does not affect the identification arguments in these papers. In fact, based on existing results for *linear* network models, this variation might facilitate identification in circumstances where shocks are correlated across units (see Appendix A). We do not pursue this, however, because the available variation in the data may not provide precise estimates of this correlation.

Appendix A provides a detailed discussion of the biases that might arise if the independence assumptions or other key assumptions in our empirical model were violated. We consider four potential issues: contemporaneous correlations in the shocks between countries, correlations over time and the endogeneity of financial linkages, endogenous default decisions, and internal amplification mechanisms with different impacts across countries. The preliminary regression estimates in Section 4.2, which include country and quarter fixed effects, suggest that any such biases may not affect our results qualitatively.

4 Data

We use data from multiple sources, on bilateral financial claims, total foreign debt, CDS spreads, and GDP, for thirteen European countries, to estimate the model. Table 1 lists the countries included in the sample. The data are quarterly and extend from 2005-Q3 to 2011-Q3. After that time the Greek statistical agency (Hellenic Statstical Authority) suspended seasonal adjustment of GDP, and this study period also avoids the restructuring of Greek sovereign debt in March 2012. Prior to mid-2005, data on CDS spreads for several sovereigns in our sample are not readily available.

The empirical network of debt holdings among the thirteen countries is constructed from data from the BIS and IMF. The BIS provides the total claims held by banks headquartered in one country on entities in another country, at a quarterly frequency.[24] These data come from the central banks of BIS member countries, which collect information on the balance sheet composition of the banks in their jurisdiction, and report to the BIS the aggregate breakdown of the banks' assets according to the country of the issuer of the security. This gives a directed network of aggregate claims among the BIS member countries (and onto non-member countries). However, these represent all financial claims, not just sovereign debt. The IMF reports the dollar amount of a sovereign's debt held by foreign creditors, also at a quarterly frequency. While this gives the amount of a sovereign's debt held abroad, it does not provide the nationalities of the various foreign creditors holding that sovereign's debt. So, to construct our empirical network of sovereign debt holdings, we weight the external sovereign debt amounts reported by the IMF according to the country-specific shares of total

[24]We use the BIS data on consolidated claims on an ultimate risk basis. See Appendix B for discussion.

claims reported by the BIS (see Appendix B for details).[25]

Figure 1 presents a visual representation of our constructed network in 2011-Q1 (the underlying amounts are reported in Appendix Table A-1). The arrows represent the total amounts of claims that banks headquartered in one country (the "creditor country") have on the sovereign debt of another. These amounts are normalized by the size of the economy of the creditor country, using annual GDP in 2004, to reflect the relative exposures. Darker arrows indicate larger proportional amounts, and aggregate claims worth less than one percent of the creditor country's 2004 GDP are not shown. Many countries have substantial aggregate claims on each other, so the arrows can be bi-directional as for example between Austria (AT) and Italy (IT). The algorithm that creates this visual representation places more strongly connected countries in the center and more weakly connected countries in the periphery.[26] Thus Germany (DE) and France (FR) are located near the center because they have substantial claims (outward arrows) and debts (inward arrows) with many other countries. We also see that France and Portugal (PT) have large holdings of sovereign debt from Italy and Greece (GR), respectively, relative to their own 2004 GDP: 28.4% for France from Italy and 12.2% for Portugal from Greece (Appendix Table A-1). In fact, our measure indicates that the exposure to Greek debt (relative to own GDP) was larger for Portugal than for any other country at that time. This will be relevant for our results on the spillover effects in Section 5.2.

The data on CDS spreads come from Credit Market Analytics (CMA). All spreads are for five-year CDS contracts, referencing the sovereign entity and denominated in US dollars. We transform these spreads into implied risk-neutral default probabilities, using the US Treasury yield curve and assuming a 40% recovery rate.[27] From these we compute the time series of

[25]This assumes that the foreign sovereign debt holdings of a country's financial institutions are proportional, on average, to their total foreign asset holdings from each other country. Also, because there are several BIS reporting countries not included in our sample, we are allowing some portion of the sovereigns' debt to be held by countries outside the network in our model (e.g., the United States).

[26]There is not a unique visual representation of the network, however, as it is a projection of an $N \times N$ matrix into two dimensions. Different algorithms (and different initializations) produce different visual representations. Nevertheless the qualitative features are reasonably stable.

[27]We follow the credit risk literature in analyzing risk-neutral default probabilities. See, for example, Ang

annualized, risk-neutral solvency probabilities, for each sovereign in each quarter.[28]

Finally, the GDP data come from the OECD's Quarterly National Accounts database. We use the annualized, seasonally adjusted measure in fixed PPP. Quarterly GDP growth rates are computed and decomposed with a principal components analysis, as described in Section 3. In addition, the common component of the growth rate is detrended by subtracting the average quarterly growth rate for each country over the period from 1995 to 2004.

In Table 2 we provide summary statistics for the solvency probabilities (p_{it}), total claims ($\sum_{j \neq i} c_{ijt}$), and total debt (D_{it}). The average risk-neutral solvency probability is 0.987, but many sovereigns have averages above 0.99 with relatively little variation. The lowest average solvency probabilities and highest standard deviations are seen, as would be expected, for Greece, Ireland, and Portugal, followed by Spain and Italy. The total normalized claim amounts ($\sum_{j \neq i} c_{ijt}$) vary greatly across sovereigns. Ireland, the Netherlands, and Belgium hold large amounts of sovereign debt of other European countries (relative to the size of their own economies), while Greece and Finland have comparatively negligible holdings. Most other countries have claims worth between one third and one half of their 2004 GDP. The normalized debt amounts, which are similar to debt-to-GDP ratios except that GDP is held constant, show the expected differences across countries, with an average close to one. (The GDP variables used in estimation are summarized in Appendix Table A-2.)

4.1 Assessment of the Constructed Network

Our measure of the financial linkages among sovereigns assumes that the allocation of sovereign debt from one country to banks headquartered in other countries is proportional to the allocation of all financial assets captured in the BIS data. We use this constructed network rather than more direct measures of foreign sovereign debt holdings because the latter are not consistently available for the countries in our sample. However, to assess the

and Longstaff (2013). Risk-neutral default probabilities reflect both the objective default probability and a risk premium. As such, they capture the impact of credit risk on a sovereign's cost of borrowing.

[28]See Appendix B for further details.

validity of our constructed network, we compare it with other data from the BIS and from the European Banking Authority (EBA), which are available either for a subsample of countries or at particular points in time.

The BIS data on bilateral foreign claims are available by the sector of the counterparty, including the public sector, for six of the countries in our sample starting in 2010-Q4. Separately, the EBA released data from bank stress tests, which list exposures to the sovereign debt from each country for a sample of large banks. These banks account for a large portion of the banking system in Europe, so adding across the banks headquartered in one country gives a good estimate of the total claims held by banks in that country on the sovereign debt from each other country. The 2011 EBA stress test used data on these exposures as of December 31, 2010. Accordingly, we can make a comparison between these EBA data, and the BIS data on claims on public sector counterparties, against our constructed network, using 2010-Q4. Appendix Table A-3 presents the correlations between these alternative measures and our constructed measure, overall and for each country. The overall correlation with our measure is 0.91 for the BIS public sector debt data and 0.88 for the EBA stress test data, which gives us confidence that our constructed network is reasonably accurate.

4.2 Descriptive Linear Regressions

Now, as a descriptive exercise, we estimate a series of naïve linear regressions using the variables that appear in our network model. These regressions do not account for the joint determination of credit risk in a payment equilibrium, so the coefficients do not have a causal interpretation. Rather, the purpose of this exercise is to assess the variation in the data that identifies our structural parameters. Accordingly, the coefficients should be interpreted simply as conditional correlations. The main specification is

$$p_{it} = a_0 + a_1 t + b \sum_{j \neq i} c_{ijt} p_{jt} + cD_{it} + d_1 Y_{it} + d_2 \Delta Y_{it}^c + d_3 \Delta Y_{it}^r + u_i + v_{it}, \qquad (7)$$

where a_0, a_1, b, c, d_1, d_2, and d_3 are coefficients, and u_i and v_{it} are country fixed effects and random error terms, respectively. The coefficient b reflects the conditional correlation between sovereign i's solvency probability (p_{it}) and the weighted average of its debtor's solvency probabilities (p_{jt}), weighted by the financial linkages (c_{ijt}). This is the same cross-moment that identifies the estimate of γ in our network model, although here the estimate of b is potentially biased by the simultaneity of p_{it} and $p_{jt}, j \neq i$.

The results of this exercise are shown in Table 3. First we estimate (7) with only the time trend and weighted average of debtor solvency probabilities on the right-hand side. When we add the other variables (column 2), the coefficient on the debtor solvencies drops substantially, from 0.040 to 0.026. To interpret these magnitudes, the latter coefficient says that an increase of 100 basis points (bps) in the weighted average of the solvency probabilities among a country's debtors is associated with a 2.6 bps increase in its own solvency probability. Columns 3 and 4 replace the linear time trend ($a_0 + a_1 t$) with time period fixed effects ($a_t 1_t$), which is feasible here because the fixed effects difference out in a linear regression. The coefficients are qualitatively similar to the prior estimates, although the magnitude of the coefficient on the debtor solvency probabilities falls to 0.018 in column 4.[29] The overall similarity indicates that a linear time trend fits the data reasonably well and should not affect our results qualitatively, and furthermore the key relationship between creditor and debtor solvency probabilities is robust to unrestricted time effects.

The regression with time period fixed effects also demonstrates how the financial linkages provide a crucial source of variation to estimate the effect of the contagion mechanism in our model. Because the overall correlation in solvency probabilities at a point in time is absorbed by the fixed effects a_t, it is the variation in financial linkages c_{ijt} across countries at a point in time that yields the estimate of b. This indicates how the estimate of γ in our

[29] As noted in the introduction, Acharya, Drechsler, and Schnabl (2014) find quite similar magnitudes for the association between individual bank CDS rates and foreign sovereign CDS rates in Europe, also using BIS data to weight the exposures to each foreign sovereign. In a specification with time and bank fixed effects, for example, they estimate that a 10% increase in the weighted average of foreign sovereign CDS rates is associated with a 0.2% increase in domestic bank CDS rates.

network model is fundamentally identified by the extent to which differential comovements in solvency probabilities are explained by differential financial linkages. It is thus useful to see that this relationship remains largely intact when time fixed effects are included.

5 Empirical Results

5.1 Estimates and Model Fit

We now present the results from our equilibrium network model, given by equation (6). Parameter estimates and the marginal effects of the associated variables are listed in Table 4. The estimates using the fixed recovery function (Panel A) and the proportional recovery function (Panel B) are nearly identical. The key parameter is γ, which governs the spillovers among sovereigns. It gives the effect of repayments received from other sovereigns, on a sovereign's own (risk-neutral) solvency probability. The average marginal effect is 0.029, which is similar in magnitude to the coefficient on debtor solvencies in the above regressions (Table 3, column 2). The signs of the effects of total debt and the GDP variables are also as expected (recall αD enters (6) negatively).

Figure 2 plots the observed and predicted solvency probabilities to illustrate the dispersion in the data and show the model's fit.[30] In particular, we plot each sovereign's "observed" risk-neutral solvency probability, derived from its 5-year CDS spread, against the model's predicted solvency probability, for each of the 293 quarterly observations in our sample. For most countries the observed and predicted solvency probabilities are quite close to 1. However the figure shows the notable exceptions to this, mainly for Greece, Ireland, and Portugal, and to a lesser extent for Spain and Italy. The model predictions match their empirical counterparts very well, as seen from the fact that most observations fall close to the 45° line. The correlation between the observed and predicted values is 0.966.

[30] This uses the fixed recovery function, but nearly identical results are obtained with either specification.

5.2 Default Simulations

We now use the estimated model to simulate the effect of a default by one sovereign on the risk of default by other sovereigns in the network. This quantifies the potential effects of the balance sheet mechanism for contagion during this period of the European crisis. To simulate the default of a given sovereign j in period t, we fix the solvency indicator for that sovereign at zero ($s_{jt} = 0$) and recompute the solvency probabilities for all other sovereigns using the estimated version of (6). We do this separately for defaults by Greece, Portugal, Italy, or Spain, in each quarter in our study period. These simulations should be considered separately for each quarter, because the model is static and the balance sheet mechanism is immediate, with no cumulative or long-run effects. All of the simulations presented here use the fixed recovery function (3), but nearly identical results are obtained with the proportional recovery function because the parameter estimates are so similar.

Figures 3 and 4 present the results, plotting the increases in the default probabilities ($1 - p_{it}$, in bps) for other selected sovereigns, given a default by the indicated sovereign. Figure 3-A shows that a Greek default could have had substantial impacts on Ireland and Portugal via the balance sheet mechanism: it would have raised their default probabilities by up to 60 bps for Ireland and just over 100 bps for Portugal (from baseline predictions of 400-500 bps and 600-750 bps, respectively, in the relevant quarters).[31] These results can be compared, loosely, to existing estimates of spillover effects from shocks to Greek CDS spreads, based on vector autoregressions of European sovereign CDS spreads. For example, Alter and Beyer (2014) and Brutti and Sauré (2015) report spillover effects from Greek to Portuguese sovereign CDS spreads, of 20% to 30% in relative terms (i.e., a 10% increase in Greek CDS spreads would lead to a 2-3% increase in Portuguese CDS spreads).[32] Our estimate of a 100 bps absolute effect on Portugal in 2011-Q1 translates to a 16% relative

[31] Official seasonally adjusted GDP data are not available for Greece in 2011-Q2 or Q3, so there are no simulations for that period. The reduction in the spillover effect on Ireland in the two preceding quarters is due to a large drop in Irish holdings of Greek debt, from 0.102 to 0.012 in our constructed measure.

[32] See table A4.1 in Alter and Beyer (2014) and table A3 in Brutti and Sauré (2015).

effect.[33] Thus, accounting for the simultaneity of credit risk, we find a somewhat smaller but still important spillover effect from Greece to Portugal via the balance sheet mechanism.

Figure 3-B shows that Portugal in turn posed a modest risk to Spain: for example, a default in 2011-Q3 would have increased Spain's default probability by 36 bps (6%, relative to the baseline prediction of 580 bps). This illustrates the possibility of a chain of contagion from Greece, to Portugal, to Spain, even though the holdings of Greek sovereign debt in Spain were quite small (0.3% of Spain's 2004 GDP, Appendix Table A-1). Portugal did not pose much risk to other sovereigns via the balance sheet mechanism, largely because no countries other than Spain had substantial holdings of its sovereign debt (e.g., worth no more than 1.1% of their 2004 GDP, Appendix Table A-1).

Defaults by Italy or Spain, shown in Figure 4, would have had somewhat larger spillover effects on certain sovereigns, which is natural given their larger amounts of external sovereign debt. However the effects are not universally larger, because the spillovers mainly impact financially vulnerable sovereigns whose banks have substantial holdings of debt from the defaulting country. This indicates the value of simulations based on a model like ours: they reflect both the cross holdings of sovereign debt and the financial vulnerability of the other countries holding the debt, both of which exhibited tremendous variation across Europe at this time. Irish banks had substantial holdings of debt from Italy and (to a lesser extent) Spain, so the simulated spillover effects on the default probability for Ireland are large: over 200 bps from Italy in 2010-Q2 and over 60 bps from Spain in 2011-Q3. On the other hand, the spillover effects from Spain to Portugal are slightly smaller than those from Greece to Portugal, because Portuguese banks had somewhat smaller holdings of debt from Spain than from Greece (according to our measure). Also, Portuguese banks had relatively minimal holdings of debt from Italy. Among the less vulnerable sovereigns, only France appears to have been at risk of contagion from the direct loss mechanism, in the event of a default by

[33] The relative increase over Portugal's baseline default probability of 750 bps is 13% (=100/750), which is in response to a relative increase in Greece's default probability of 84% (because its baseline predicted default probability in that quarter was 16%), hence a relative effect of 16% (=0.13/0.84)

Italy. This is because French banks had enormous holdings of Italian sovereign debt (e.g., equal to 28% of France's GDP in 2011-Q1, Appendix Table A-1). The spillover effect from Italy to France reaches 100 bps in 2011-Q3, nearly a 50% increase over France's baseline default probability of 215 bps.

5.3 Systemic Risk Measure

These default simulations provide a natural measure of the systemic risk posed by each sovereign, based on the total predicted spillover effects from one sovereign's default onto the other sovereigns in the network. This measure is the total expected losses implied by the increased default probabilities among the other sovereigns in the network, normalized by the total external debt of the sovereign with the initial default. In other words, this gives the expected spillover losses per unit of external debt—or, how potentially contagious is a sovereign's debt per unit.

Our measure is computed as follows. Given a default by sovereign j in period t, we use the above simulations to calculate the change in solvency probabilities among the other sovereigns in the network. Let $\hat{p}_{it}$ be the baseline predicted solvency probability for some country i in period t using the estimated model, and let $\tilde{p}_{it}(j)$ be the simulated solvency probability under the counterfactual that country j defaults. These simulated probabilities reflect both the direct effects of the loss of repayments from sovereign j, and any indirect effects from the further losses of repayments from any other sovereigns (k, etc.) that default in this counterfactual. (The higher order sequences of losses are naturally included because each simulation finds a new payment equilibrium given a default by sovereign j.) Then, given the baseline and simulated probabilities, the expected spillover losses from sovereign i due to an initial default by sovereign j is $[\hat{p}_{it} - \tilde{p}_{it}(j)]D_{it}$. We add these losses across all countries, and divide by the total external debt of sovereign j to normalize, which yields our measure:

$$\lambda_{jt} \equiv \frac{1}{D_{jt}} \sum_{i \neq j} [\hat{p}_{it} - \tilde{p}_{it}(j)] D_{it}. \tag{8}$$

This gives the *expected spillover losses per unit of debt* of country j. Because it is normalized by the total amount of debt, this measure emphasizes the effects of a country's position in the network of financial linkages—i.e., who its creditors are, and how sensitive those creditors are to losses—rather than simply the magnitude of the losses in the initial default.

This measure is analogous to the Katz-Bonacich centrality measure that other authors have used to quantify the systemic importance of each entity in a financial network (e.g., Cohen-Cole, Patacchini, and Zenou, 2011; Denbee, Julliard, Li, and Yuan, 2021; Bonaldi, Hortaçsu, and Kastl, 2015). The difference is that our measure is calculated from a nonlinear model, while the Katz-Bonacich measure applies to linear models. Nevertheless the purpose is the same: these measures capture both direct and indirect spillovers; they can be used to identify which entities pose the greatest threat at a point in time, and to analyze the evolution systemic risk over time.

Figure 5 plots λ_{jt} for the most at-risk sovereigns (Panel A) and for five large European economies (Panel B). The magnitudes of these expected spillovers are not terribly large: for each $1 of external debt directly lost in a default, the expected losses from additional defaults by other sovereigns are less than 2.5 cents. The levels and trends are generally similar among all the countries in both panels. The foreign debt from Greece and Portugal had somewhat higher potential for contagion, as did the debt from Germany, followed by Ireland, France, Italy, and Spain, and last the United Kingdom.[34]

Figure 6 shows λ_{jt} for smaller European economies (Panel A) and the weighted average among all the sovereigns in our sample (Panel B). Austria's foreign debt had the greatest potential for contagion according to this measure, with expected spillover losses per unit of debt that were roughly double those of any other sovereign. This turns out to be the case

[34]The UK had lower potential spillover losses than other countries because a relatively large proportion of its debt was held outside Europe (e.g., by the United States). That debt is included in the normalization, but it cannot generate spillover losses from countries outside the network in our sample.

because Italy held a relatively large share of Austria's debt, and Italy was more sensitive to losses because of its higher baseline risk of default. Finally, the weighted average of these spillover losses, which uses the total external debt amounts (D_{jt}) as weights, rose to about 0.6 cents per dollar during the recession of 2008-09, and leveled there until the end of 2010, when it began to rise as the sovereign debt crisis widened in Europe.

A natural concern is whether these expected spillover losses incorporate market beliefs about the likelihood of a bailout for a sovereign at risk of default. Indeed we think it is reasonable to assume that the observed CDS spreads do reflect market beliefs about possible bailouts. Accordingly, this measure should be interpreted as reflecting market expectations about losses that might occur despite efforts to bail out a sovereign (e.g., as in the case of Greece in March 2012). This applies both to the sovereign with the initial default and to the other sovereigns where there might be further bailouts to try to prevent additional defaults.

5.4 Overall Effects on Cost of Borrowing

As a final exercise, we consider how the risk of contagion from the balance sheet mechanism may have affected the cost of borrowing for European sovereigns at this time. To do this, we compare simulations that eliminate the spillover effects in our model against our baseline predictions which fit the observed data, where the risk of contagion was present.[35] Because CDS spreads and bond yields are approximately proportional to default risk, the decreases in default probabilities in these simulations indicate what fraction of the cost of borrowing could be attributed to this form of contagion.

To compute the counterfactual predictions, we set the model so that there are no losses when a default occurs: i.e., the recovery rate is 100% ($\delta = 1$). Then, for each sovereign in each quarter, we can compare the counterfactual predictions, $\hat{p}_{it}^{\mathrm{NC}}$ ("NC" for "no contagion") against the baseline predictions $\hat{p}_{it}$. The proportion of a sovereign's borrowing cost (e.g., its credit spread or bond yield) attributed to the contagion risk from the balance sheet

[35] See Bahaj (2020) for a sophisticated analysis in a similar spirit, which assesses the total effect of contagion risk (via any mechanism) on sovereign borrowing costs, by eliminating the effects of systemic shocks.

mechanism is then measured as $(\hat{p}_{it}^{\text{NC}} - \hat{p}_{it})/(1 - \hat{p}_{it})$.

Table 5 reports the results for 2011-Q1. Overall these proportional changes are small, averaging 2.3% across the thirteen sovereigns (weighted by their external debt amounts, D_{it}). However we see substantial heterogeneity in this measure. The differences across sovereigns are not simply a function of their overall credit risk; rather, the impact of this form of contagion risk depends on a sovereign's linkages in the network. Perhaps somewhat surprisingly, we find that this form of contagion risk had the largest proportional effect on the borrowing costs of France, accounting for 5.44% of that country's total credit spread in 2011-Q1. The effects were also non-trivial, around 3% of the total credit spread, for Austria, Germany, Netherlands, and Portugal. Thus, in a proportional sense, the risk of contagion from the balance sheet mechanism appears to have had larger effects on the cost of borrowing for the more financially secure European sovereigns.

6 Conclusion

In this paper, we build upon the rich theoretical literature on financial networks to construct a network model of credit risk among thirteen European sovereigns. Using data on sovereign CDS spreads and the cross-holdings of sovereign debt from 2005 to 2011, we estimate the model and conduct counterfactual experiments to quantify the spillover effects from a direct, "balance sheet" mechanism for contagion. Our empirical framework and approach, along with the contagion measure that we develop, could be used to quantify expected spillovers of credit risk in other financial networks.

Our estimates imply that credit markets perceived the risk of contagion from direct losses in a sovereign default to be small overall. On average, our estimates imply that the spillovers from this form of contagion accounted for less than five percent of the total cost of borrowing for the sovereigns in our analysis, even at the peak of the crisis. However we find evidence of substantial effects on certain sovereigns, particularly Portugal.

There are many other channels through which contagion may operate, and some may well have contributed more to the transmission of credit risk during the European sovereign debt crisis. Assessing the quantitative importance of other contagion mechanisms, that are precisely defined in the context of a theoretical model, will be important for understanding the risks from different channels, and for designing effective policies to manage future crises.

References

Acemoglu, Daron, Asuman Ozdaglar, and Alireza Tahbaz-Salehi, 2015, Systemic risk and stability in financial networks, *American Economic Review* 105, 564–608.

Acharya, Viral, Itamar Drechsler, and Philipp Schnabl, 2014, A pyrrhic victory? bank bailouts and sovereign credit risk, *The Journal of Finance* 69, 2689–2739.

Allen, Franklin, and Ana Babus, 2009, Networks in finance, in Paul Kleindorfer, Yoram (Jerry) Wind, and Robert Gunther, ed.: *The Network Challenge: Strategy, Profit, and Risk in an Interlinked World* . pp. 367–382 (Wharton School Publishing).

Allen, F., and D. Gale, 2000, Financial contagion, *Journal of Political Economy* 108, 1–33.

Alter, Adrian, and Andreas Beyer, 2014, The dynamics of spillover effects during the european sovereign debt turmoil, *Journal of Banking & Finance* 42, 134–153.

Alter, Adrian, and Yves S Schüler, 2012, Credit spread interdependencies of european states and banks during the financial crisis, *Journal of Banking & Finance* 36, 3444–3468.

Ang, Andrew, and Francis A Longstaff, 2013, Systemic sovereign credit risk: Lessons from the us and europe, *Journal of Monetary Economics* 60, 493–510.

Arellano, Cristina, Yan Bai, and Sandra Lizarazo, 2017, Sovereign risk contagion, Working Paper 24031 National Bureau of Economic Research.

Avdjiev, Stefan, Christian Upper, and Karsten von Kleist, 2010, Highlights of international banking and financial market activity, *BIS Quarterly Review* p. 13.

Bahaj, Saleem, 2020, Sovereign spreads in the euro area: Cross border transmission and macroeconomic implications, *Journal of Monetary Economics* 110, 116–135.

Benjamin, David, and Mark LJ Wright, 2009, Recovery before redemption: A theory of delays in sovereign debt renegotiations, *Available at SSRN 1392539*.

Benzoni, Luca, Pierre Collin-Dufresne, Robert S. Goldstein, and Jean Helwege, 2015, Modeling credit contagion via the updating of fragile beliefs, *The Review of Financial Studies* 28, 1960–2008.

Blume, Lawrence E., William A. Brock, Steven N. Durlauf, and Yannis M. Ioannides, 2011, Identification of social interactions, in Jess Benhabib, Alberto Bisin, and Matthew O. Jackson, ed.: *Handbook of Social Economics*vol. 1 of *Handbook of Social Economics* . chap. 18, pp. 853–964 (North-Holland).

Bolton, Patrick, and Olivier Jeanne, 2011, Sovereign default risk and bank fragility in financially integrated economies, *IMF Economic Review* 59, 162–194.

Bonaldi, Pietro, Ali Hortaçsu, and Jakub Kastl, 2015, An empirical analysis of funding costs spillovers in the euro-zone with application to systemic risk, Princeton university working paper.

Bramoullé, Yann, Habiba Djebbari, and Bernard Fortin, 2009, Identification of peer effects through social networks, *Journal of econometrics* 150, 41–55.

Brock, William A, and Steven N Durlauf, 2007, Identification of binary choice models with social interactions, *Journal of Econometrics* 140, 52–75.

Brutti, Filippo, and Philip Sauré, 2015, Transmission of sovereign risk in the euro crisis, *Journal of International Economics* 97, 231–248.

Caporin, Massimiliano, Loriana Pelizzon, Francesco Ravazzolo, and Roberto Rigobon, 2018, Measuring sovereign contagion in europe, *Journal of Financial Stability* 34, 150–181.

Cohen-Cole, E., E. Patacchini, and Y. Zenou, 2011, Systemic risk and network formation in the interbank market, Discussion paper, CEPR Discussion Papers.

Constancio, V, 2012, Contagion and the european debt crisis, *Financial Stability Review* pp. 109–121.

De Bruyckere, Valerie, Maria Gerhardt, Glenn Schepens, and Rudi Vander Vennet, 2013, Bank/sovereign risk spillovers in the european debt crisis, *Journal of Banking & Finance* 37, 4793–4809.

Denbee, Edward, Christian Julliard, Ye Li, and Kathy Yuan, 2021, Network risk and key players: A structural analysis of interbank liquidity, *Journal of Financial Economics*.

Eaton, J., and M. Gersovitz, 1981, Debt with potential repudiation: Theoretical and empirical analysis, *The Review of Economic Studies* 48, 289–309.

Eisenberg, Larry, and Thomas H Noe, 2001, Systemic risk in financial systems, *Management Science* 47, 236–249.

Elliott, Matthew, Benjamin Golub, and Matthew O. Jackson, 2014, Financial networks and contagion, *American Economic Review* 104, 3115–53.

Farhi, Emmanuel, and Jean Tirole, 2018, Deadly embrace: Sovereign and financial balance sheets doom loops, *The Review of Economic Studies* 85, 1781–1823.

Forbes, Kristin J, 2012, The “Big C”: Identifying and mitigating contagion, in *Proceedings-Economic Policy Symposium-Jackson Hole* pp. 23–87. Federal Reserve Bank of Kansas City.

Glasserman, Paul, and H. Peyton Young, 2015, How likely is contagion in financial networks?, *Journal of Banking & Finance.*

——— , 2016, Contagion in financial networks, *Journal of Economic Literature* 54, 779–831.

Gofman, Michael, 2017, Efficiency and stability of a financial architecture with too-interconnected-to-fail institutions, *Journal of Financial Economics* 124, 113–146.

Gouriéroux, C., J.-C. Héam, and A. Monfort, 2012, Bilateral exposures and systemic solvency risk, *Canadian Journal of Economics* 45, 1273–1309.

Kalbaska, Alesia, and Mateusz Gatkowski, 2012, Eurozone sovereign contagion: Evidence from the cds market (2005–2010), *Journal of Economic Behavior & Organization* 83, 657–673.

Kallestrup, René, David Lando, and Agatha Murgoci, 2016, Financial sector linkages and the dynamics of bank and sovereign credit spreads, *Journal of Empirical Finance* 38, 374–393.

Kodres, Laura E, and Matthew Pritsker, 2002, A rational expectations model of financial contagion, *The journal of finance* 57, 769–799.

Krauth, Brian V, 2006, Simulation-based estimation of peer effects, *Journal of Econometrics* 133, 243–271.

Lee, Lung-fei, Ji Li, and Xu Lin, 2014, Binary choice models with social network under heterogeneous rational expectations, *The Review of Economics and Statistics* 96, 402–417.

Lee, Lung-fei, Xiaodong Liu, and Xu Lin, 2010, Specification and estimation of social interaction models with network structures, *Econometrics Journal* 13, 145–176.

Longstaff, F.A., J. Pan, L.H. Pedersen, and K.J. Singleton, 2011, How sovereign is sovereign credit risk?, *American Economic Journal: Macroeconomics* 3, 75–103.

Merler, Silvia, Jean Pisani-Ferry, et al., 2012, Hazardous tango: sovereign-bank interdependence and financial stability in the euro area, *Financial Stability Review* 16, 201–210.

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

Rogers, L. C. G., and L. A. M. Veraart, 2013, Failure and rescue in an interbank network, *Management Science* 59, 882–898.

Soetevent, Adriaan R., and Peter Kooreman, 2007, A discrete-choice model with social interactions: with an application to high school teen behavior, *Journal of Applied Econometrics* 22, 599–624.

Sturzenegger, Federico, and Jeromin Zettelmeyer, 2008, Haircuts: Estimating investor losses in sovereign debt restructurings, 1998–2005, *Journal of International Money and Finance* 27, 780–805.

Yue, Vivian Z, 2010, Sovereign default and debt renegotiation, *Journal of International Economics* 80, 176–187.

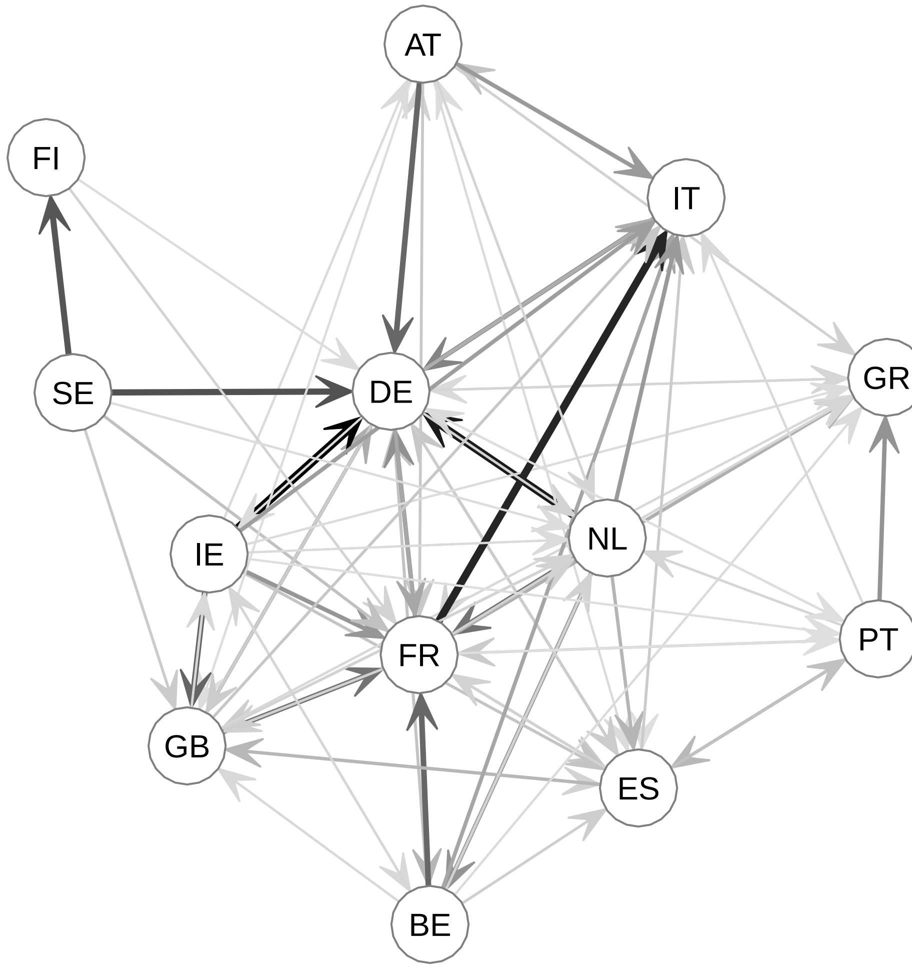

Figure 1: Network Graph, 2011-Q1. The figure illustrates the network structure of aggregate sovereign debt holdings in the first quarter of 2011. Countries are represented by their two letter abbreviation in Table 1. Arrows represent the aggregate bank holdings in one country of the sovereign debt of another. Darker, thicker arrows indicate larger amounts in proportion to the creditor country's GDP in 2004. Holdings worth less than 1% of the creditor country's GDP are not shown.

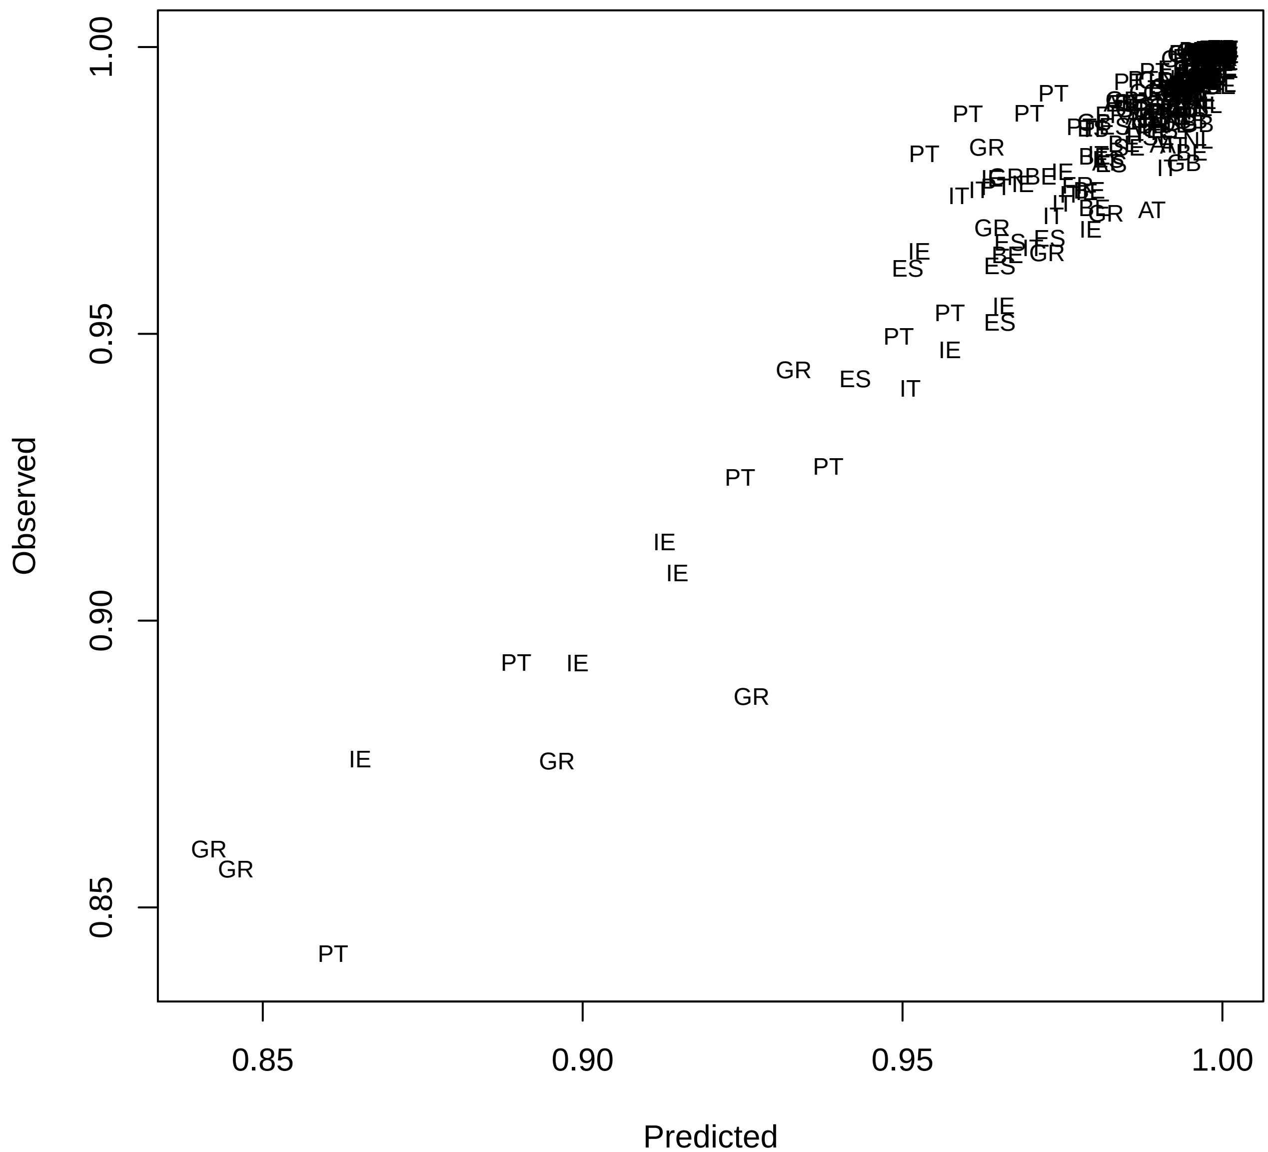

Figure 2: Predicted and Observed Solvency Probabilities. The figure plots the predicted and observed risk-neutral quarterly solvency probabilities for each country at each quarter in our sample. Observed solvency probabilities are obtained with a transformation of 5-year CDS contract prices, as described in Section 4 and Appendix B. Predicted solvency probabilities are generated from the estimated network model, specified in equation (6) with recovery function (2). Country abbreviations are listed in Table 1.

A.

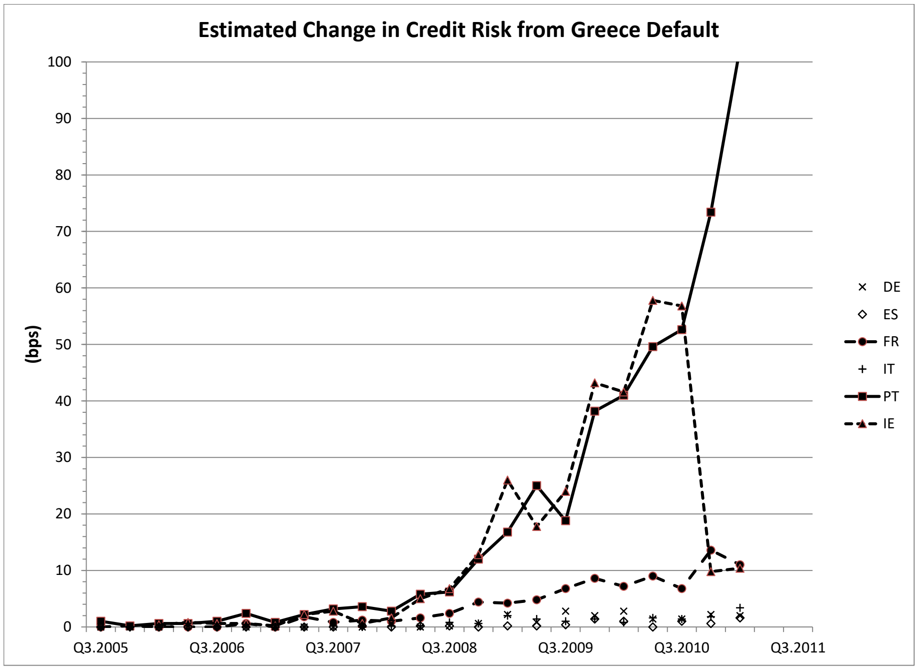

B.

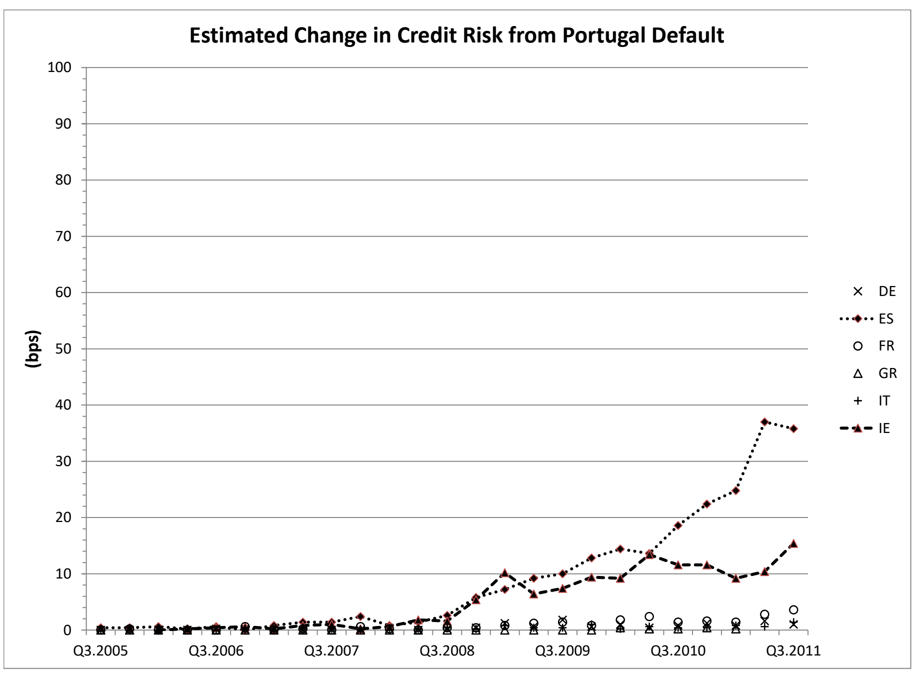

Figure 3: Simulated Default Probabilities: Greece and Portugal. The figure shows the change in the risk-neutral probability of default for selected sovereigns, assuming a default in Greece (panel A) or Portugal (panel B), in simulations using the estimated model. See Section 5.2 for details.

A.

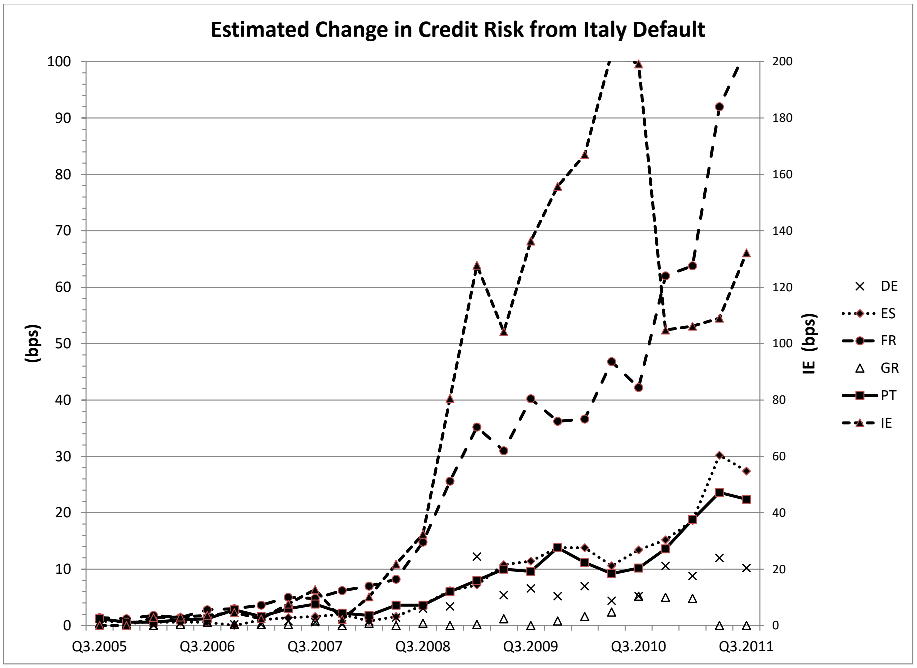

B.

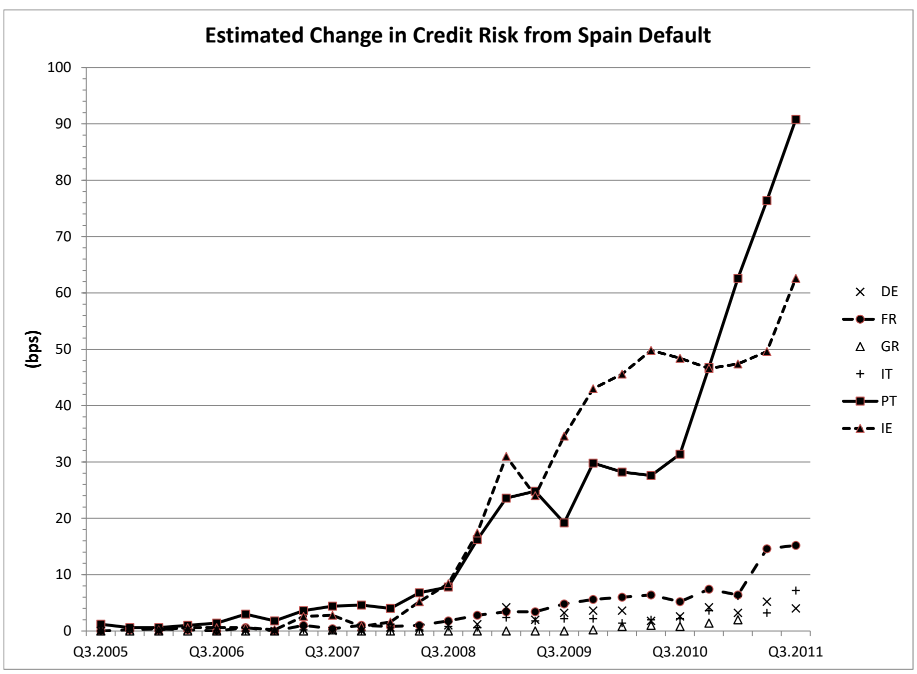

Figure 4: Simulated Default Probabilities: Italy and Spain. The figure shows the change in the risk-neutral probability of default for selected sovereigns, assuming a default in Italy (panel A) or Spain (panel B), in simulations using the estimated model. See Section 5.2 for details.

A.

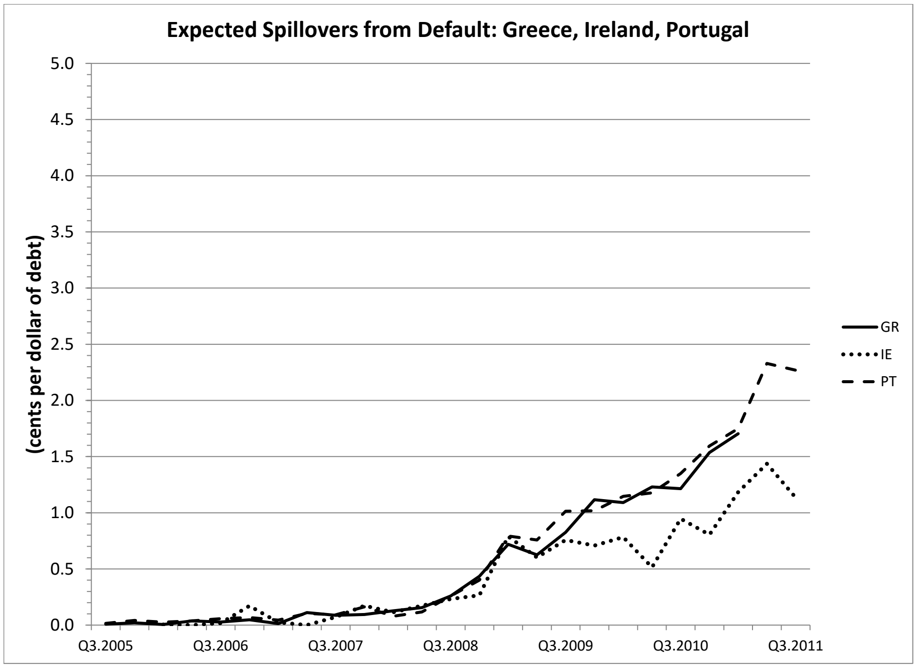

B.

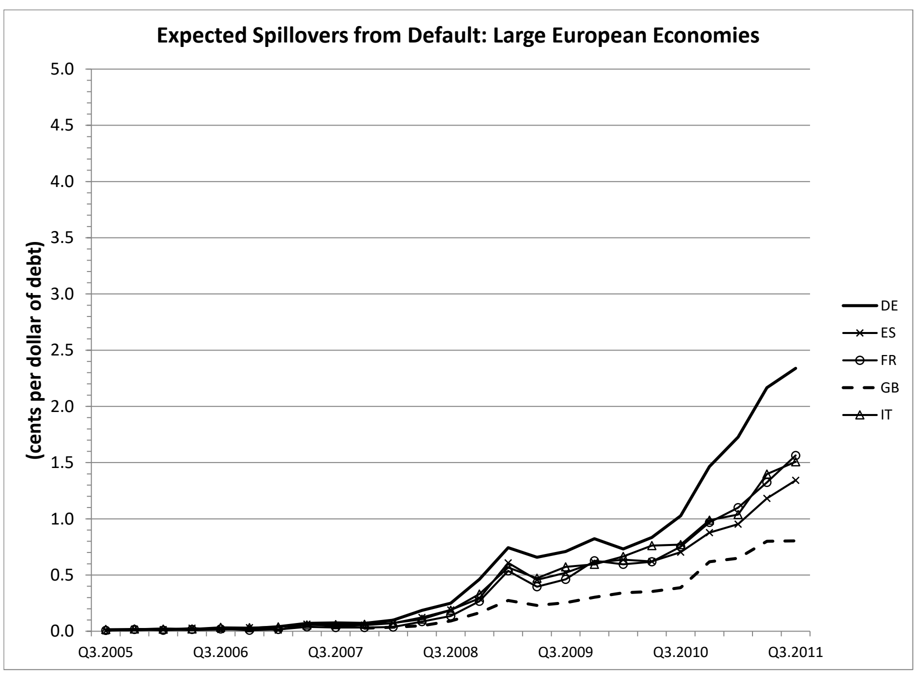

Figure 5: Expected Spillover Losses per Unit of Debt: At-Risk Sovereigns and Large Economies. The figure plots our measure of expected losses due to increased risk of default at other countries, conditional on a default by the country indicated in each series, per unit of that country's external debt. See equation (8) for the definition of this measure.

A.

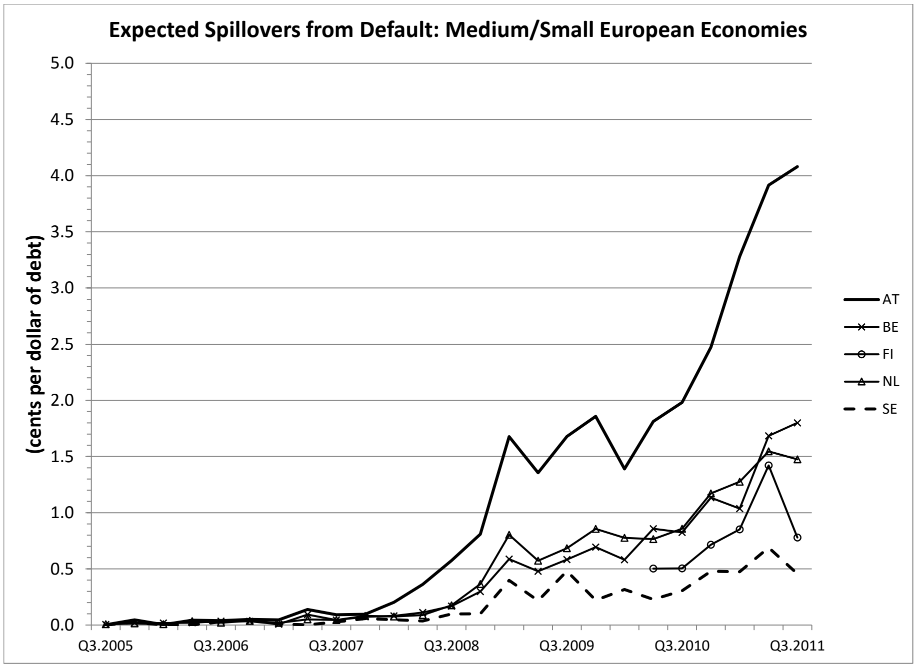

B.

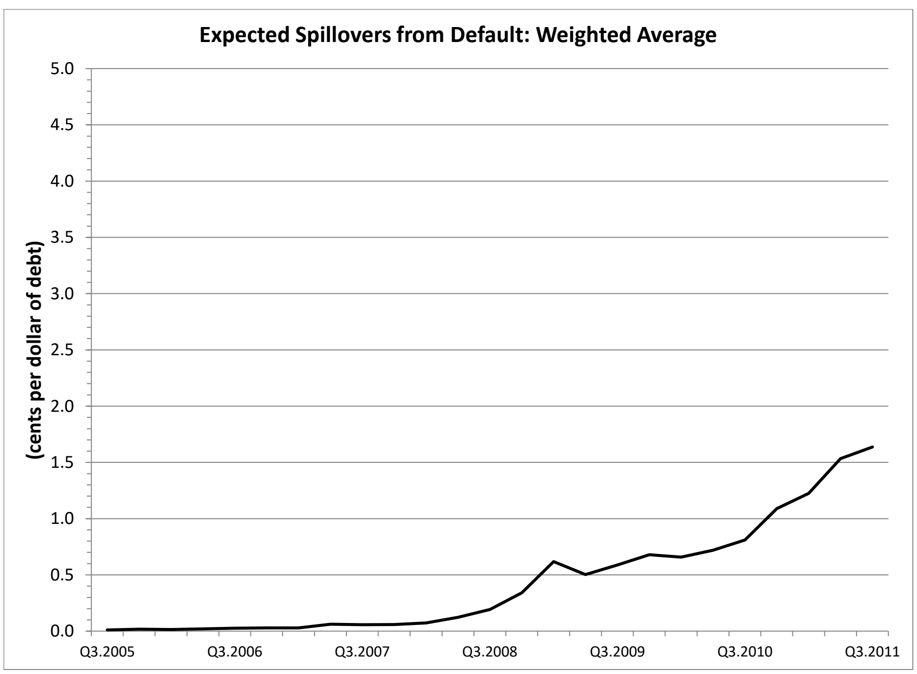

Figure 6: Expected Spillover Losses per Dollar of Debt: Medium or Small Economies, and Weighted Average. The figure plots our measure of expected losses due to increased risk of default at other countries, conditional on a default by the country indicated in each series, per dollar of that country's foreign debt. See equation (8) for the definition of this measure. Weighted average is among all countries in the sample, weighted by each country's total external debt.

Table 1: **List of Sovereigns**

Austria	AT	Italy	IT
Belgium	BE	Netherlands	NL
Finland	FI	Portugal	PT
France	FR	Spain	ES
Germany	DE	Sweden	SE
Greece	GR	United Kingdom	GB
Ireland	IE		

Notes: The table lists the names and abbreviation codes for the sovereigns in our sample.

Table 2: **Summary Statistics for Main Estimation Variables**

	Solvency Prob.		Total Claims		Total Debt	
Country	Mean	(SD)	Mean	(SD)	Mean	(SD)
All	0.9870	(0.024)	0.42	(0.31)	0.94	(0.32)
AT	0.9927	(0.008)	0.32	(0.06)	0.93	(0.11)
BE	0.9912	(0.011)	0.76	(0.22)	1.30	(0.16)
DE	0.9964	(0.004)	0.30	(0.04)	0.94	(0.13)
ES	0.9855	(0.017)	0.19	(0.04)	0.58	(0.15)
FI	0.9942	(0.002)	0.07	(0.01)	0.72	(0.06)
FR	0.9946	(0.006)	0.48	(0.17)	1.03	(0.18)
GB	0.9906	(0.005)	0.34	(0.06)	0.79	(0.16)
GR	0.9675	(0.048)	0.02	(0.01)	1.31	(0.25)
IE	0.9701	(0.038)	1.04	(0.11)	0.79	(0.38)
IT	0.9865	(0.015)	0.19	(0.08)	1.44	(0.16)
NL	0.9955	(0.004)	0.86	(0.08)	0.77	(0.14)
PT	0.9751	(0.040)	0.20	(0.06)	0.84	(0.18)
SE	0.9950	(0.004)	0.43	(0.07)	0.61	(0.05)

Notes: Sample averages and standard deviations of the listed variables are given for the entire panel of countries ("All") and then separately for each country. Solvency probabilities are risk-neutral probabilities derived from 5-year CDS contracts. Total claims are the sum of bilateral financial claims ($\sum_{j\neq i} c_{ijt}$), constructed from BIS and IMF data. Total debt is total external debt from the IMF. Claims and debt are normalized by each country's 2004 GDP. See Section 4 and Appendix B for further details.

Table 3: **Linear Regressions for Solvency Probabilities**

Variable	(1)	(2)	(3)	(4)
Debtor solvencies	0.040*	0.026*	0.029*	0.018*
$(L \cdot p)$	(0.011)	(0.009)	(0.011)	(0.009)
Own debt	- - -	-0.051*	- - -	-0.050*
(D)		(0.009)		(0.010)
GDP level	- - -	0.211*	- - -	0.354*
(Y)		(0.032)		(0.055)
GDP growth, common	- - -	0.001	- - -	0.001
(ΔY^c)		(0.001)		(0.003)
GDP growth, residual	- - -	0.003*	- - -	0.003*
(ΔY^r)		(0.001)		(0.001)
Time control	t	t	1_t	1_t
R^2	0.484	0.640	0.552	0.738
N	293	293	293	293

Notes: The dependent variable is the solvency probability for country i in period t: p_{it}. Each column is a separate regression. All regressions include country fixed effects. See equation (7) for the complete specification. Standard errors are in parentheses; * p-value < 0.05.

Table 4: **Estimated Parameters and Marginal Effects in the Network Model**

Param.	Value	Marg. Eff.	Variable
A. Fixed recovery rate			
γ	1.0547	0.0289	Payments received (R)
α	0.1680	0.0046	Total foreign debt (D)
β_1	4.2926	0.1176	GDP level (Y)
β_2	0.0848	0.0023	GDP growth, common (ΔY^c)
β_3	0.0338	0.0009	GDP growth, residual (ΔY^r)
B. Proportional recovery rate			
γ	1.0546	0.0290	Payments received (R)
α	0.1680	0.0046	Total foreign debt (D)
β_1	4.2925	0.1180	GDP level (Y)
β_2	0.0848	0.0023	GDP growth, common (ΔY^c)
β_3	0.0337	0.0009	GDP growth, residual (ΔY^r)

Notes: The table shows estimates of the parameters in equation (6) using the fixed recovery function (2) (panel A) or the proportional recovery function (3) (panel B), and marginal effects of the associated variables. Note that α enters the model negatively. Marginal effects are computed as the average of the marginal effects for each observation. The equilibrium payments received (R), foreign debt (D), and GDP level (Y) are normalized by the country's 2004 GDP.

Table 5: **Proportional Effects on Borrowing Costs in 2011-Q1**

Country	Pct of Borrowing Cost Due to Contagion
AT	3.02
BE	2.09
DE	2.81
ES	0.93
FI	0.00
FR	5.44
GB	1.64
GR	0.04
IE	1.73
IT	0.33
NL	3.05
PT	3.20
SE	0.85

Notes: The table presents results from a counterfactual simulation in which spillover effects are eliminated, by setting the recovery rate to 100% ($\delta = 1$). The reported values are percentage changes in borrowing costs (measured using default probabilities) in the first quarter of 2011, relative to the baseline predictions from the model.

Appendices

A Analysis of Potential Biases

Here we provide a discussion of the biases that might arise if certain key assumptions in our empirical model were violated. We consider four potential issues: contemporaneous correlations in the financial shocks among countries, correlations over time and the endogeneity of financial linkages, endogenous default decisions, and internal amplification mechanisms. Our focus is on the bias in the estimate of γ, the parameter that governs the magnitude of spillovers from a default. The informal analyses below suggest that the likely biases are upward, however our preliminary regression estimates, which include country and quarter fixed effects, indicate that any such biases do not affect our results qualitatively.

Correlations in financial shocks: As noted in Section 3.1, it may be possible to allow for a contemporaneous correlation in the financial shocks among countries, but for statistical precision we estimate the model under the assumption of independence. If the shocks X_{it} and X_{jt} were in fact correlated, the observed relationship between p_{it} and p_{jt} would naturally reflect this correlation in addition to the true effects of repayments between countries i and j. Because the repayments are an increasing function of the shocks (e.g., X_{it} enters positively in (5)), the estimate of γ would be biased in the same direction as the correlation in these shocks. Hence, a positive correlation in the financial shocks among countries would result in an upward bias in the estimate of γ.

Based on existing results for linear network models (e.g., Bramoullé, Djebbari, and Fortin, 2009; Lee, Liu, and Lin, 2010), we speculate that a contemporaneous correlation in the shocks would be separately identifiable from the endogenous spillover effect in our nonlinear model as well. (This is not the case for nonlinear models with complete graphs, where the spillover effects are uniform within groups—see Krauth (2006).) Specifically, we believe that if ϵ_{it} were decomposed into a common and idiosyncratic component, such as u_t and v_{it}, the variance of u could be identified separately from the parameters in (6). This follows from similar logic as the identification of nonlinear panel models with random effects. All the variables in (6) exhibit variation across countries at a point in time, including the claims that influence R_{it}. Hence the distribution of a common shock should be identifiable. However, to our knowledge, such results on identification with correlated unobservables are not currently available for our class of nonlinear network models. Brock and Durlauf (2007) discuss various conditions to achieve partial identification in nonlinear models with *complete graphs*, and Lee, Li, and Lin (2014) consider a nonlinear model with *incomplete information*, but these do not directly apply to our class.

Endogeneity of financial linkages: A bias would arise if our constructed linkages, c_{ijt}, were endogenous; i.e., if they were correlated with the shocks, X_{it}, for any reason. Because the financial claims are established in the previous period, this would require that the shocks be correlated over time.[36] In that case, banks in country i might reduce their holdings of debt from sovereign j if a low value of the shock $X_{j,t-1}$ is realized, because that predicts a low value of X_{jt}, and hence a higher probability of default for country j in period t. This process would yield a positive correlation between the claims of i on j (c_{ijt}) and the current shock for j (X_{jt}), which determines the solvency of j (s_{jt}) and thereby affects the repayments to i (R_{it}). Then, because the predicted repayments in our model, given by equation (6), do not account for the lagged shocks ($X_{j,t-1}$), there would be an error term for the difference between the correct predictions and our predictions (heuristically, $\mathrm{E}[R_{it}|X_{j,t-1},\ldots] - \mathrm{E}[R_{it}|\ldots]$). This error term would have a positive correlation with our predicted values of R_{it} because of the positive correlation between c_{ijt} and X_{jt} described above. Therefore, this positive correlation would generate an upward bias in the estimate of γ.

Endogenous default: If default were endogenous, the decision rule for each country would be a function of the state variables known at the time of the payment equilibrium: the network-wide matrix and vectors $\mathbf{C}_t$, D_t, Y_t, and X_t. Conceptually, we could incorporate such a decision rule into the solvency condition with a policy function $\pi_i(\mathbf{C}_t, D_t, Y_t, X_t)$ that adjusts the threshold for default away from the fixed values given by the parameters $\bar{\pi}_i + \bar{\pi}_t$ in our model (bars added here for clarity).[37] Then $-\pi_i(\mathbf{C}_t, D_t, Y_t, X_t)$ would appear in (6) as an error term, and so the question is how it would be correlated with the predicted repayments based on the exogenous default rule in our model. If we assume that a country does not receive payments on its claims when it defaults and goes into autarky, then the relative value of default should be decreasing in the (true) equilibrium repayments. This suggests that the threshold $\pi_i(\mathbf{C}_t, D_t, Y_t, X_t)$ would be decreasing in R_{it} (as derived from the state variables and the other decision rules $\pi_j(\cdot)$), making default less likely when R_{it} is larger. This relationship also holds for our predicted repayments based on an exogenous default rule. Hence there would be a positive correlation between the R_{it} we use and this error term, $-\pi_i(\cdot)$, which would generate an upward bias in the estimate of γ.

[36] As noted earlier in the paper, none of the empirical papers that estimate structural models of spillovers in financial networks allow for shocks that are correlated over time.

[37] This assumes the putative equilibrium strategies would have a single crossing property in X_{it}.

Amplification mechanisms: To reflect internal amplification mechanisms with potentially different impacts across countries and over time, we could add parameters that vary the effect of repayments over i and t, as in $(\bar{\gamma} + \gamma_{it})R_{it}$ (bar added for clarity). Then $\gamma_{it}R_{it}$ would be an error term in (6), and so the question is whether these deviations would be systematically positive or negative when our predicted R_{it} is relatively higher or lower. If we suppose that banks are more leveraged on average when they have larger debt holdings, then the sensitivity to losses (γ_{it}) should be greater when the total holdings ($\sum_{j\neq i} c_{ijt}$) are larger. Accordingly, there would be a positive correlation between our predicted repayments and the error term $\gamma_{it}R_{it}$, which would generate an upward bias in the estimate of $\bar{\gamma}$.

B Construction of Variables

Here we describe how we construct the network of financial linkages used in our analysis and how we transform the observed spreads on 5-year CDS contracts into risk-neutral solvency probabilities.

Constructing the Network of Financial Linkages

The BIS reports asset holdings of financial institutions headquartered in one country according to the country of the ultimate counterparty, at a quarterly frequency. This measure includes all financial assets, not just sovereign debt, that are held by the financial sector. Specifically, we use the BIS consolidated international banking statistics on an ultimate risk basis. The data provided on an ultimate risk basis are more appropriate for our purposes than the data on an immediate borrower basis. For example, in a 2010 Quarterly Review issued by the BIS (Avdjiev, Upper, and von Kleist (2010)), it is stated:

> "*The BIS consolidated international banking statistics on an ultimate risk basis are the most appropriate source for measuring the aggregate exposures of a banking system to a given country. Unlike the BIS consolidated international banking statistics on an immediate borrower basis, they are adjusted for net risk transfers. For example, suppose that a Swedish bank extends a loan to a company based in Mexico and the loan is guaranteed by a US bank. On an immediate borrower basis, the loan would be considered a claim of a Swedish bank on Mexico, as the immediate borrower resides in Mexico. On an ultimate risk basis, however, the loan would be regarded as a claim of a Swedish bank on the United States since that is where the ultimate risk resides.*"

We define b_{ijt} as the total value of financial claims held by banks in country i on entities in country j at date t, as reported by the BIS on an ultimate risk basis. We then define an adjusted claims measure, the "BIS weight," as follows:

$$\text{BISwgt}_{ijt} = \frac{b_{ijt}}{\sum_{k=1}^{N} b_{kjt} + \sum_{k=N+1}^{\#\text{BIS}} b_{kjt}}.$$

The data from the BIS include other countries beyond the 13 European countries in our sample. As indicated by the second term in the denominator above, we include those other countries to compute our adjusted claims measure (#BIS is the total number of BIS reporting countries). Thus the BIS weight gives the proportion of all claims on entities in country j (held by banks in any BIS reporting country) that are held by banks in country i.

Then to construct the network of sovereign debt holdings, we use this adjusted claims measure to weight each sovereign's total external sovereign debt, and thereby allocate it among the countries in our sample (and outside the sample). The measure of a sovereign j's debt that is externally held, D_{jt}^{foreign}, comes from data provided by the IMF. Finally, we compute the measure of sovereign i's aggregate claims on sovereign j at date t as follows:

$$c_{ijt} = \text{BISwgt}_{ijt} \cdot D_{jt}^{\text{foreign}} / Y_{i,2004}.$$

This includes the normalization for sovereign i's 2004 GDP ($Y_{i,2004}$). This is the measure of financial linkages used to estimate the model.

For a simple, concrete example, suppose the BIS data report that 40% of the total financial claims issued by entities located in country A are held by institutions located in country B and 60% are held by institutions located in country C. Additionally, suppose the IMF reports that of the debt issued by the government of sovereign A, $200 billion is held by foreign creditors. Our construction using the BIS weights would then assume that $80 billion of sovereign A's debt is held by country B and $120 billion is held by country C. Finally, if country B's 2004 GDP was $400 billion and country C's was $1.2 trillion, the normalized measures of their claims on sovereign A would be 0.20 and 0.10 respectively.

Imputing Risk-Neutral Solvency Probabilities from CDS Spreads

We use spreads on 5-year CDS contracts to impute risk-neutral solvency probabilities for the sovereigns in our sample. CDS contracts provide insurance against a credit event of a reference entity, which in our case is a sovereign. The purchaser of protection obtains the right to sell bonds issued by the underlying entity, at their face value, to the seller of the CDS contract. In exchange, the purchaser of the CDS contract makes periodic payments to the seller until the occurrence of a credit event by the reference entity or the maturity of the contract.

Let P_t denote the risk-neutral probability of the underlying sovereign remaining solvent (without a credit event) up to date t (i.e., the cumulative survival probability). The present value of the contingent payments received by the buyer of protection can be expressed as

$$\sum_{t=1}^{T}(1-\delta)(P_{t-1}-P_t)d_t$$

where δ is an assumed recovery rate on the sovereign bonds and d_t is the risk-free discount factor. Similarly, the present value of the fixed payments made by the buyer of protection can be expressed as

$$\sum_{t=1}^{T} d_t P_{t-1} S$$

where S is the fixed payment rate (commonly referred to as the CDS spread).

We obtain a time series of CDS spreads for each of the 13 sovereigns in our sample using data from CMA. The data are reported at a monthly frequency, and we take the simple average within each quarter to construct the quarterly series. The most complete data are available for 5-year CDS contracts denominated in US dollars. Accordingly, we assume a constant hazard rate in order to impute quarterly, risk-neutral solvency probabilities from these CDS spreads. Under this assumption and letting $(1-\bar{p})$ denote the per period probability of default, so that $\bar{p}$ is the per period solvency probability, we have $P_t = \bar{p}^t$. By no arbitrage, the present value of the fixed and contingent payments must be equal, which yields

$$\sum_{t=1}^{T}(1-\delta)[\bar{p}^{t-1}-\bar{p}^t]d_t = \sum_{t=1}^{T} d_t \bar{p}^{t-1} S$$

Following the literature and estimates from a sample of historical sovereign defaults, we assume a recovery rate of $\delta = 0.4$. We compute the discount factor, d_t, using empirical yields on US Treasuries. Thus, given data on the CDS spread, S, we can impute the risk-neutral quarterly solvency probability, $\bar{p}$, using the equation above. Note that this can be repeated for each sovereign at each date in our sample, providing the panel of implied, risk-neutral quarterly solvency probabilities, p_{it}, required for our estimation.

Table A-1: Aggregate Bank Holdings of Foreign Sovereign Debt in 2011-Q1

Bank HQ Country	Holdings from Each Sovereign as Percentage of Own Country's 2004 GDP												
	AT	BE	DE	ES	FI	FR	GB	GR	IE	IT	NL	PT	SE
AT	- -	0.6	18.6	1.5	0.2	4.4	1.3	3.0	0.3	11.1	2.6	0.4	0.3
BE	0.6	- -	5.0	3.5	0.1	18.6	2.0	1.6	2.2	9.4	2.9	0.8	0.2
DE	2.7	0.8	- -	3.7	0.3	9.2	3.8	2.4	1.3	8.3	2.7	1.1	0.5
ES	0.3	0.2	4.4	- -	0.1	2.9	6.9	0.3	0.3	3.9	0.7	5.3	0.1
FI	0.2	0.1	1.4	0.5	- -	2.8	0.3	0.0	0.1	0.6	0.6	0.1	0.8
FR	0.8	7.0	12.6	4.2	0.3	- -	3.0	8.0	0.5	28.4	3.0	1.1	0.3
GB	0.3	0.9	8.5	2.7	0.2	16.8	- -	1.9	2.0	4.4	2.7	0.9	0.4
GR	0.0	0.0	1.6	0.1	0.0	0.7	0.9	- -	0.1	0.3	0.6	0.0	0.0
IE	1.6	2.0	33.8	4.9	0.0	11.6	19.0	1.4	- -	10.5	1.6	1.1	0.6
IT	5.1	0.2	14.3	1.0	0.0	2.9	0.5	0.7	0.2	- -	0.6	0.2	0.1
NL	1.2	12.0	29.9	7.1	0.4	17.0	3.5	2.4	1.0	11.1	- -	0.8	0.4
PT	0.1	0.1	1.8	6.7	0.0	3.9	0.5	12.2	0.7	1.8	2.4	- -	0.0
SE	0.3	0.6	21.8	0.7	21.2	5.4	4.0	0.1	0.2	0.5	1.4	0.1	- -

Notes: The table reports our constructed measure of the aggregate holdings of banks headquartered in each country (listed by row) of sovereign debt from each other country (listed by column), in the first quarter of 2011. The data are taken from the BIS and IMF and transformed as described in Appendix B. Gross nominal values are normalized by the 2004 GDP of the home country, and are listed above as percentages. These are the amounts represented in the network graph in Figure 1.

Table A-2: Summary Statistics for GDP Variables

	GDP Level		GDP Growth			
			Commmon		Residual	
Country	Mean	(SD)	Mean	(SD)	Mean	(SD)
All	1.08	(0.03)	-0.46	(0.91)	0.00	(0.71)
AT	1.09	(0.03)	-0.16	(0.79)	0.01	(0.45)
BE	1.08	(0.02)	-0.24	(0.71)	0.02	(0.31)
DE	1.10	(0.03)	0.07	(1.09)	0.05	(0.46)
ES	1.11	(0.02)	-0.48	(0.63)	-0.04	(0.41)
FI	1.07	(0.01)	-0.26	(0.51)	0.33	(1.07)
FR	1.07	(0.02)	-0.33	(0.58)	0.02	(0.25)
GB	1.07	(0.03)	-0.90	(0.95)	-0.10	(0.54)
GR	1.08	(0.04)	-0.96	(0.38)	0.02	(1.20)
IE	1.10	(0.04)	-1.31	(1.30)	-0.10	(1.63)
IT	1.03	(0.03)	-0.39	(0.88)	-0.01	(0.41)
NL	1.11	(0.03)	-0.32	(0.74)	0.00	(0.46)
PT	1.09	(0.02)	-0.60	(0.63)	-0.05	(0.55)
SE	1.07	(0.04)	-0.26	(1.15)	0.03	(0.73)

Notes: Sample averages and standard deviations of the listed variables are given for the entire panel of countries ("All") and then separately for each country. GDP levels are normalized by each country's 2004 GDP. GDP growth is decomposed into a common component and a residual via a principal components analysis described in Section 3. The common component of GDP growth is also detrended by subtracting the average quarterly growth rate for each country over the period from 1995 to 2004.

Table A-3: Comparison with Alternative Measures of Claims on Sovereign Debt in 2010-Q4

	Correlations with Our Measure	
Bank HQ Country	EBA 2011 Stress Test	BIS Public Sector Claims
All	0.883	0.907
AT	0.955	
BE	0.450	0.573
DE	0.832	0.813
ES	0.816	0.865
FI	0.677	
FR	0.901	0.947
GB	0.947	0.848
GR	0.812	
IE	0.491	
IT	0.964	0.983
NL	0.964	
PT	0.732	
SE	0.979	

Notes: The table shows the correlations between our constructed measure of aggregate bank holdings of sovereign debt from each other country, and two other measures of these holdings available from the BIS and EBA. The comparison is made using data for 2010-Q4 based on the timing of an EBA stress test. The first row shows the correlations across all observations and subsequent rows show the correlations for each country in the holdings of sovereign debt from each other country. See Section 4.1 for further details.