NBER WORKING PAPER SERIES

CONTAGION IN THE EUROPEAN SOVEREIGN DEBT CRISIS

Brent Glover
Seth Richards-Shubik

Working Paper 20567
http://www.nber.org/papers/w20567

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 2014

We thank Ana Babus, Irina Balteanu, Steven N. Durlauf, Matthew Elliott, Michael Gofman, Christian Julliard, 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, and the Financial and Economics Networks Conference at the University of Wisconsin for helpful comments. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Contagion in the European Sovereign Debt Crisis
Brent Glover and Seth Richards-Shubik
NBER Working Paper No. 20567
October 2014
JEL No. D85,F34,F36,G01,L14

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 model, based on the recent theoretical literature on contagion in financial networks, and estimate it with data on sovereign credit default swap spreads and the detailed structure of financial linkages among thirteen European sovereigns from 2005 to 2011. Simulations from the estimated model show that a sovereign default generates only small spillovers to other sovereigns. These results imply that credit markets do not demand a significant premium for the interconnectedness of sovereign debt in Europe.

Brent Glover
Tepper School of Business
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213
gloverb@andrew.cmu.edu

Seth Richards-Shubik
H. John Heinz III College
School of Public Policy and Management
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213
and NBER
sethrs@andrew.cmu.edu

"Nothing less than the future of Europe is at stake. ... Immediate help is needed to ensure the financial stability of the eurozone. *This must be done to avoid a chain-reaction to the European and international financial system, and contagion to other eurozone states.* There is no alternative."

– Angela Merkel to the Bundestag, May 5, 2010, in support of a €110 billion bailout package for Greece[1]

"But above all, the most important thing is the construction of the firewall of this European fund [the EFSF] *to prevent contagion taking place elsewhere.*"

– David Cameron to Parliament, Oct. 26, 2011, prior to a European summit where leaders agreed to expand the lending capacity of the European Financial Stability Facility (EFSF) to €1 trillion[2]

1 Introduction

Policymakers and the media have raised major concerns about financial contagion in the ongoing sovereign debt crisis in Europe. The basic fear is that the default of one sovereign would have spillover effects that result in subsequent defaults or increased borrowing costs for other sovereigns. Given the interconnectedness of global financial markets, the potential for such spillovers across countries seems self-evident, and as the above quotes show, the risk of contagion has been used as a justification for massive bailout efforts.[3] Since the crisis began in 2008, almost €500 billion have been provided to eight sovereigns from various European funding mechanisms and the International Monetary Fund (IMF).

The extent to which the purported financial spillovers among sovereigns are economically important, however, remains an open question. Existing assessments of contagion in the European debt crisis are limited and primarily descriptive in nature, mainly examining

[1]Source: `www.telegraph.co.uk/finance/financialcrisis/7683271/Merkel-plea-to-save-Europe-as-contagion-hits-Iberia.html`; emphasis added here and below.

[2]Source: `www.theguardian.com/politics/2011/oct/26/eu-debt-crisis-contagion-cameron-pmqs`

[3]The Vice-President of the European Central Bank, Vítor Constâncio, made this point even more explicitly in an address in October 2011: "Contagion is crucial for policy-making. This is in particular the case because it usually constitutes an externality, in the economic meaning of the term. ... There is a market failure that policy should try to address" (Constancio 2012).

correlations in bond yields and credit default swap (CDS) spreads. The true magnitude of the potential spillovers from a sovereign default is difficult to measure, in part due to data limitations, but also because credit outcomes in an interconnected network are jointly determined and therefore mutually endogenous. Nevertheless, quantifying the risk of such spillovers is necessary to assess whether the benefits of a sovereign bailout outweigh the costs.

In this paper 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 model based on the recent theoretical literature on contagion in financial networks (e.g., Acemoglu, Ozdaglar, and Tahbaz-Salehi forthcoming; Elliott, Golub, and Jackson forthcoming; Glasserman and Young forthcoming), and estimate it with data on sovereign CDS spreads and the detailed structure of financial linkages among thirteen European sovereigns from 2005 to 2011. In this framework the spillovers from a default occur via direct losses to assets (e.g., loans or bonds) held by creditors, in what could be referred to as a "balance sheet" mechanism for contagion. The equilibrium solution for solvency and repayments among the sovereigns in the network expressly accounts for the joint determination of asset values based on this mechanism.

We use data from the Bank for International Settlements (BIS) and IMF to construct an empirical network of financial linkages among the sovereigns in our sample for each quarter. Additionally, we impute risk-neutral market expectations for the sovereigns' default probabilities in each quarter using the spreads on their 5-year CDS. Combining these series with data on the sovereigns' GDP, we estimate the parameters of our network model of contagion.

After estimating the model with these data, we use a series of simulations to quantify the potential spillovers from a sovereign default. Specifically, we consider counterfactuals where we simulate the default of one sovereign and compute the predicted change in the risk-neutral default probabilities of the other sovereigns in the network. From this we develop a novel measure of the contagion risk posed by each country. Our measure has a relationship with centrality measures commonly used in network analysis, but it also has a more direct economic interpretation as the *expected spillover losses per dollar of debt*. With this measure we show how contagion risk has risen over time, and we arrive at a potentially surprising result for the country with the greatest potential for contagion per unit of debt: Austria. Much of Austria's debt is held by Italy, a financially vulnerable sovereign with substantial external debt. Thus, while Austria's default probability is low, the model predicts relatively high spillover losses in the event of an Austrian default.

Yet while we document an increase in contagion risk over time and find informative heterogeneity across sovereigns, the magnitude of the predicted losses resulting from these spillovers are, for the most part, economically small. On average the predicted losses due to contagion account for only one percent of the total expected losses implied by the sovereign CDS spreads in our sample. These results indicate that credit markets do not demand a significant premium for spillover losses resulting from a sovereign default. In other words, the financial interconnectedness of European sovereign debt holdings does not appear to have an economically significant effect on the sovereigns' cost of borrowing.

Given these results, it is particularly important to consider factors that might bias our estimates toward a finding of no contagion. For this it is useful that our empirical model is directly related to a class of models found in the microeconometric literature on social interactions. The conditions for the identification of endogenous spillover effects are well understood in that literature (Manski 1993; Blume et al. 2011). The basic challenges of simultaneity and the "reflection" problem, which arise in models with interconnected agents, are resolved by having data on individual network linkages, as we do (Bramoullé, Djebbari, and Fortin 2009). However, another possible source of bias comes from the potential endogeneity of these linkages. Here we follow the theoretical literature on financial contagion, and the empirical work related to this literature (discussed below), by treating the network of financial linkages as exogenous. Hence any unobserved factors that determine both financial linkages and credit risk would bias our results. We consider this issue in detail and show that any bias is likely to be upward, so it would not affect our overall conclusion that the potential spillovers from the balance sheet mechanism are relatively small.

Essentially, the basic empirical fact that drives our results is that the differential financial linkages among countries explain relatively little of the differential comovements in sovereign credit risk. Any model where the transmission of risk is in some way related to the aggregate financial linkages among countries would yield a small estimate for the spillover effects because of this feature of the data. Other transmission mechanisms exist that do not involve financial linkages, such as changes in investor risk preferences or changes in beliefs regarding the likelihood of a particular event, such as a sovereign default.[4] These mechanisms may deserve further attention in future work, both theoretically and empirically. While our results suggest that credit markets perceived small spillover effects from one channel for contagion, it is possible that other channels might have larger economic effects.

Existing assessments of contagion in the European debt crisis do not account for the joint determination of credit risk in equilbrium, and they typically do not use data on the financial linkages among countries. One approach analyzes comovements in sovereign bond yields or CDS spreads across countries. The correlations that remain after filtering out long-run trends are attributed to contagion (see Constancio (2012) for a summary).[5] Another approach

[4]See, for example, Benzoni et al. (2012) and Kodres and Pritsker (2002).

[5]Recent work in the finance literature has also used sovereign CDS spreads to examine the time series properties and comovement of sovereign credit risk more generally. Longstaff et al. (2011) find that global factors explain a large portion of the common variation in sovereign CDS spreads. Dieckmann and Plank (2012) study CDS spreads for a sample of developed countries and find strong comovement that increased

estimates regression models of sovereign credit risk that include both the macroeconomic fundamentals of the individual country along with some measure of the average credit risk among other European sovereigns. The coefficient on the latter is considered to represent the effect of contagion (e.g., Arghyrou and Kontonikas 2012; Beirne and Fratzscher 2013).

Our empirical approach is more closely related to a handful of recent papers that estimate structural models of spillovers in interbank networks (Cohen-Cole, Patacchini, and Zenou 2011; Denbee et al. 2014; Bonaldi, Hortaçsu, and Kastl 2014).[6] We follow the key assumptions made in these papers, specifically that financial linkages established in a previous period are exogenous and that unobserved shocks are independent over time. These assumptions make it possible to apply a structural network model to the available data, as they allow each time period to be treated separately.

In addition, two recent papers on the European debt crisis provide rationales to apply a model such as ours to this context. Bolton and Jeanne (2011) develop a model of contagion based on bank holdings of foreign sovereign debt, which is used as collateral. This is similar to the balance sheet mechanism as expressed in our model, and clearly the extent of the spillovers between two countries would relate to the amount of these holdings. Our empirical network is accordingly constructed to represent the total bank holdings in one country of the sovereign debt from another. Acharya, Drechsler, and Schnabl (forthcoming) provide theory and evidence for a feedback loop in credit risk between banks and the central government within each country.[7] This shows how a chain of contagion could emerge among multiple

significantly following the financial crisis beginning in 2008. Ang and Longstaff (2013) similarly find strong commonality in Eurozone sovereign CDS spreads. These authors do not attempt to isolate the role of contagion in these comovements, however.

[6] There is also a literature that conducts simulation studies with calibrated models of interbank networks. See Gofman (2013), for example, and Upper (2011) for a survey.

[7] Their empirical analysis of domestic spillovers between banks and the central government also includes bank exposures to credit risk from foreign sovereigns, using BIS data on aggregate financial linkages as weights. However the identification of the spillover from foreign sovereigns is not the focus of their analysis. In particular, their statistical model does not account for a possible feedback from banks to sovereigns.

sovereigns, even though the central governments themselves do not function explicitly as financial intermediaries. In addition to these papers on the European crisis, Elliott, Golub, and Jackson (forthcoming) specifically suggest that their model could be applied to a network of countries, and they use BIS data on aggregate financial linkages to provide a quantitative illustration.

Finally, we should note how our paper relates to the international macroeconomics literature on sovereign debt. This literature, beginning with the seminal contribution of Eaton and Gersovitz (1981), typically studies a sovereign's optimal borrowing and default decisions in dynamic, incomplete markets models. Aguiar and Gopinath (2006) use a quantitative model of debt and default in a small open economy framework and find an important role for a stochastic trend in growth rates for emerging economies. Arellano (2008) uses a similar model framework to study the interaction of sovereign credit risk with output fluctuations and interest rates. This literature generally considers the debt issuance and default decisions of a single sovereign in isolation. A notable recent exception is Arellano and Bai (2013), who develop a model in which two sovereigns borrow from a common lender, making defaults and renegotiations interdependent. Our work is distinct and complementary to this literature, in that we consider a network of countries which are both borrowers and lenders, and our focus is mainly an empirical assessment of the risk of contagion in the recent sovereign debt crisis in Europe.

The remainder of the paper is organized as follows. Section 2 presents the framework for our network model of contagion among sovereigns. We then develop an empirical version of the model and discuss estimation and identification in Section 3. The data are described in Section 4, and in Section 5 we present the empirical results. First, we show the parameter estimates and model fit. We then use the estimated model to simulate counterfactual sovereign defaults and examine the contagion effects of these shocks propagating through the network. Section 6 concludes.

2 Theoretical Framework

Our framework follows a number of recent papers in the theoretical literature on contagion in financial networks.[8] 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 (Eisenberg and Noe 2001).[9] 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 "balance sheet" mechanism of contagion.

Applying this framework to our context, each country is treated as a single, aggregate financial entity, and countries are connected through their aggregate financial claims on each other.[10] Because our focus is on sovereign debt, this aggregate approach relies on the close

[8]See, for example, Gouriéroux, Héam, and Monfort (2012), Acemoglu, Ozdaglar, and Tahbaz-Salehi (forthcoming), Elliott, Golub, and Jackson (forthcoming), and Glasserman and Young (forthcoming). These papers generally build on the framework proposed in Eisenberg and Noe (2001). For a more general survey of models of contagion in financial networks see Allen and Babus (2009).

[9]The network of financial linkages is not endogenized in these models. Instead, these papers study the risk of contagion and the implications of different canonical network topologies on this risk.

[10]As noted earlier, Elliott, Golub, and Jackson (forthcoming) specifically suggest that this framework could be applied to countries. When they introduce their model they name "countries, banks, or firms" as examples of the entities in the network (p. 4), and they later provide an empirical illistration of the model using BIS data on the aggregate financial linkages among six European countries.

connection in credit risk between the central government and banks within each country (Acharya, Drechsler, and Schnabl forthcoming), and the fact that banks hold substantial amounts of foreign sovereign debt (Bolton and Jeanne 2011). Losses in the value of bank holdings of foreign debt thus impact the central government, which in turn affects its ability to make payments on its own sovereign debt.

Formally, 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.[11] In each period, sovereigns hold debt claims on each other that were established in a previous period. The face value of country i's gross, aggregate claims on country j, payable at date t, is denoted l_{ijt}. These bilateral claims are collected into a matrix L_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 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., $D_{it} \geq \sum_{j \neq i} l_{jit}$). A sovereign's output, $Y_{it} \in \mathbb{R}^+$, 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 countries 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 (l_{ijt}). Solvency is denoted with indicators s_{it}. If sovereign j is solvent in period t ($s_{jt} = 1$), then sovereign i receives the full value of its claims on j, i.e., l_{ijt}. If, on the other hand, a country defaults, its creditors receive a proportion of their claims δl_{ijt}, where $\delta \in [0, 1)$ is a fixed, exogenous recovery rate. This

[11] This is a limitation, but as we discuss in Section 3.1 and Appendix A, it does not appear to qualitatively impact the estimate of the spillovers from this mechanism for contagion. Other empirical analyses of spillovers in financial networks similarly apply static models to repeated observations on a single set of players, thereby treating each period independently (Cohen-Cole, Patacchini, and Zenou 2011; Denbee et al. 2014; Bonaldi, Hortaçsu, and Kastl 2014).

assumption of a fixed recovery rate is common in the credit risk literature, and the value we choose is consistent with historical recovery rates for sovereign defaults.[12] Given this fixed recovery rate, the contingent payment that country i receives for its claims on j in period t can then be written as $l_{ijt}[\delta + (1-\delta)s_{jt}]$. The total repayments received from other countries in the payment equilibrium in period t is thus

$$R_{it} \equiv \sum_{j \neq i} l_{ijt}[\delta + (1-\delta)s_{jt}]. \tag{1}$$

A sovereign is solvent if, with these repayments and the shocks Y_{it} and X_{it}, it has sufficient assets to pay its debts. Accordingly, the solvency of each sovereign is determined as

$$s_{it} = \mathbb{1}\left\{R_{it} + Y_{it} + X_{it} > D_{it}\right\}. \tag{2}$$

A payment equilibrium establishes a vector of repayments $(R_{it})_{i=1}^{N}$, or equivalently a vector of solvency indicators $(s_{it})_{i=1}^{N}$, that solve the system of equations defined by (2).

Depending on the values of Y_{it} and X_{it} across all countries, there may be multiple solutions to (2). Similar to the model in Elliott, Golub, and Jackson (forthcoming), this is a consequence of the discrete loss that occurs with a default. When there are multiple solutions (i.e., multiple equilibria), we follow Elliott, Golub, and Jackson (forthcoming) and select the "best-case" equilibrium in which the fewest countries default.[13] 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

[12] Ang and Longstaff (2013) assume a recovery rate of 50%, and Longstaff et al. (2011) and Pan and Singleton (2008) assume 25%. In a sample of historical sovereign debt restructurings, Sturzenegger and Zettelmeyer (2008) estimate a range of recovery rates from 30 - 75%. The discrete losses that occur with a fixed recovery rate can be motivated as a consequence of the renegotiations involved in a sovereign default. See, for example, Yue (2010) and Benjamin and Wright (2009).

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

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 (2) to determine which countries would, in fact, default; reduce the repayment amounts based on these defaults; use (2) to determine if any additional countries would default; repeat this process until no further countries would default.[14]

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 (l_{ijt}) and total debts (D_{it}), which were established in the previous period.
1. Output (Y_{it}) and financial shocks (X_{it}) are realized.
2. Solvency (s_{it}) is 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), Denbee et al.

[14] Elliott, Golub, and Jackson (forthcoming) and Eisenberg and Noe (2001) use similar algorithms.

(2014), and Bonaldi, Hortaçsu, and Kastl (2014), which similarly estimate structural models of spillovers in financial networks. All of 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. The network of financial linkages can then be considered exogenous if actions in a previous period define the network, as above or in Denbee et al. (2014), 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 have chosen instead to follow the above papers in the network literature and estimate a static model of the payment equilibrium, to serve as a starting point to assess the quantitative importance of one proposed mechanism for contagion.

3 Empirical Approach

Our goal is to estimate an empirical version of the solvency condition given in (2), which can then be used to quantify the potential spillovers from a sovereign default that arise from the balance sheet mechanism of 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 be solvent. In particular, we use the observed spreads on sovereign CDS contracts to impute a sovereign's risk-neutral default probability.

We suppose that the CDS spreads at the end of period $t-1$ reflect the market's assessment of the 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 $t-1$ (i.e., in step 4 of period $t-1$). We denote these conditional expectations p_{it}, and they are formally defined as follows:

$$p_{it} \equiv \mathrm{E}\left[s_{it} | L_t, (D_{jt}, Y_{j,t-1}, X_{j,t-1})_{j=1}^{N}\right]. \tag{3}$$

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 (i.e., the vector of solvency indicators, $(s_{jt})_{j=1}^{N})$ over this distribution.

To adapt the generalized solvency condition in (2) to work with our 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 (l) 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.[15] Thus the empirical version of the solvency condition is specified as

$$s_{it} = \mathbb{1}\left\{\gamma R_{it} - \alpha D_{it} + \beta Y_{it} + X_{it} > \pi_i + \pi_t\right\}. \tag{4}$$

The parameters γ and α express the proportions of claims and debts that are payable each period, and β gives the proportion of aggregate output that is available to the central gov-

[15]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.

ernment 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 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 (4) over their joint distribution and thereby compute the expectations in (3). 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 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} is assumed to be normal with variance σ^2_Y. Thus, the forecasted distribution of aggregate output for sovereign i in period t is

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

These are 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 some baseline level of aggregate output. Thus, we effectively allow for larger shocks, on average, in countries with larger economies.[16] Otherwise, the output and financial shocks are assumed to be independent across countries and over time, which follows Cohen-Cole, Patacchini, and Zenou (2011), Denbee et al. (2014), and Bonaldi, Hortaçsu, and Kastl (2014).[17]

Applying these specifications, the network-wide vector of conditional expectations in (3), which we refer to as the *solvency probabilities*, can be expressed as follows:

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

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 $(1\{\dots\}_{i=1}^N)$ in the integral gives the vector of solvency indicators $((s_{it})_{i=1}^N)$ as a function of the vectors of observables and shocks. The interdependencies across countries arise because the indicators s_{jt} are embedded in each R_{it}. 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. Also, because β is not separately identified from β_1, β_2, and β_3, we set $\beta = 1$. Consequently, the parameters β_1, β_2, β_3 are interpreted as the combination of the forecast for future output and the relationship between output and solvency. Finally, we use a simple linear trend to capture any changes in the default threshold

[16] This assumption also appears in the theoretical literature we draw from (e.g., Glasserman and Young forthcoming).

[17] Appendix A considers the biases that could arise if these independence assumptions are violated. There we show, among other things, that that a positive correlation in the shocks among countries could only result in an upward bias in the estimate of γ, while our main concern is with a downward bias.

over time, so that π_t is specified as θt.[18] This yields the ultimate specification that we take to the data:

$$(p_{it})_{i=1}^{N} = \int 1\left\{\gamma R_{it} - \alpha D_{it} + \beta_1 Y_{i,t-1} + \beta_2 \Delta Y^{c}_{i,t-1} + \beta_3 \Delta Y^{r}_{i,t-1} + \epsilon_{it} > \pi_i + \theta t\right\}_{i=1}^{N} \cdot \prod_{i=1}^{N} \phi(\epsilon_{it}) \mathrm{d}\epsilon_{it} \qquad (5)$$

The integral is computed via simulation.[19] For each vector of pseudo-random draws of $(\epsilon_{it})_{i=1}^{N}$, we solve the system of equations defined by (4) for the vector of solvency indicators. (When there are multiple solutions we select the best-case equilibrium, as described in Section 2). The average of these vectors of indicators across all draws provides an approximation of the vector of solvency probabilities above.

We estimate the parameters in (5) 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.[20] The identification of the model is discussed next.

3.1 Identification

To consider identification, our empirical model can be understood within a class of models used to study social and spatial interactions. The class consists of static equilibrium models where best responses are nonlinear functions of the realized actions of other players (i.e., these models are based on simultaneous-move games of complete information). In our case it is

[18] 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 (Neyman and Scott 1948).

[19] We use antithetic acceleration to improve the precision of the simulator (Stern 1997).

[20] The source of the econometric error between the predicted probabilities and the empirical probabilities is left unspecified. However to provide some intuition for our estimation procedure, we can compare (5) with a generalized linear model. If (5) were a single-equation, non-equilibrium model, it could be treated as a GLM where the econometric error yields a quasi-binomial distribution of the observed solvency probabilities. Maximum likelihood estimation would be accomplished with iteratively reweighted least squares (a Fisher scoring algorithm). Our estimation procedure does not weight the observations as they would be in that approach, but those weights would favor the same observations (i.e., those with relatively low empirical solvency probabilities) that drive our estimates.

the solvency outcomes in (4) 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 within a network. Also both employ selection rules when multiple equilibria are present, as do we—in our case, motivated by the theoretical literature (e.g., Elliott, Golub, and Jackson forthcoming).

The results in Krauth (2006) show clearly that a model such as ours is semi-parametrically identified under the assumption that the shocks (ϵ_{it}) are independent across countries and over time. Furthermore we speculate that the rich network data available in our context would make it possible to identify a contemporaneous correlation in these shocks. The main difference between our model and those in Krauth (2006) and Soetevent and Kooreman (2007) is that in our setting 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 results obtained for linear network models in Bramoullé, Djebbari, and Fortin (2009) and Lee, Liu, and Lin (2010) suggest that the additional information on network structure would serve to identify a correlation in the shocks separately from the endogenous spillover effect.[21] However, to our knowledge, the analogous results on identification are not yet available for nonlinear network models.[22] In light of this we maintain the assumption of independent shocks, but

[21]For example, 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 (5). This follows from similar logic as for the identification of a nonlinear panel model with random effects. All the variables in (5) exhibit variation across countries at a point in time, including the claims that influence R_{it}. Hence the distribution of a common unobservable should be identifiable.

[22]Brock and Durlauf (2007) establish the generic identification of nonlinear models with group-based interactions. Lee, Li, and Lin (2014) consider a nonlinear network model, but it is based on a game of incomplete information.

we consider the bias that would arise if this assumption were violated.

Appendix A provides a detailed analysis of the biases that would arise if the independence assumption or other key assumptions in our empirical model were violated. We consider four potential issues: correlations in the shocks between countries, the endogeneity of financial linkages, endogenous default decisions, and internal amplification mechanisms with different impacts across countries. Our focus is on the bias in the estimate of γ, the parameter that governs the magnitude of spillovers from a default. In each case we show that the likely bias is upward, so these departures from our model would not affect the overall conclusion that credit market perceptions of the potential spillovers from the balance sheet mechanism are relatively small.

4 Data

In this section we discuss the data used to estimate the network model described above. We combine data provided by the Bank for International Settlements (BIS) and International Monetary Fund (IMF) to construct an empirical network of bilateral, aggregate financial linkages among sovereigns. We construct this network for a set of European sovereigns for each quarter over the period from 2005-Q3 to 2011-Q3. This is combined with data on sovereign credit default swaps (CDS) and GDP to estimate the specification in (5). Table 1 lists the thirteen sovereigns included in our sample.

The central banks of BIS member countries collect data on the balance sheet composition of the banks in their jurisdiction. They aggregate these data and report to the BIS the breakdown of banks' assets according to the country of the issuer of the security. For the BIS member countries, this provides a network, at a quarterly frequency, of the claims held by the banks of one country on entities in another.[23] 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. While this gives the amount of a sovereign's debt held abroad, it is an aggregated measure that does not provide the nationalities of the various foreign creditors holding a given sovereign's debt. Thus, to construct our empirical network of sovereign debt claims, we weight the external sovereign debt amounts reported by the IMF according to the shares reported by the BIS.[24] See Appendix B for details on the variable construction.

Figure 1 gives a visual representation of our constructed network in 2011-Q1. The arrows represent financial claims that one country has on the sovereign debt of another. These amounts are normalized by the size of the economy of the creditor country, using 2004 GDP, to reflect their relative exposures. Darker arrows indicate larger proportional amounts, and claims worth less than one percent of the creditor's 2004 GDP are not shown. Nearly all countries have some amount of claims on each other, and so arrows can be bi-directional, such as 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.[25] 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 claims (relative to their

[23] We use the BIS data on consolidated claims on an ultimate risk basis.

[24] Note that this assumes the foreign sovereign debt holdings of a country's financial institutions are proportional to their total foreign asset holdings. For a 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 that the IMF reports that of the debt issued by the government of sovereign A, $50 billion is held by foreign creditors. Our construction of the network would then assume that $20 billion of sovereign A's debt is held by country B and $30 billion is held by country C. Also note that there are several BIS reporting countries that are not included in our sample, and so we are allowing for a portion of the sovereigns' debt to be held by countries outside our network.

[25] 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 (or different initializations) produce different visual representations. Nevertheless the qualitative features are reasonably stable.

own economies) on the sovereign debt of Italy (IT) and Greece (GR), respectively. This will be relevant to some of the results seen in Section 5.2.

We collect CDS spreads from Credit Market Analytics (CMA) for each of the thirteen sovereigns in our sample. All spreads are on five-year CDS contracts, referencing the sovereign entity and denominated in US dollars. We transform the CDS spreads to compute implied risk-neutral default probabilities.[26] Specifically we use the 5-year sovereign CDS spreads and the U.S. Treasury yield curve to compute the time series of quarterly solvency probabilities for each sovereign, p_{it}.[27] Last, in addition to the financial linkages and CDS spreads, we collect data on countries' GDP. We use quarterly GDP data that is annualized, seasonally adjusted, and measured in fixed PPP, taken from the OECD's Quarterly National Accounts database. Quarterly growth rates are 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 1995–2004.

In Table 2 we provide summary statistics for variables used in the estimation of the model. 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 ($\sum_{j \neq i} l_{ijt}$) 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),

[26] We follow the credit risk literature in analyzing risk-neutral default probabilities. See, for example, Ang 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, which is our ultimate object of interest.

[27] Note that this transformation of CDS spreads to risk-neutral solvency or default probabilities assumes a 40% recovery rate and a discount factor derived from the current Treasury yield. See Appendix B for details on how we impute a sovereign's risk-neutral solvency probability from its CDS spread.

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.

4.1 Assessment of the Constructed Network

Our measure of the financial linkages among countries assumes that the allocation of external sovereign debt to foreign banks is proportional to the allocation of all financial assets from a given country. We use this constructed network rather than more direct measures of claims on sovereign debt because the latter are not consistently available for the countries in our sample. However, to assess the validity of our constructed network, we compare it with other data from the BIS and from the European Banking Authority (EBA), which are available for a subsample of countries or at specific 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 has released data from its bank stress tests, which show the exposures to the sovereign debt from each country for a sample of large banks. These banks account for a large portion of the financial sector 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-1 presents the correlations between these alternative measures and our constructed measure, overall and for each country as a creditor.

The overall correlations with our measure are 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

As a final 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 illustrate the variation in the data that identifies our stuctural parameters. Thus, the coefficients should be taken simply as conditional correlations. The main specification is

$$p_{it} = a_0 + a_1 t + b \sum_{j \neq i} l_{ijt} p_{jt} + c D_{it} + d_1 Y_{it} + d_2 \Delta Y_{it}^c + d_3 \Delta Y_{it}^r + u_i + v_{it}, \tag{6}$$

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 expresses 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 (l_{ijt}). This is the same cross-moment that identifies the estimate of γ in our network model, although here the estimate of b is obviously biased from the simultaneity of p_{it} and $p_{jt}, j \neq i$.

The results of this exercise are shown in Table 3. First we estimate (6) 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. This is quite a small association, which is at least suggestive that the true spillover effects are not large. 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 reasonable 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. The overall similarity indicates that a linear time trend fits the data reasonably well and should not affect the results qualitatively, although there may be a modest upward bias in the estimate of γ in our equilibrium network model (but as noted in Section 3.1 our concern is mainly with downward biases).

When time period fixed effects are included, the financial linkages provide a crucial source of variation to estimate b because the overall correlation in solvency probabilities at a point in time would be absorbed by the fixed effects a_t. This is why we say the estimate of γ ultimately depends on the extent to which differential comovements in solvency probabilities are explained by differential financial linkages. Thus it is reassuring to see that the cross-moment expressed with b, which drives the estimate of γ, remains largely intact when time fixed effects are used in place of a linear time trend.

5 Empirical Results

5.1 Estimates and Model Fit

We now present the results from our network model, expressed in equation (5). Parameter estimates and the marginal effects of the associated variables are listed in Table 4. The key parameter is γ, which gives the effect of repayments received on financial claims on the (risk-neutral) probability that a sovereign remains solvent each quarter. The average marginal effect of repayments is 0.021, which is similar in magnitude to the coefficient on debtor solvencies in the above regressions (Table 3, column 2). As expected, however, the estimate from our equilibrium model is smaller than the naïve regression estimate because here we properly account for the joint determination of solvency in the network. The effect of a sovereign's total foreign debt load is given by α. Its marginal effect implies that an increase in the normalized debt load equal to one standard deviation of this variable (0.32, Table 2) would raise the probability of default by 25 bps on average.[28] The effects of the GDP variables are as expected.

Figure 2 plots the observed and predicted solvency probabilities to illustrate the dispersion in the data and show the model's fit. 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 both 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 0.965.

5.2 Simulations and Contagion

Using the estimated version of equation (5), we can simulate the short-run effect of the default of one sovereign on the risk of default of other sovereigns in the network. This provides one assessment of the potential for contagion based on the balance sheet mechanism. We then construct a measure of the expected spillover losses due to this increased probability of default

[28]This marginal effect is lower than the analogous coefficient in the linear regressions, but we see a more similar magnitude in a naïve nonlinear regression model. (We estimated a version of (6) as a GLM with a probit link function. These results are not shown but are available on request.)

of other sovereigns, which can be used to quantify the systemic risk from each sovereign. We also compute the aggregate amount of these spillover losses by combining the expected spillovers from each sovereign, weighted by its baseline probability of default. This provides a measure of the total potential for contagion based on the balance sheet mechanism and shows how it has evolved over time. Finally, we simulate the change in solvency probabilities if there were no spillover effects, which provides a simple illustration of the impact of this form of contagion risk on a sovereign's cost of borrowing.

To simulate the default of a given sovereign j in period t, we fix the solvency indicator for that country at zero ($s_{jt} = 0$) and recompute the solvency probabilities for all other sovereigns according to the estimated version of (5). The four panels in Figures 3 and 4 plot the increase in the (risk-neutral) default probabilities for selected sovereigns, given a default in Greece, Portugal, Italy, or Spain. These simulations should be considered separately for each period, as there are no cumulative or long-run effects expressed in the model.

There are notable differences in the impact of a default by one of these four sovereigns. Greece poses a somewhat substantial threat to Ireland and Portugal, reducing their solvency probabilities by up to 60 bps with a default (Figure 3-A).[29] Portugal also poses a small threat to Ireland (but not reciprocally to Greece), and is more of a threat to Spain in the last four quarters in our sample (Figure 3-B). The asymmetric relationship between Greece and Portugal comes from the fact that Portugal held a relatively large amount of Greek debt at that time, but not vice versa.

Figure 4-A shows that a default in Italy would increase the default probability in France by up to 60 bps. This is substantial relative to the baseline default probability in France,

[29] Official 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 in our constructed network, from 0.102 to 0.012 in the normalized measure.

which ranges from 100-200 bps at the end of our sample period. Spain would not have such an impact on any large economy, but a Spanish default would increase the default risk in Portugal by up to 50 bps (Figure 4-B). Both Italy and Spain would also impact Ireland, although the effect from Italy is about three times larger (notice that Ireland is plotted on the right-hand axis in Figure 4-A). Overall, the difference between the effects of a default in Italy and Spain is not surprising, given that Italy's foreign debt was three times larger than Spain's in 2010 and 2011.

To account for the large differences across sovereigns in the amount of their foreign debt, we construct a measure that normalizes for the total amount of external debt of the country with the initial default. This measure gives the expected spillover losses due to any additional defaults, per dollar of foreign debt of the initial country. It 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 et al. 2014; Bonaldi, Hortaçsu, and Kastl 2014). As with that measure, our measure can be used to analyze systemic risk and identify which sovereigns pose the greatest threat.[30]

Our measure is defined 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 original 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 country j, and any indirect effects from the further losses of repayments from other countries (k, etc.) that default because they are not repaid either. This includes any higher order sequences of losses because a new payment

[30]The Katz-Bonacich centrality measure does not apply directly in our case because our model is nonlinear. However our measure similarly captures all the higher order (i.e., indirect) effects of an initial shock.

equilibrium is determined for the scenario in which country j defaults. Then, given the baseline and simulated probabilities, the expected spillover losses from country i due to the default of country j is $[\hat{p}_{it} - \tilde{p}_{it}(j)]D_{it}$. We add these losses across all countries, and divide by the total foreign debt of country j to normalize, which yields our measure:

$$\lambda_{jt} \equiv \frac{1}{D_{jt}} \sum_{i \neq j} [\hat{p}_{it} - \tilde{p}_{it}(j)]D_{it}. \tag{7}$$

This gives the *expected spillover losses per dollar of debt* of country j. Our measure can be thought of as expressing the "contagiousness" of a sovereign's foreign debt. It captures differences in the potential for spillovers that arise from 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 from the total amount of its foreign debt.

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 large: for each \$1 of debt directly lost in default, the expected losses from additional defaults at other countries are less than 2 cents. The levels and trends are generally similar among all the countries in both panels. Greece and Portugal have relatively more contagious debt, as does Germany, followed by Ireland, France, Italy, and Spain, and last the United Kingdom.[31] 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 has the highest potential for contagion, with expected spillover losses per dollar of debt that are roughly double those of any other sovereign (note the different scale for this plot). This turns out to be the case because Italy holds a relatively large share of Austria's debt, and Italy is relatively sensitive

[31] The UK has lower potential spillover losses than other countries because a relatively large proportion of its debt is held outside Europe (e.g., by the United States). This debt is included in the normalization but is not counted toward spillover losses within the sample of European sovereigns.

to losses on its claims because it has a somewhat higher risk of default. Finally, the weighted average of the spillovers, which uses the total foreign debt amounts (D_{jt}) as weights, rose to almost 0.6 cents per dollar during the recession of 2008-09, then leveled at 0.4 cents until the end of 2010, when it began to rise steadily as the sovereign debt crisis widened in Europe.

To show the aggregate spillover losses, we combine the (unnormalized) expected spillovers from each sovereign weighted by the baseline probability that it would default. Thus the aggregate expected spillover losses in period t are

$$\sum_j (1-\hat{p}_{jt}) \sum_{i \neq j} [\hat{p}_{it} - \tilde{p}_{it}(j)] D_{it}.$$

To put this amount in context, we also compute the total losses implied by the baseline default probabilities: $\sum_j (1-\hat{p}_{jt}) D_{jt}$. Moreover, given the fit of our model, we know that these are very close to the total losses implied by the risk-neutral default probabilities imputed from the observed CDS spreads. These series are plotted in Figure 7. The total expected losses exceed \$400 billion in 2011-Q3 (as before, each quarter should be considered independently). At the same time the expected losses due to contagion are only \$4 billion. Throughout our sample period, the expected losses due to contagion represent about one half to one percent of the total expected losses.

A natural concern is whether these expected 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, the total expected losses should be interpreted as expectations for losses that may occur despite efforts to bail out a sovereign (e.g., as in the case of Greece in March 2012). Similarly, the expected losses due to contagion would incorporate beliefs about further bailouts to prevent additional defaults. These beliefs about bailouts should affect the total expected losses and the expected spillover losses in roughly equal fashion. Our finding that the spillovers represent a very small portion of the total losses would not be impacted unless market beliefs about the likelihood of a bailout are drastically different when countries are at risk due to contagion rather than their own internal factors.[32]

Finally, to provide some indication of the extent to which the risk of contagion impacts sovereigns' cost of borrowing, we simulate default probabilities under a scenario where there are no spillovers from a default and compare these with the baseline default probabilities. Specifically, we set the model so that a sovereign experiences no loss in repayments if one of its debtors defaults and recompute the solvency probabilities using the estimated parameters. Table 5 shows the results of this exercise for 2011-Q1. In most cases the difference between the baseline prediction and the no-spillover simulation are negligible. The largest differences are for Ireland and Portugal, which have a reduction in their risk-neutral default probabilities of 15 bps. This is a small fraction of their total default probabilities (less than 2%). The exact correspondence converting these default probabilities into interest rates on sovereign debt would depend on several factors, but the changes in the interest rates should be roughly commensurate. This indicates that the risk of contagion due to the balance sheet mechanism does not have an economically signficant effect on sovereigns' cost of borrowing.

6 Conclusion

In this paper we evaluate the economic importance of contagion in the European sovereign debt crisis. Building on the recent theoretical literature on financial networks, we construct a network model of credit risk among thirteen European sovereigns. Using data on sovereign

[32]Specifically, to have a downward bias on our estimate of spillover losses, the market would need to expect greater effort to bail out countries, such as Ireland and Portugal, that have substantial holdings of foreign sovereign debt, compared with Greece. In that case the solvency probability for Ireland and Portugal would not be as strongly related to the solvency probability of Greece, despite their relatively large holdings of Greek debt. We do not think this scenario is very plausible.

CDS spreads and the cross-holdings of sovereign debt, we estimate this model and conduct counterfactual experiments to quantify the spillover effects from a direct, "balance sheet" mechanism for contagion.

Our estimates imply that credit markets perceived the potential spillovers from a sovereign default to be small in magnitude. On average, the predicted losses due to contagion account for only one percent of the total expected losses implied by the sovereign CDS spreads in our sample of countries. This suggests that the interconnectedness of sovereign debt holdings in Europe does not appear to have an economically significant effect on sovereigns' cost of borrowing. Moreover, simple regression estimates show that comovements in sovereign credit risk have little relationship to the aggregate financial linkages between countries. As a consequence, any model where the transmission of risk is related to these linkages would likely find small spillovers from this channel for contagion, in the European crisis.

These results may provide some guidance toward evaluating the net benefits of sovereign bailouts. Much of the justification for bailouts in the recent crisis relied on an assumption that there would be large contagion effects resulting from a sovereign default. In contrast, our estimates suggest that the risk of contagion—from direct losses to the value of debt holdings, in particular—was relatively small. This provides one assessment of an important channel for contagion, which may be useful to consider among the many factors that bear on the benefits and costs of a sovereign bailout.

A Analysis of Potential Biases

Here we provide a detailed analysis of the biases that could arise if certain key assumptions in our empirical model were violated. We consider four potential issues: correlations in the financial shocks between countries, the endogeneity of financial linkages, endogenous default decisions, and internal amplification mechanisms with different impacts across countries. Our focus is on the bias in the estimate of γ, the parameter that governs the magnitude of spillovers from a default, and in each case we show that the likely bias is upward.

Correlations in financial shocks. Although we speculate that a correlation in the financial shocks among countries could be identified, we estimate the model under the assumption that they are independent. If, in fact, the shocks were correlated, such as X_{it} and X_{jt}, the observed correlation between p_{it} and p_{jt} would reflect this in addition to the true effects of repayments between countries i and j (i.e., $\gamma l_{ijt}[\delta+(1-\delta)s_{jt}]$ and $\gamma l_{jit}[\delta+(1-\delta)s_{it}]$). Because the repayments are an increasing function of the shocks (via s_{it} and s_{jt}), 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 between countries would result in an upward bias.[33]

Endogeneity of financial linkages. In general, a bias would arise if the observed claim amounts l_{ijt} were correlated with the unobservables in period t. Because these claims are established at the end of the previous period, this would require that the unobservables, specifically the financial shocks, are correlated over time.[34] Then, in some model of sovereign debt holdings (which is beyond the scope of this paper), banks in country i might reduce their holdings of debt from country j if a low value of the shock $X_{j,t-1}$ is observed, because this 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 (l_{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}). However, our predicted repayments in the model (5) do not account for the lagged shocks (i.e., $X_{j,t-1}$), so there would be an error term for the difference between the correct predictions and our predictions (heuristically, $\mathrm{E}[R_{it}|X_{j,t-1},\dots]-\mathrm{E}[R_{it}|\dots]$). This error term would have a positive correlation with our predicted values of R_{it} because of the positive correlation between l_{ijt} and X_{jt} described above. Therefore, this positive correlation would generate an upward bias in the estimate of γ.

[33] If this correlation were uniform among all countries, it might be possible to address with period-specific effects like π_t (although this can be problematic in a nonlinear model). If the correlations were different for different country-pairs, as in say ρ_{ij}, the bias could depend on the relationship between ρ_{ij} and l_{ijt}. It seems likely that the shocks would be more highly correlated between countries with stronger linkages, so that ρ_{ij} and l_{ijt} would be positively related, which would again yield an upward bias in the estimate of γ.

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

On the other hand, if investors in country i tend to diversify away from home when their own country experiences a negative shock, then there could be a downward bias. In this case, given a low value of the domestic shock $X_{i,t-1}$, banks in country i might increase their holdings of debt from other sovereigns. This would introduce a negative correlation between l_{ijt} and X_{it} that is not accounted for in our model. However, we think this is an unlikely scenario for two reasons. First, if shocks are correlated across countries (as discussed above), then typically there would be no relative advantage to increasing investment abroad. Second, there is evidence that banks tend to increase their "home bias" in times of crisis, even in crises which specifically relate to domestic conditions (Giannetti and Laeven 2012a,b).

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 L_t, D_t, Y_t, and X_t. Conceptually, we could incorporate such a decision rule into the solvency condition with a policy function $\pi_i(L_t, D_t, Y_t, X_t)$ that adjusts the threshold for default up or down from fixed values captured by the parameters $\bar{\pi}_i + \bar{\pi}_t$ (bars added here for clarity).[35] Then $-\pi_i(L_t, D_t, Y_t, X_t)$ would appear in (5) as an error term, and so the question is how it would be correlated with our predicted repayments based on an exogenous default rule. 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 $\pi_i(L_t, D_t, Y_t, X_t)$ 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 produce an upward bias in the estimate of γ.

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 (5), 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 more debt holdings, then the sensitivity to losses (γ_{it}) should be greater when the total holdings ($\sum_{j \neq i} l_{ijt}$) 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 $\bar{\gamma}$.

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

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 according to country of ultimate counterparty at a quarterly frequency. This measure includes all financial assets, not just sovereign debt, that is held by the financial sector. Define b_{ijt} as the value of sovereign i's financial claims on sovereign j at date t as reported by the BIS. We define an adjusted claims measure as

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

Note that the data on financial holdings provided by the BIS include countries outside of the 13 European countries in our sample. As indicated by the second term in the denominator above, we include these data to compute our adjusted claims measure (#BIS is the total number of BIS reporting countries). To construct the network of holdings used in our estimation, we use the adjusted claims to weight each counterparty's total externally held sovereign debt. 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 claims on sovereign j held at date t as

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

This includes the normalization for sovereign i's 2004 GDP ($Y_{i,2004}$). These l_{ijt}'s are the bilateral claims that are ultimately used in our estimation of the network.

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 spread in a CDS).

We obtain a time series of CDS spreads for each of the 13 sovereigns in our sample using data from CMA. We only have complete data in our sample for 5-year CDS contracts denominated in US dollars. As such, we assume a constant hazard rate for a sovereign default in order to impute risk-neutral solvency probabilities from the 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. Given data for the CDS spreads, 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.

References

Acemoglu, Daron, Asuman Ozdaglar, and Alireza Tahbaz-Salehi. Forthcoming. "Systemic risk and stability in financial networks." *American Economic Review*.

Acharya, Viral V., Itamar Drechsler, and Philipp Schnabl. Forthcoming. "A Pyrrhic Victory? Bank Bailouts and Sovereign Credit Risk." *Journal of Finance*.

Aguiar, M., and G. Gopinath. 2006. "Defaultable debt, interest rates and the current account." *Journal of International Economics* 69:64–83.

Allen, Franklin, and Ana Babus. 2009. "Networks in Finance." In *The Network Challenge: Strategy, Profit, and Risk in an Interlinked World*, Paul Kleindorfer, Yoram (Jerry) Wind, and Robert Gunther, eds. Wharton School Publishing.

Ang, Andrew, and Francis A Longstaff. 2013. "Systemic sovereign credit risk: Lessons from the US and Europe." *Journal of Monetary Economics* 60 (5):493–510.

Arellano, Cristina. 2008. "Default risk and income fluctuations in emerging economies." *American Economic Review* 98 (3):690–712.

Arellano, Cristina, and Yan Bai. 2013. "Linkages Across Sovereign Debt Markets." National Bureau of Economic Research Working Paper 19548.

Arghyrou, Michael G., and Alexandros Kontonikas. 2012. "The EMU sovereign-debt crisis: Fundamentals, expectations and contagion." *Journal of International Financial Markets, Institutions and Money* 22 (4):658–677.

Beirne, John, and Marcel Fratzscher. 2013. "The pricing of sovereign risk and contagion during the European sovereign debt crisis." *Journal of International Money and Finance* 34:60–82.

Benjamin, David, and Mark LJ Wright. 2009. "Recovery before redemption: A theory of delays in sovereign debt renegotiations." Unpublished.

Benzoni, Luca, Pierre Collin-Dufresne, Robert S Goldstein, and Jean Helwege. 2012. "Modeling credit contagion via the updating of fragile beliefs." Federal Reserve Bank of Chicago Working Paper 2012-04.

Blume, Lawrence E., William A. Brock, Steven N. Durlauf, and Yannis M. Ioannides. 2011. "Identification of Social Interactions." In *Handbook of Social Economics*, vol. 1, Jess Benhabib, Alberto Bisin, and Matthew O. Jackson, eds. North-Holland, 853–964.

Bolton, Patrick, and Olivier Jeanne. 2011. "Sovereign default risk and bank fragility in financially integrated economies." *IMF Economic Review* 59 (2):162–194.

Bonaldi, Pietro, Ali Hortaçsu, and Jakub Kastl. 2014. "An Empirical Analysis of Systemic Risk in the EURO-zone." Unpublished.

Bramoullé, Yann, Habiba Djebbari, and Bernard Fortin. 2009. "Identification of peer effects through social networks." *Journal of Econometrics* 150 (1):41–55.

Brock, William A, and Steven N Durlauf. 2007. "Identification of binary choice models with social interactions." *Journal of Econometrics* 140 (1):52–75.

Cohen-Cole, E., E. Patacchini, and Y. Zenou. 2011. "Systemic Risk and Network Formation in the Interbank Market." CEPR Discussion Paper 8332.

Constancio, V. 2012. "Contagion and the European debt crisis." *Financial Stability Review* (16):109–121.

Denbee, Edward, Christian Julliard, Ye Li, and Kathy Yuan. 2014. "Network Risk and Key Players: A Structural Analysis of Interbank Liquidity." Unpublished.

Dieckmann, S., and T. Plank. 2012. "Default Risk of Advanced Economies: An Empirical Analysis of Credit Default Swaps during the Financial Crisis." *Review of Finance* 16 (4):903–934 .

Eaton, J., and M. Gersovitz. 1981. "Debt with potential repudiation: Theoretical and empirical analysis." *The Review of Economic Studies* 48 (2):289–309.

Eisenberg, Larry, and Thomas H Noe. 2001. "Systemic risk in financial systems." *Management Science* 47 (2):236–249.

Elliott, Matthew, Benjamin Golub, and Matthew O. Jackson. Forthcoming. "Financial networks and contagion." *American Economic Review*.

Giannetti, Mariassunta, and Luc Laeven. 2012a. "The flight home effect: Evidence from the syndicated loan market during financial crises." *Journal of Financial Economics* 104 (1):23–43.

———. 2012b. "Flight Home, Flight Abroad, and International Credit Cycles." *American Economic Review* 102 (3):219–24.

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

Gofman, Michael. 2013. "A network-based analysis of over-the-counter markets." Unpublished.

Gouriéroux, C., J.-C. Héam, and A. Monfort. 2012. "Bilateral exposures and systemic solvency risk." *Canadian Journal of Economics* 45 (4):1273–1309.

Kodres, Laura E, and Matthew Pritsker. 2002. "A rational expectations model of financial contagion." *Journal of Finance* 57 (2):769–799.

Krauth, Brian V. 2006. "Simulation-based estimation of peer effects." *Journal of Econometrics* 133 (1):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 (3):402–417.

Lee, Lung-fei, Xiaodong Liu, and Xu Lin. 2010. "Specification and estimation of social interaction models with network structures." *Econometrics Journal* 13 (2):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 (2):75–103.

Manski, Charles F. 1993. "Identification of Endogenous Social Effects: The Reflection Problem." *The Review of Economic Studies* 60 (3):531–542.

Neyman, J., and Elizabeth L. Scott. 1948. "Consistent Estimates Based on Partially Consistent Observations." *Econometrica* 16 (1):1–32.

Pan, J., and K.J. Singleton. 2008. "Default and recovery implicit in the term structure of sovereign CDS spreads." *Journal of Finance* 63 (5):2345–84.

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 (3):599–624.

Stern, Steven. 1997. "Simulation-Based Estimation." *Journal of Economic Literature* 35 (4):2006–39.

Sturzenegger, Federico, and Jeromin Zettelmeyer. 2008. "Haircuts: Estimating investor losses in sovereign debt restructurings, 1998–2005." *Journal of International Money and Finance* 27 (5):780–805.

Upper, Christian. 2011. "Simulation methods to assess the danger of contagion in interbank markets." *Journal of Financial Stability* 7 (3):111–125.

Yue, Vivian Z. 2010. "Sovereign default and debt renegotiation." *Journal of International Economics* 80 (2):176–187.

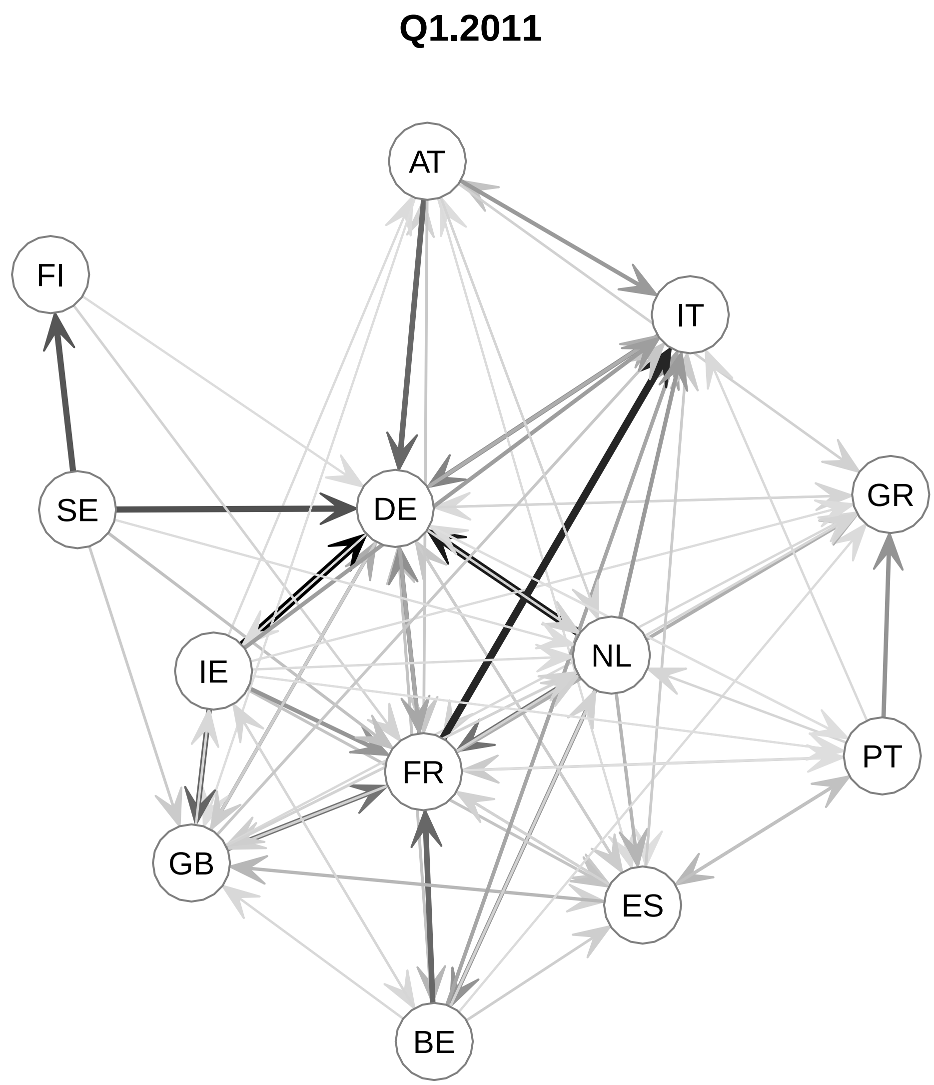

Figure 1: Network Graph, 2011-Q1. The figure displays the network structure of aggregate sovereign debt holdings in the first quarter of 2011. Countries are represented by their two letter abbreviation in Table 1. Arrows represent bank holdings from one country on the sovereign debt of another. Darker, thicker arrows indicate larger amounts in proportion to the creditor country's GDP in 2004.

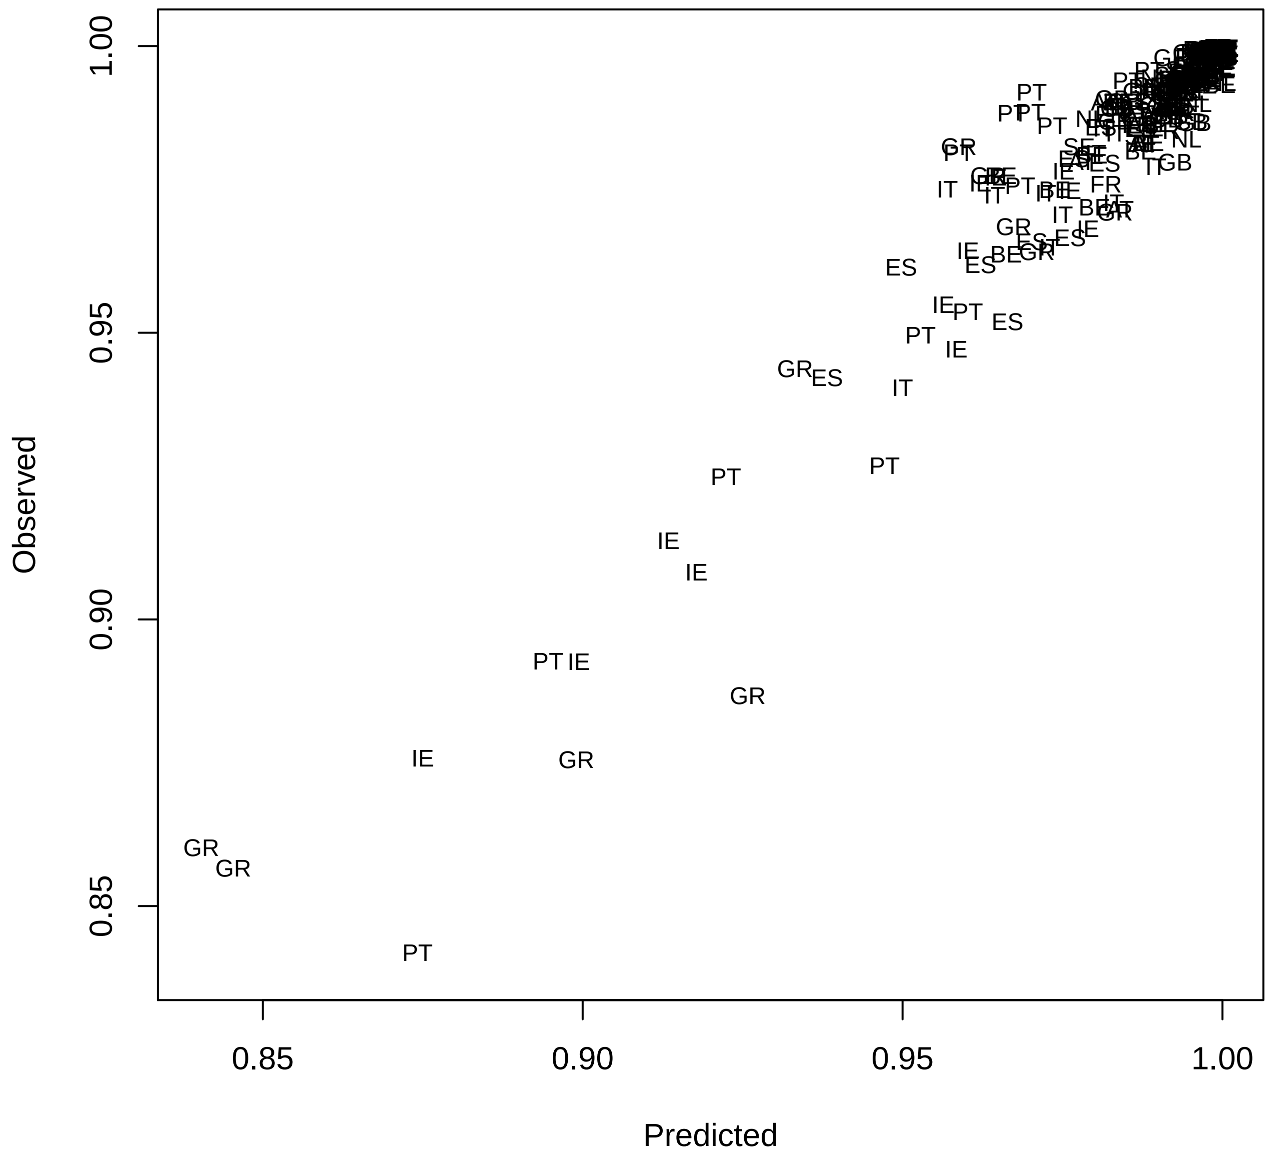

Figure 2: Predicted and Observed Solvency Probabilities. The figure plots the predicted and observed risk-neutral quarterly solvency probabilities for each country at each quarter in our sample. Observed solvency probabilities are obtained with a transformation of 5-year CDS contract prices, as described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

A.

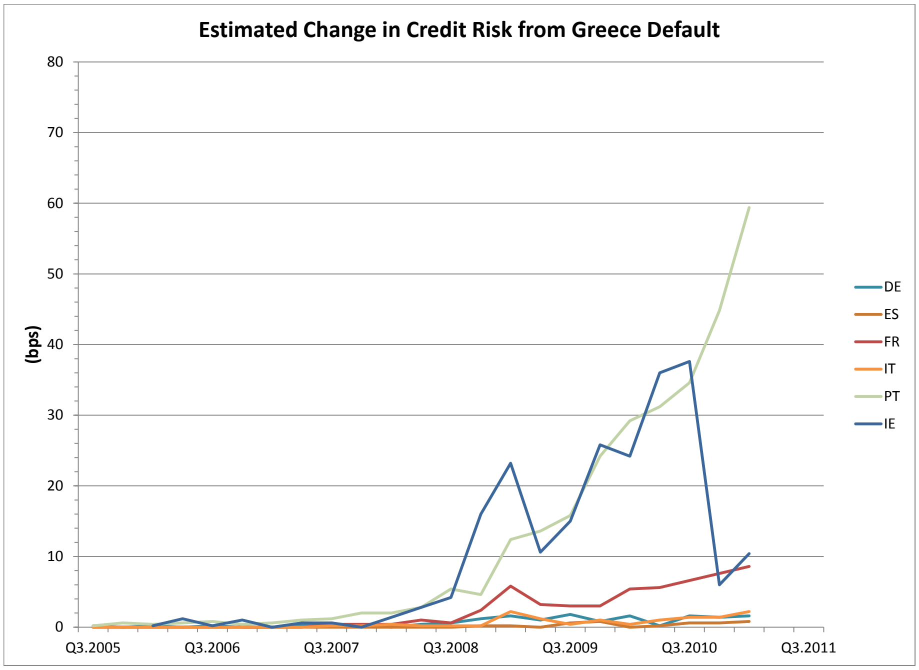

B.

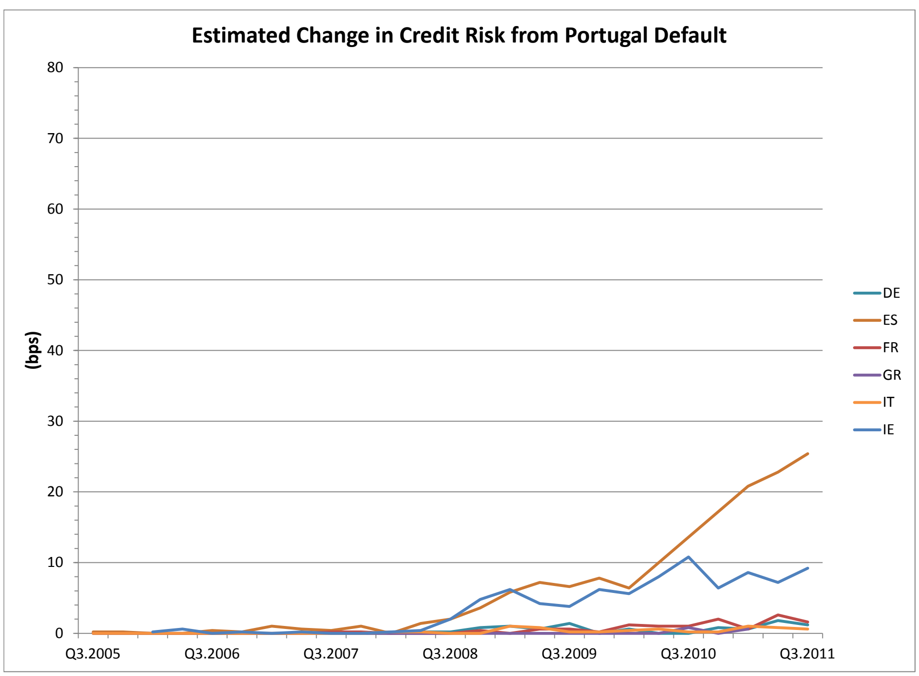

Figure 3: Simulated Default Probabilities: Greece and Portugal. The figure shows the change in the risk-neutral probability of default for selected sovereigns, assuming a default in Greece (A) or Portugal (B), in simulations using the estimated model. See Section 5.2 for details.

A.

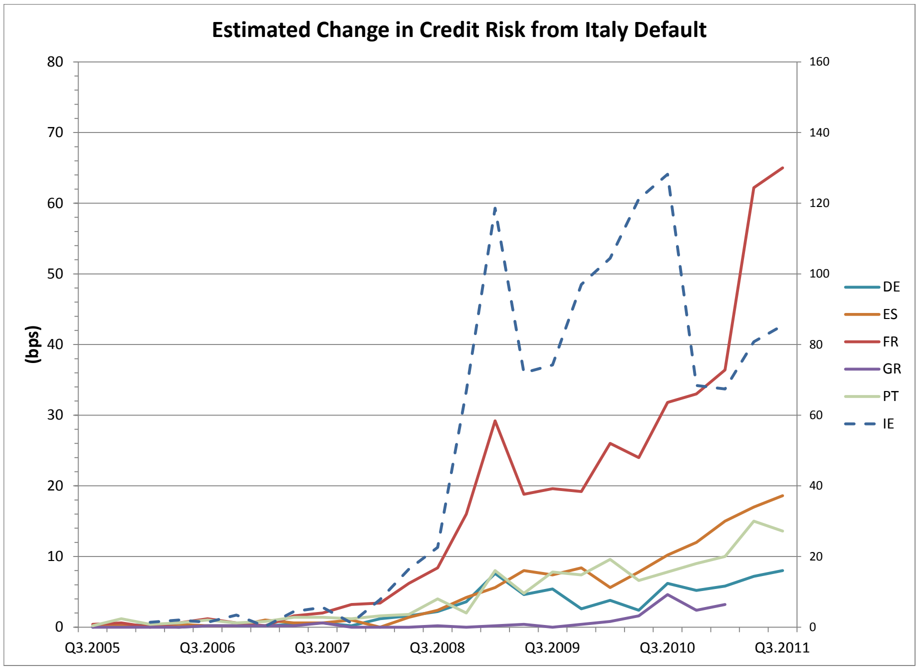

B.

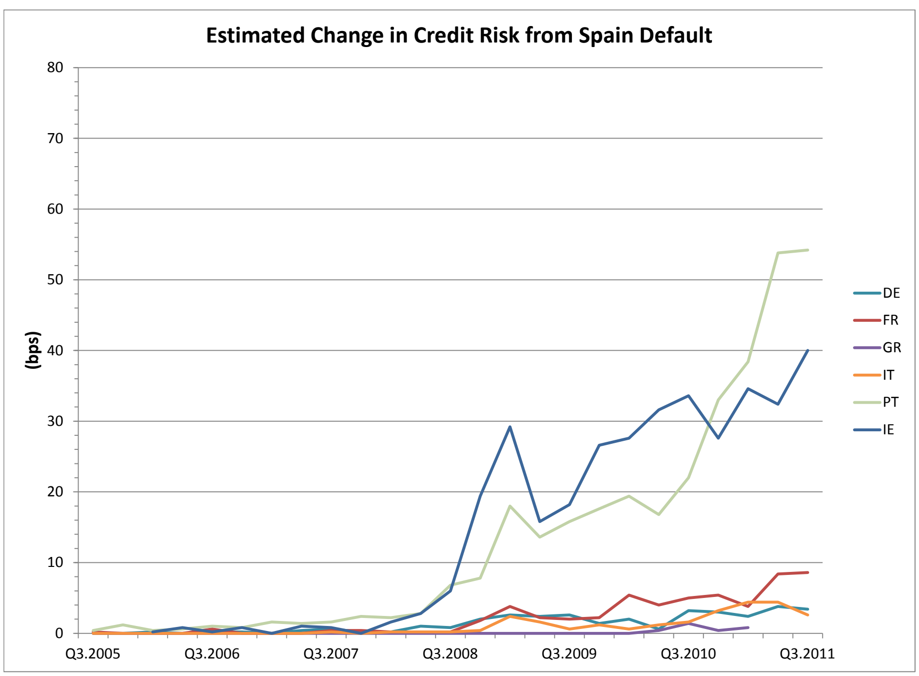

Figure 4: Simulated Default Probabilities: Italy and Spain. The figure shows the change in the risk-neutral probability of default for selected sovereigns, assuming a default in Italy (A) or Spain (B), in simulations using the estimated model. See Section 5.2 for details.

A.

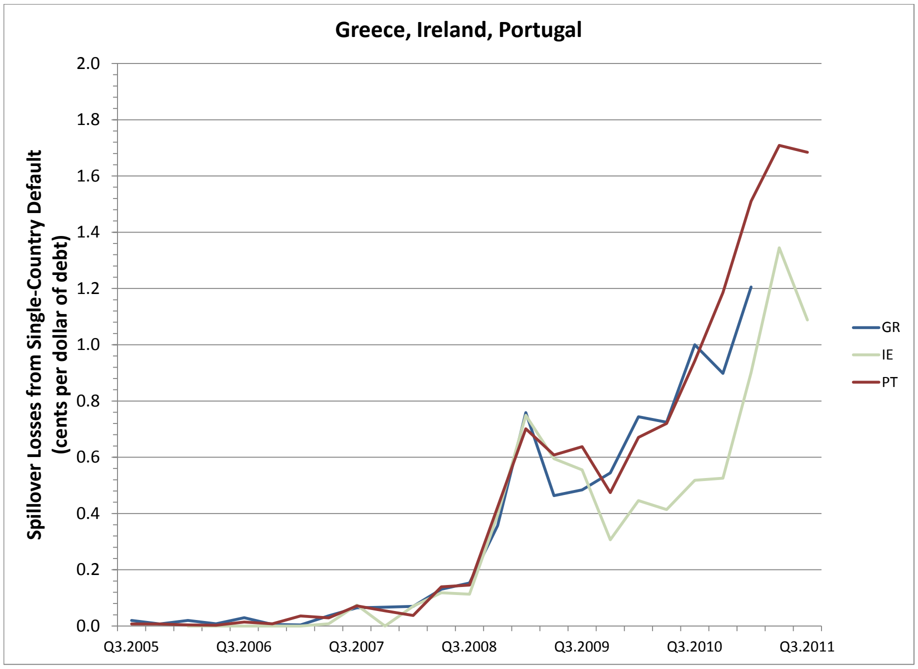

B.

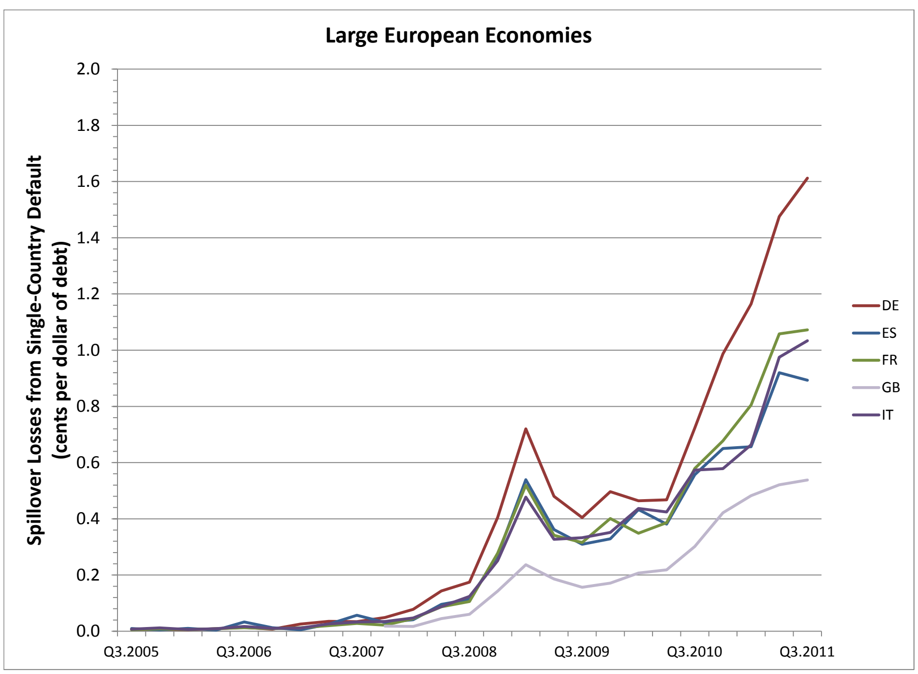

Figure 5: Expected Spillover Losses per Dollar of Debt: At-Risk Sovereigns and Large Economies. The figure plots our measure of expected losses due to increased risk of default at other countries, conditional on a default by the country indicated in each series, per dollar of that country's foreign debt. See equation (7) for the definition of this measure.

A.

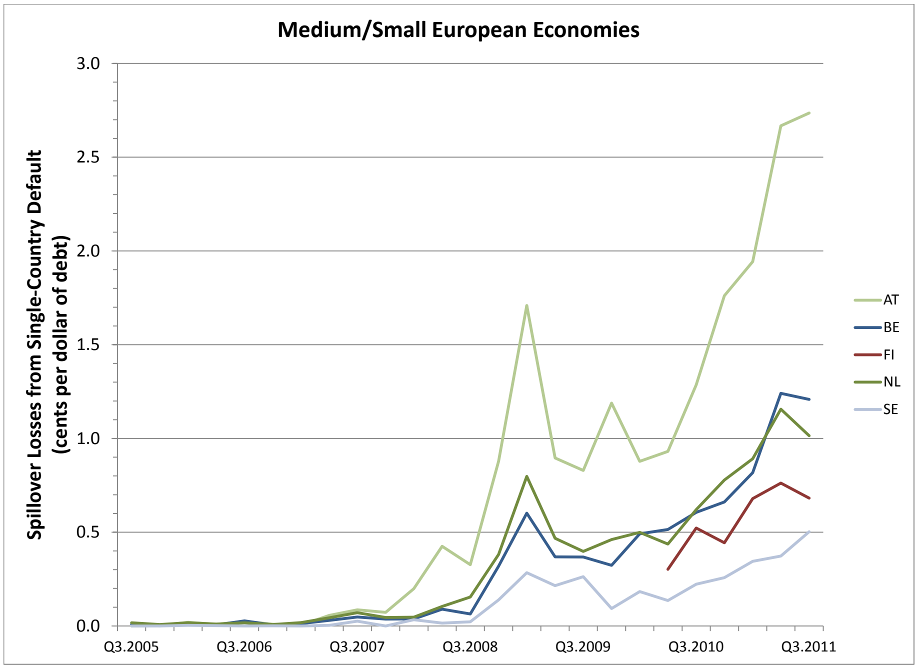

B.

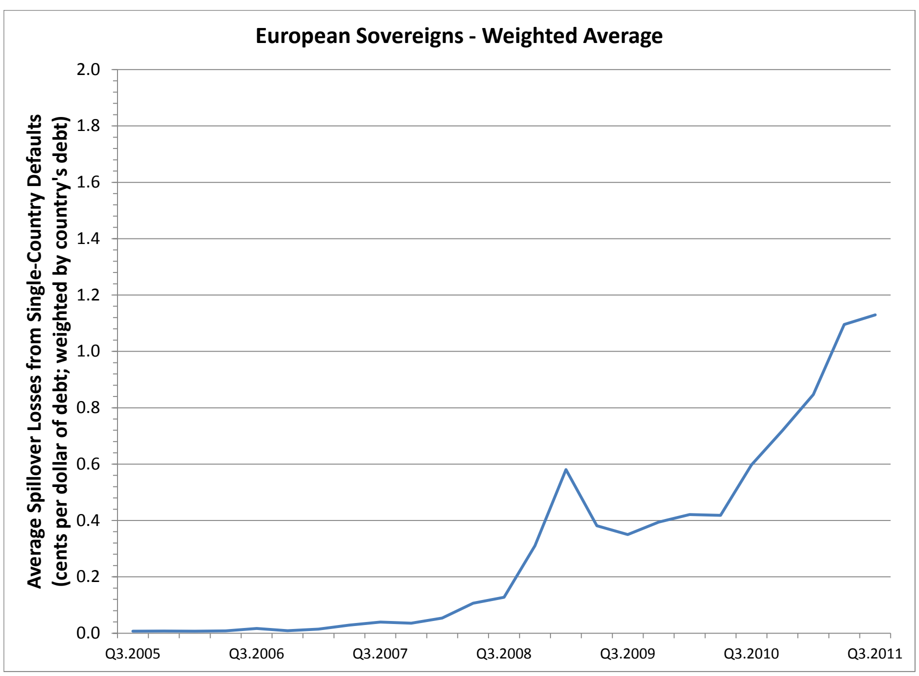

Figure 6: Expected Spillover Losses per Dollar of Debt: Medium or Small Economies, and Weighted Average. The figure plots our measure of expected losses due to increased risk of default at other countries, conditional on a default by the country indicated in each series, per dollar of that country's foreign debt. See equation (7) for the definition of this measure. Weighted average is among all countries in the sample, weighted by each country's total foreign debt.

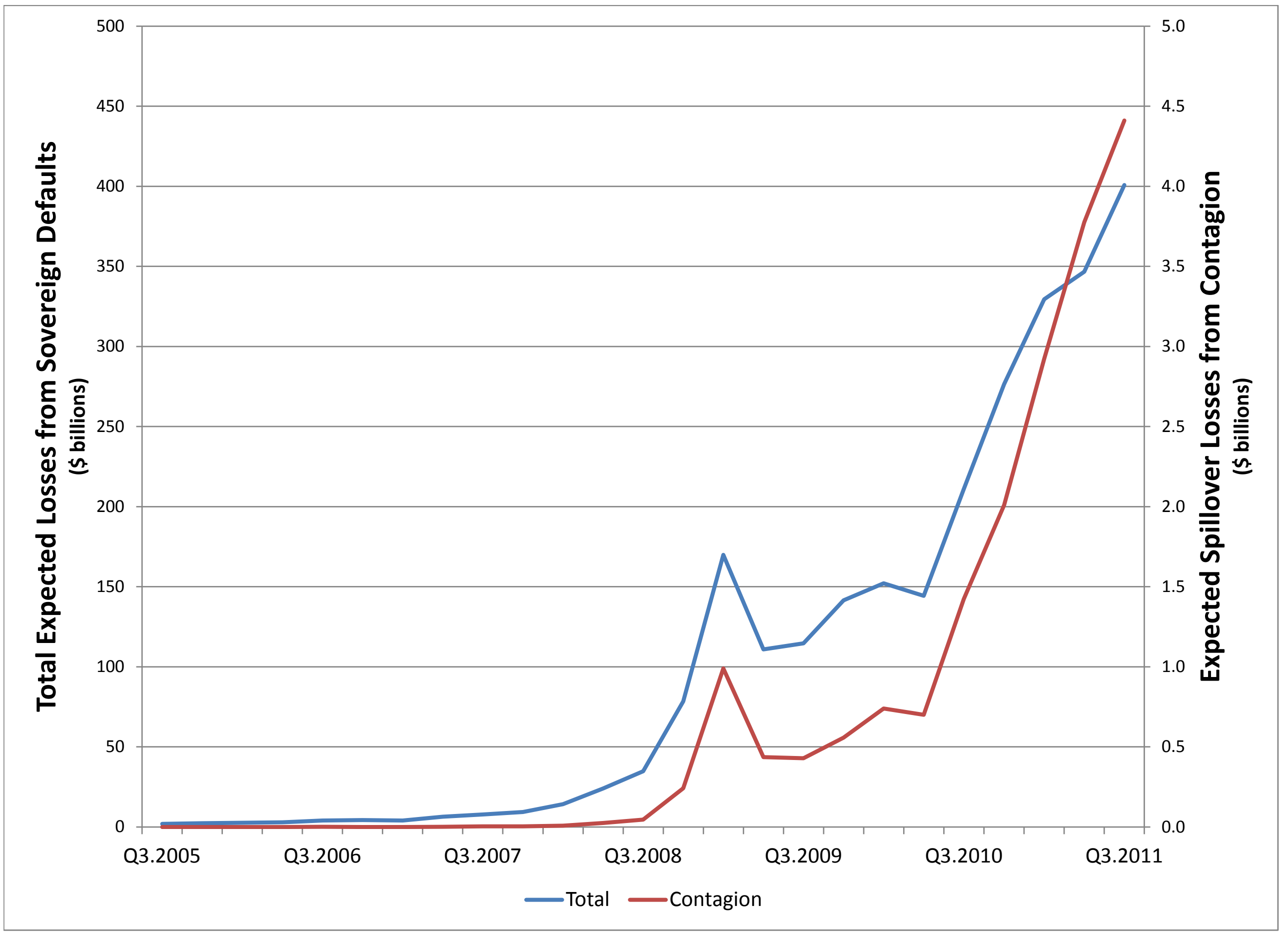

Figure 7: Total Expected Losses and Expected Losses Due to Contagion of Default. The figure plots the total amount of expected losses from all sovereigns in the sample, using the predicted solvency probabilities from the estimated model, and the amount of expected spillover losses from contagion of default. See Section 5.2 for further description.

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		

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

Table 2: **Summary Statistics for 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 ($\sum_{j \neq i} l_{ijt}$), constructed from BIS and IMF data. Total debt is total foreign 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 (6) 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
γ	0.746	0.021	Equilibrium repayments on claims (R)
α	0.280	0.008	Total foreign debt (D)
β_1	4.363	0.123	GDP level (Y)
β_2	0.135	0.004	Common component of GDP growth rate (ΔY^c)
β_3	0.042	0.001	Residual component of GDP growth rate (ΔY^r)

Notes: The table shows estimates of the parameters in equation (5) and marginal effects of the associated variables. Note that α enters the model negatively. Marginal effects are computed as the average of the marginal effect for each observation. The equilibrium repayments (R), foreign debt (D), and GDP level (Y) are normalized by the country's 2004 GDP.

Table 5: **Simulated Default Probabilities with No Spillovers, 2011-Q1**

Country	Baseline Prediction	No-spillover Simulation
AT	140.2	137.8
BE	244.6	239.4
DE	58.0	57.0
ES	386.6	384.4
FI	62.6	62.6
FR	111.6	107.0
GB	136.4	135.6
GR	1562.6	1562.0
IE	855.2	841.8
IT	345.8	343.6
NL	87.4	85.0
PT	724.4	709.4
SE	70.8	70.0

Notes: The table presents results from a counterfactual simulation in which spillover effects are eliminated. Specifically, a sovereign is assumed to experience no loss in repayments if one of its debtors defaults. The first column shows baseline risk-neutral default probabilites predicted for 2011-Q1 using the estimated model. The second column shows default probabilities in the counterfactual simulation, which uses the estimated parameter values, but sets $\delta = 1$ to eliminate the spillovers from a default. Probabilities are reported in basis points (bps).

Table A-1: Comparison with Alternative Measures of Claims on Sovereign Debt in 2010-Q4

	Correlations with Our Measure	
Country (creditor)	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 bilateral claims on sovereign debt and two other measures of bilateral claims on sovereign debt available from the BIS and EBA. The comparison is made using data for 2010-Q4. The first row shows the correlation across all countries as creditors, and subsequent rows show the correlation within each country in the amounts of their claims on other countries. (Note that each country has claims on up to $N-1$ other sovereigns.) See Section 4.1 for further details.