Credit Default Swap Spreads and Systemic Financial Risk

Stefano Giglio*
Harvard University

February 2011

Abstract

This paper presents a novel method to measure the joint default risk of large financial institutions (systemic default risk) using information in bond and credit default swap (CDS) prices. Bond prices reflect individual default probabilities of the issuers. CDS contracts, which insure against such defaults, pay off only as long as the seller of protection itself is solvent. Therefore, CDS prices contain information about the probability of joint default of both the bond issuer and the protection seller. If we consider the entire set of CDS contracts written by each financial institution against the default of each other institution we can learn about all pairwise default probabilities across the financial network. This information, however, is not sufficient to completely characterize the joint distribution function of defaults of these banks. In this paper, I show how this information can be optimally aggregated to construct bounds on the probability of systemic default events. This method enables me to measure systemic default risk without making any assumptions about the joint distribution function. Two main results emerge from the empirical application of this method to the recent financial crisis. First, I show that an increase in systemic risk in large global banks did not occur until after Bear Stearns' collapse in March 2008. Second, some of the large observed spikes in CDS spreads and bond yield spreads during this period (for example, following Lehman Brothers' default) correspond to spikes in idiosyncratic default risk rather than systemic risk.

*Contact: sgiglio@fas.harvard.edu. I owe special thanks to Effi Benmelech, John Campbell and Emmanuel Farhi for their many comments and suggestions. Also, I am very grateful to Ruchir Agarwal, Robert Barro, Hui Chen, Josh Coval, Tom Cunningham, Mikkel Davies, Ian Dew-Becker, Darrell Duffie, Ben Friedman, Bob Hall, Eugene Kandel, Leonid Kogan, Anton Korinek, David Laibson, Greg Mankiw, Ian Martin, Parag Pathak, Carolin Pfueger, Monika Piazzesi, Martin Schneider, Tiago Severo, Andrei Shleifer, Alp Simsek, Ken Singleton, Holger Spamann, Jeremy Stein, Bob Turley, Tom Vogl, Charles-Henri Weymuller and seminar participants at the Harvard Macro Lunch, the Harvard Finance Lunch, the MIT Sloan Finance Lunch, the Stanford Institute for Theoretical Economics, the Program for Evolutionary Dynamics at Harvard, the Federal Reserve Bank of Chicago, London Business School, NYU Stern, Dartmouth Tuck, Stanford GSB, Northwestern Kellogg, Princeton, MIT Sloan, Chicago Booth, Berkeley Haas, Yale SOM, Carnegie Mellon Tepper and Wharton for many helpful suggestions.

1 Introduction

This paper seeks to measure the probability that several large financial institutions default within a short time. During periods of financial distress, this probability is not negligible because large financial intermediaries are highly interconnected and are exposed to common shocks. The central role of these institutions in the global economy makes this issue especially important. In particular, in this paper I focus on measuring the probability of default of *at least* r large financial institutions. I refer to these joint default events as *systemic default risk of degree* r.

Measuring systemic risk is a difficult task. On the one hand, measures that are based directly on the books of financial institutions are mostly backward-looking in nature. They are limited by the complexity of the balance sheet, the risks involved and the availability of data. On the other hand, measures based on the historical distribution of returns require estimating joint tail probabilities from limited time series, which is possible only under strong parametric assumptions.

The most widely used measures of systemic risk, like the one proposed in this paper, are instead market-based. They try to circumvent the limitations outlined above by aggregating individual default probabilities of financial institutions obtained from the prices of traded securities like credit default swaps (CDSs).[1] These market-based measures are forward-looking and reflect the information set of market participants; they measure risk-neutral rather than objective default probabilities.

Figure 1 plots two simple examples of such market-based measures of default risk in the financial sector: the average 5-year bond yield spread and the average 5-year CDS spread of the largest 15 financial institutions by CDS activity[2] between 2004 and 2010. The yield spread (the yield on a firm's bonds in excess of the risk-free rate) and the CDS spread (the cost of insuring against the firm's default) both reflect the probability that a firm defaults. The idea behind these measures is that an increase in systemic risk in the financial sector should cause the risk of default of each institution to increase. This should result in an increase of both the average yield spread and the average CDS spread. These measures suggest an increase in systemic risk starting in August 2007, followed by several episodes in which systemic risk spiked (such as around March 2008, September 2008 and then March 2009), and a final drop after April 2009.

The existing market-based measures can be misleading for two reasons. First, they involve strong modeling assumptions in aggregating the individual risks of financial intermediaries into estimates of systemic risk. For example, the CDS-based measure reported in Figure 1 is only informative about the joint distribution of defaults under strong assumptions about the relationship between marginal and joint default probabilities. Second, existing measures based on the prices of securities traded over the counter (OTC) ignore counterparty risk. As I show below, ignoring counterparty risk introduces a bias that increases precisely when the financial system is distressed.

In this paper I propose a novel market-based measure of systemic risk, based on a combination

[1] A credit default swap is an insurance contract against the default of a firm, for example a financial institution. The CDS spread corresponds to the yearly insurance premium. See section 2 for details on the contract.

[2] The Figure replicates the Counterparty Risk Index, produced by Credit Derivative Research. The index was created to allow buyers of CDS protection to easily hedge their counterparty risk exposures, and therefore includes the spreads of the most active counterparties in the CDS market. These 15 institutions alone constitute more than 90% of the CDS protection sold.

of bond prices and CDS spreads. I exploit the pricing of counterparty risk in CDS contracts to learn about the joint default risk of pairs of institutions. Because this information set is not rich enough to completely characterize the full joint distribution function across the network, I use linear programming to construct the tightest possible bounds on systemic risk consistent with the observed prices. This allows me to measure systemic risk without making modeling assumptions about the relationship between risks of individual institutions and joint risks. In addition, it allows me to track the contribution to systemic risk of each bank over time.

The starting point for the analysis is the presence of counterparty risk in CDS contracts. When the seller of protection manifests higher risk of default, the value of the default insurance (the CDS spread) decreases, particularly if the two defaults are correlated. Therefore, the spread of a CDS written by a financial institution against the default of a bond reflects both the probability of default of the bond issuer, called the *reference entity,* and the risk of joint default with the seller of protection. The price of the bond instead reflects only the marginal probability of default of the firm which issued the bond.[3] Combining bond and CDS prices enables us to infer the joint default probability of the two entities.

A standard way to see this is to look at the *bond/CDS basis.* By buying the bond and insuring it with the corresponding CDS of the same maturity, one obtains a risk-free debt security as long as there is no counterparty risk in the CDS contract. An approximate arbitrage relation then implies that in the absence of counterparty risk, the yield spread on the bond issued by i over the risk-free rate $(y^i - r^F)$ should be equal to the corresponding CDS spread z_i written on i. That is, the bond/CDS basis, defined as the difference between the two $(z^i - (y^i - r^F))$, should be zero. Counterparty risk, by lowering the CDS spread z^i without affecting the yield spread of the bond, produces a negative basis. The bond/CDS basis therefore contains information about the joint default risk of the reference entity and the protection seller. Figure 1 shows that the average bond/CDS basis for financial institutions is negative, as expected, and varies significantly over time.

The entire set of bond prices of financial institutions and CDS spreads written on them by other financial institutions allows us to learn all marginal and pairwise probabilities of default across the financial network. In general, this information set is not sufficient to completely pin down systemic default risk. Standard approaches to measuring systemic risk overcome this problem by imposing modeling assumptions, such as Gaussian or Student-t copulas, that allow them to obtain point estimates of the joint default probabilities.

In this paper, instead, I construct the tightest upper and lower bounds on the average monthly probability of joint default of at least r financial institutions (for different values of r) consistent with the observed bond and CDS prices. These bounds are obtained making *no assumptions* about the joint default distribution and only use marginal and pairwise default information inferred from security prices. Yet, because they are constructed to be the tightest possible bounds for systemic risk, they prove to be quite informative about joint default risk.

Three limitations affect the construction of the bounds. First, the presence of an unobserved

[3] Of course, other factors affect bond prices as well. The most important one, liquidity, is explicitly accounted for in the paper.

liquidity process in the bond market confounds the filtering of individual default probabilities out of CDS spreads. Nevertheless, interesting results can be obtained under minimal assumptions on the liquidity process. In particular, I only impose a lower bound for this process, calibrated from the time series of bond and CDS prices of each bank, or from the cross-section of financial and nonfinancial firms of comparable credit rating. Second, for every reference entity, I observe only an average of the CDS quotes posted by the main counterparties, so my information set is smaller than the ideal one.[4] Third, this paper obtains risk-neutral, not objective, default probabilities.[5] The risk-neutral probabilities are interesting per se because they reveal the perception of the markets about the severity of these states of the world. In addition, they can be considered upper bounds on the objective default probabilities, because default states are bad states with high marginal utility.[6] Finally, I show strong implications for objective probabilities under mild assumptions on the utility function.

This analysis, applied to the period between January 2004 and June 2010, shows that when we consider the information contained in both bond *and* CDS prices, we can bound systemic risk to be much lower than if we used bond prices or CDS spreads alone. Contrary to other measures of systemic risk, which report a sharp increase in systemic risk already in 2007, we can exclude a large increase in systemic default risk before Bear Stearns' failure in March 2008. Moreover, we can show that observed spikes in CDS spreads and bond yields in the month before Bear's collapse and in September 2008, after Lehman's default, do not correspond to spikes in systemic risk (as standard measures report), but reflect sharp increases in *idiosyncratic* default risk of one or a small number of banks. Instead, systemic risk seems to increase smoothly after an initial jump related to Bear's failure.

While the optimal bounds are a complex function of the constraints imposed by the observed prices, the intuition for the main results is straightforward. If systemic risk had been high in 2007 and early 2008, the price of insurance against large dealers - purchased from other large financial institutions - should have dropped considerably. But this did not happen. The relatively high equilibrium spreads during this period impose a tight upper bound on the amount of systemic risk perceived in financial markets. Equivalently, the bond/CDS basis – which reflects joint default risks – did not increase enough during 2007 and early 2008 to be consistent with an increase in systemic default risk. This paper allows for a decomposition between idiosyncratic and systemic default risk that would be impossible to achieve without the information contained in the basis.

Finally, the methodology presented in this paper allows me to study the configuration of the financial network over time, as well as the contribution to systemic risk of each institution at the upper bound on systemic risk. This analysis shows that markets anticipated by more than a month a sharp increase in the joint default probability of two key institutions, Lehman Brothers and Merrill Lynch, which in fact ended up facing default on the same weekend (13-14 September 2008). Similarly,

[4] This limitation is also the reason I look at the group of dealers who are counterparties to most contracts (more than 90% of the market). This way, I can be confident I am including most of the dealers from whom the CDS quotes are obtained.

[5] Anderson (2009) underlines the differences between the two by comparing risk-neutral default processes obtained from CDS spreads with objective processes obtained using historical data on defaults.

[6] This argument assumes that the marginal investor is risk averse and internalizes the risks of these securities.

these banks' *contribution to systemic risk* (i.e. the probability of a systemic event in which they would be involved) increased sharply well before their collapse.

The paper proceeds as follows. After a brief literature review, section 2 presents an introduction to credit default swaps and counterparty risk. Section 3 presents the theory of the optimal probability bounds, and section 4 discusses the main issues of implementation and the data. Section 5 presents the empirical results. In section 6, I run a series of robustness tests. Section 7 concludes.

1.1 Related literature on measuring systemic risk

The literature on measures of systemic risk in the financial sector is large and has grown further following the financial crisis of 2007-2009. Four categories of papers can be identified, based on different methods they use to quantify systemic risk and to study the relative contribution of each bank to this risk.

First, structural approaches look directly at the books of financial institutions in order to learn about the distribution of joint shocks. Lehar (2005) uses the Merton (1974) model to obtain the time-series of the market value of banks' assets, using observed equity prices and balance sheet information. He produces a measure of joint default risk under the assumption of multivariate normal distribution of returns. Gray et al. (2008) propose using a similar approach to measure systemic risk not only within the financial sector but also across sectors and countries.

The structural approach requires strong assumptions about the liability structure of financial institutions, as well as about the marginal and joint distribution of risks. To overcome these difficulties, reduced-form approaches look at the historical distribution of returns. For example, Acharya et al. (2010) compute a measure of individual contributions to systemic risk from individual banks' equity returns during periods of negative returns for the financial sector as a whole. Similarly, Adrian and Brunnermeier (2009) use quantile regression to estimate the VaR of the financial sector as a whole conditional on each individual bank experiencing a loss in its asset values.

A reduced-form approach looking at historical returns has the disadvantage of trying to learn about tail events from a limited time series of returns. This requires strong assumptions about the tail behavior of return distributions. A third branch of the literature tries to avoid this problem by looking directly at the probabilities of tail risks implied by derivatives whose price is very sensitive to these precise risks. Lacking a traded security that directly reflects the joint default risk of the largest financial institutions, papers in this category typically extract marginal default risk information from CDS spreads, and make inference about joint default risk by aggregating them using a certain copula together with estimates of correlations. Examples are Huang, Zhou and Zhu (2009) and Avesani, Pascual and Li (2006), which assume Normal and Student-t copulas, and Segoviano and Goodhart (2009), which employs the CIMDO copula (Segoviano (2008)). These papers estimate a time-series of systemic risk that closely resembles the simple measures plotted in Figure 1, with large spikes around March and September 2008 and a steep increase in systemic risk starting in August 2007. These results are in sharp contrast with the ones presented in this paper, which does not need to assume any specific copula for the joint distribution of defaults but uses all the information in bond and CDS prices to learn about pairwise default risk.

Finally, several other papers have proposed measures of systemic risk that do not result in estimates of joint default risks. For example, Kritzman, Li, Page and Rigobon (2010) propose using the fraction of the variance explained by the first principal components of CDS spreads of large financial institutions as a measure of systemic risk in the financial sector. Other papers, instead, use individual or aggregated financial indicators to empirically predict financial crises, often in a cross-country setting. Examples of this approach are Poghoshan and Cihak (2009), Cihak and Schaeck (2007), Demirguc and Detragiache (1998 and 1999), and Gonzales-Hermosillo (1999).

2 Credit Default Swaps and Counterparty Risk

This section discusses the sources of counterparty risk in CDS contracts. I first describe the main characteristics of CDS contracts. Then, I discuss why counterparty risk in these contracts arises mainly from the possibility of double default of the reference entity and the counterparty. This risk can be sizable if defaults of financial institutions are not independent. Finally, I argue that collateral provisions, when present, are unlikely to eliminate this risk. Therefore, counterparty risk should be priced in CDS spreads.

2.1 The Credit Default Swaps Market

Credit default swaps are credit derivatives that allow the transfer of the credit risk of a firm between two agents for a predetermined amount of time. In a typical CDS contract, the protection seller offers the protection buyer insurance against the default of an underlying bond issued by a certain company (the reference entity). In the event of default by the reference entity, the seller commits to buy the bond for a price equal to its face value from the protection buyer.[7] In exchange for the insurance, the buyer pays a quarterly premium, called the *CDS spread*, quoted as an annualized percentage of the notional value insured. If default occurs, the contract terminates, and the quarterly payments are interrupted. If default does not occur during the life of the contract, the contract terminates at its maturity date.

While in general these contracts are traded over the counter and can be customized by the buyer and the seller, in the recent years they have become more standardized, following the guidelines of the International Swaps and Derivatives Association (ISDA). The CDS market is quite liquid, with low transaction costs to initiate a contract with a market maker on short notice, and with numerous dealers posting quotes (see Blanco et al. (2003) and Longstaff et al. (2005)). Reliable quotes for the 5-year maturity CDS can be obtained through several financial data firms (Bloomberg, Datastream, Markit).

The CDS market has grown quickly in the last few years. Notional exposures grew from about \$5 trillion in 2004 to around \$60 trillion at its peak in 2007, and despite the financial crisis, the total

[7] In practice, the terms of the CDS could involve physical delivery of the defaulted bond or cash settlement. In the former case, usually any bond of equal seniority can be delivered. For example, for the CDS written on a senior unsecured bond, any other senior unsecured bond of the firm could be delivered. In addition, the credit event could include restructuring or a downgrade of the reference bond. These clauses have a potential effect on the price of the CDS, discussed in Appendix C.

notional exposure is still around $40 trillion. The main reason for this growth in gross terms is that, due to the high liquidity of the CDS market, the easiest way to adjust the exposure to credit risk has been to enter new CDS contracts (possibly offsetting the existing ones) rather than operating directly in the bond market or cancelling CDS agreements already in place. At the center of this network of CDS contracts, a few main dealers operated with very high gross and low net exposures, emerging as the main counterparties in the market. For example, Fitch Ratings[8] states that in 2006 the top 10 counterparties (all broker/dealers) accounted for about 89% of the total protection sold. With the crisis, the market concentrated even more, after the disappearance of some of its key players.[9]

2.2 Counterparty Risk

Traded over the counter, a CDS contract involves counterparty risk: the protection seller might default during the life of the CDS and therefore might not be able to comply with the commitments implied by the contract.[10] In this case, the holders of CDS claims would still recover part of the expected payments due under the contract. Like other derivatives, CDS claims are treated pari passu with senior unsecured bonds, but they are also protected by "safe harbor" provisions, which exempt them from automatic stay of the assets of the firms, so that they can immediately seize any collateral that has been posted for them. In addition, positions across different derivatives with the same counterparty can be netted against each other. The latter potentially increases the recovery in case of counterparty default, but only if the buyer finds herself with large enough out-of-the-money positions with the seller when the seller defaults, thus hedging counterparty risk.[11]

In the case of early termination of the contract due to seller default, the seller has to compensate the buyer for the replacement cost of the contract, i.e. the cost of initiating a new insurance contract with another protection seller. This claim is small as long as the default risk of the reference entity does not jump substantially when the seller defaults, relative to the original terms of the contract. The larger the change in the CDS spread of the reference entity when the seller defaults, the larger the claim of the buyer against the defaulted counterparty. In the extreme case, where the default of the seller occurs simultaneously with the default of the reference entity, the payment due under the contract would be equal to the full insurance payment.

A simple two-period example of the pricing of bonds and CDSs can be useful to understand the role of counterparty risk. Consider a group of N dealers which have each issued a zero-coupon bond with a face value of $1 maturing at time 1, and the CDS contract written at time 0 by each of them against the default of each other dealer. Call A_i the event of default of institution i at time 1. Call $P(A_i)$ the probability of default of bank i, and $P(A_i \cap A_j)$ the probability of joint default of i and j

[8] Fitch Ratings, 2006, Global Credit Derivatives Survey.

[9] Fitch Ratings, 2008, Global Credit Derivatives Survey.

[10] The role of counterparty risk in CDS spreads has been studied by Hull and White (2001), Jarrow and Turnbull (1995), Jarrow and Yu (2001), and more recently, in the context of rare disaster risk, by Barro (2010). To reduce counterparty risk - which stems mainly from the OTC nature of the contract - there are now several proposal to create a centralized clearinghouse. For a detailed discussion, see Duffie and Zhu (2010).

[11] The possibility of entering offsetting contracts with the counterparty, rather than canceling existing ones, will work in this direction.

at time 1. All probabilities are risk-neutral. Call R the expected recovery rate on the bond in case of default, and suppose that in the event of joint default the CDS claim recovers a fraction $S \geq R$. Finally, assume that the risk-free rate between periods 0 and 1 is zero.

In this setting, the price of the bond issued by i, p_i, is determined as:

$$p_i = (1 - P(A_i)) + P(A_i)R \tag{1}$$

If there is no counterparty risk in the CDS contract, the insurance premium z_i, or CDS spread, paid at time 0 to insure that bond is:

$$z_i = P(A_i)(1 - R)$$

It is easy to see that between the bond and the CDS there is a theoretical arbitrage relation (Longstaff et al. (2005)): $z_i = 1 - p_i$. Consider now the case in which there is counterparty risk in the CDS contract. Then, the spread paid to buy insurance from j against i's default will be:

$$\begin{aligned} z_{ji} &= \left[P(A_i) - P(A_i \cap A_j)\right](1 - R) + P(A_i \cap A_j)(1 - R)S \\ &= \left[P(A_i) - (1 - S)P(A_i \cap A_j)\right](1 - R) \end{aligned} \tag{2}$$

since the buyer of protection obtains the full payment $(1 - R)$ only if the reference entity defaults alone, otherwise only a fraction S of it. Note that the spread z_{ji} decreases with the probability of joint default $P(A_i \cap A_j)$; the arbitrage relation with the bond is broken. It is also important to realize that the order of magnitude of counterparty risk could in theory be as high as the spread itself. While in models where defaults are independent we have $P(A_i \cap A_j) = P(A_i)P(A_j)$, most observers of the crisis would agree that defaults of major dealers are far from independent, and therefore the probability of the joint default can be of a much larger order of magnitude.

In this simple two-period example, I have considered only two cases of counterparty risk: the case of simultaneous default of the seller and the reference entity, with the corresponding loss of $(1-R)(1-S)$ to the protection buyer, and the case of default by the seller alone, with no loss to the protection buyer. In reality, it is possible that the defaults of the seller and the reference entity do not occur simultaneously, yet the buyer of protection incurs losses of the same order of magnitude as if they did. This can happen, for example, if the seller's default triggers a jump in the default probability of the reference entity, which might end up defaulting only some time later. In this case, the contract would be highly in the money immediately after the seller's default. Similarly, the buyer might suffer a loss if the default of the reference entity triggers the subsequent default of the counterparty: for example, because the seller did not adequately hedge the credit risk of the reference entity. In all these cases, the two defaults do not happen simultaneously, but they are connected in such a way that the protection buyer still suffers a potentially large loss on her claim. I refer to all these cases as *double default* cases.[12]

[12] Of course, it is also possible that the value of the contract increases when the seller defaults, but the reference entity does not default for a while, or does not default at all. In some cases, this might still induce a loss to the

2.3 Collateral Agreements and Pricing of Counterparty Risk

In order to protect the buyers against counterparty risk, some (but not all) CDS contracts involve a collateral agreement, under which collateral calls are tied mechanically to changes in the value of the CDS contract, as well as to downgrades of the rating of the protection seller. Typically, margins on CDSs are adjusted at a daily or weekly frequency. While helpful in reducing counterparty exposure, standard collateral agreements cannot eliminate the counterparty risk coming from double default. In this section I discuss the main reasons for this, and I provide additional evidence in Appendix A.

First, according to the ISDA Margin Survey 2008, only about 66% of the nominal exposure in credit derivatives (of which CDSs are the most important type) had a collateral agreement at all in 2007 and 2008; this number was even lower in the years before. In addition, as reported in the ISDA Survey, collateral agreements were employed much less frequently when the counterparty was a large dealer.

Second, several documents reveal that often collateral posted was even lower than the current value of the position. Even the buyer that most aggressively called for collateral during the crisis, Goldman Sachs, was not covered completely on its CDS exposures with other large dealers (in particular, AIG), let alone the potential exposure in case of sudden default of the reference entities.[13] As several documents show, smaller buyers of CDSs had much less collateral posted than the value of their positions - large dealers were usually able to obtain more collateral than they had to post when dealing with smaller counterparties.

Third, the nature of collateral posting is such that even for buyers who call enough collateral to fully cover the current value of their positions, the extent to which counterparty risk is reduced depends on the jump properties of the default events. As long as defaults are relatively well anticipated (there are no jumps), adjustment of collateral to the current value of the exposure can remove almost all counterparty risk. However, especially for financial intermediaries, defaults often occur suddenly and over very short periods of time (e.g. over the weekend), so that the buyer might not be able to obtain enough collateral to cover all the losses in time.

The Lehman bankruptcy is an interesting example of this. Until the weekend of September 13th-14th, during which Lehman collapsed and two other large financial institutions were bailed out (Merrill Lynch and AIG), the default risk of these financial institutions was deemed to be low, as reflected by low CDS spreads and high credit ratings. The joint shock to the three institutions happened suddenly, so that collateral adjustment would have been small. As it turned out, a *double default* event did *not* materialize, because of the government bailout - therefore, buyers of Merrill and AIG CDSs from Lehman did not experience large losses. However, these events show that the risk of simultaneous collapse of several banks was relevant, and that standard collateralization practices would not have prevented large losses to buyers of CDS contracts, had the government

protection buyer. However, as long as the default risk of the reference entity remains of the same order of magnitude, the value of the contract will be sufficiently small that the collateral posted will allow the recovery of most of it. In section 4 I describe more in detail the exact pricing model I use.

[13] The documents refer specifically to a large amount ($22bn) of CDS protection bought by Goldman from AIG on super-senior tranches of CDOs, but arguably similar practices were used on all credit derivatives instruments.

decided not to intervene.[14]

Finally, note that collateral calls themselves, if large enough, can cause the default of the protection seller. This contributed to the collapse of AIG in September 2008: the government had to step in to prevent collateral calls from bringing down the firm. The collateral demand can create an additional channel for joint default of the seller and the reference entity, and thereby increase the possibility of double default.

Therefore, the presence of collateral agreements – when in place at all – improves but does not solve the problem of counterparty risk related to double default. Buyers of CDSs were aware of this residual counterparty risk, as shown in documents reported in Appendix A. For example, Barclays Capital issued a report[15] in February 2008 precisely on the effect of counterparty risk on CDS prices. Buyers of CDSs frequently believed that the best way to reduce their counterparty exposure was to buy additional CDSs protection *against their counterparty* - which directly increased the total cost of buying CDS protection.[16] The presence of counterparty risk even in collateralized CDS contracts would then lower the value of the CDS insurance and therefore the spread buyers would be willing to pay for it.

It is worth discussing here a recent study of counterparty risk pricing by Arora, Gandhi and Longstaff (2009), who have access to quotes posted by different counterparties. They document that dealers with high default risk (as measured by the spread of the CDS written on them) posted quotes systematically lower than the other dealers for the same reference entity, and especially so after Lehman's bankruptcy. However, they also report that for most reference entities, the difference is very small, in the order of a few basis points. While at first sight this result seems to imply that counterparty risk is not priced in CDS contracts, this is not necessarily the case, for several reasons.

First, the study looks at the relation between the price quoted by each dealer and its *marginal* default risk. However, what matters for CDS pricing is the *joint* default risk of the reference entity and the counterparty; the results I present in the paper show that there can be a significant difference between marginal and joint default risk[17], and this may translate in a weak relation between marginal default risk and the quotes posted by the dealers. Second, the study looks at the variation of quotes around the daily average, and therefore it filters out all components of counterparty risk which are common to all dealers. Given that the financial crisis affected several large financial institutions at the same time, this effect could be quite large, if not dominant. In addition, as discussed in Section 4, the bounds I construct are based only on the *average* quote and therefore do not depend on the cross-sectional variation around it. On this respect, the results obtained in this paper complement those in Arora et al. (2009): their study focuses on the pricing of counterparty risk around the daily

[14]For example, a buyer of a 5-year Lehman CDS a month before its default would have been in the money, on Friday September 12th, for about 15 cents on the dollar (the present discounted value of the change in spread from about 350bp to 700bp during the previous month). Similarly, a buyer of Merrill CDS would have been in the money for 5 cents on the dollar. So, even if the buyer *did* have a collateral agreement, and if she *did* call for collateral up to the current value of the contract (both unrealistic assumptions), she would have not had more than 15 and 5 cents of collateral on those contracts on September 15th.

[15]Barclays Capital, 2008, "Counterparty Risk in Credit Markets".

[16]Of course, the same mechanism would have additional indirect costs related to the residual counterparty risk in the exposure on the latter CDS.

[17]See for example Figure 7.

average, while this paper concentrates on the average level of counterparty risk. This also implies that if for any reason dealers tend to post quotes that are not too far from the average, the results of this paper will still hold[18]. Finally, in this paper I allow – but not impose – average counterparty risk to explain part of the difference between bond yields and CDS spread. So, to the extent that in fact counterparty risk was *not* priced in these securities, this should be picked up by my methodology.

3 Probability bounds: theory

This section develops the theory of the probability bounds on systemic default events for a network of N institutions in which bond prices and CDS spreads are observed. I start with an introductory example that explains the main ideas. Then, I show how to use linear programming theory to solve the general problem, and I derive some properties of the optimal bounds. In this section, I assume that bond prices are affected only by the default risk of the issuer. I leave for section 4 a detailed discussion of the implementation of the bounds, including the adjustments needed to take into account the effect of liquidity on bond prices.

3.1 Probability bounds on systemic risk: an introductory example

Consider a two-period setting, and suppose that the financial sector consists of only three intermediaries - banks 1, 2 and 3. Since they are the only intermediaries in the market, protection against the default of $i \in I \equiv \{1, 2, 3\}$ must be bought from a bank $j \in I \backslash i$, i.e. one of the other two intermediaries. As shown by the pricing formulas (1) and (2), if we observe all bond prices p_i and all CDS spreads z_{ji}, we can learn the marginal default probabilities of each bank as well as the pairwise default probabilities for each pair (i, j) of banks (assuming we know the recovery rates). Because this information set contains the default probabilities of one or at most two institutions, but contains no direct information on the probability of the default event of all three institutions, call this a *probability information set of order two.* As an example, we might infer that

$$P(A_i) = 0.2 \ \ \forall i$$

$$P(A_1 \cap A_2) = P(A_2 \cap A_3) = 0.07, \ P(A_1 \cap A_3) = 0.01$$

As mentioned above, in this paper I define systemic risk *of degree r* the probability of joint default of at least r financial intermediaries, P_r. With only three banks, we obtain the following three measures of systemic risk:

$$P_1 = P(A_1 \cup A_2 \cup A_3)$$

[18] If instead only dealers with particularly high counterparty risk tend to post quotes that are relatively too high – for example to avoid signaling their high risk – the average quote will not reflect the true average counterparty risk, but will be biased upwards. Note that a higher average reported quote can be mapped into a higher recovery rate S in case of double default. Section 6 shows that the main results of the paper are robust to implied recovery rates of up to 90% – therefore, they will be robust to a significant bias in the way dealers report quotes. Essentially, for the main results of the paper to hold, all is required is that *some* (even only a part) of the *average* counterparty risk is reflected in the average CDS spread we observe.

$$P_2 = P((A_1 \cap A_2) \cup (A_2 \cap A_3) \cup (A_1 \cap A_3))$$

$$P_3 = P(A_1 \cap A_2 \cap A_3)$$

All these definitions involve unions and intersections of the three default events, and therefore for all values of r these are probabilities *of order three* (higher than the information set).

At first sight, one might think that if we observed *all* bond prices and *all* CDS spreads, thus learning $P(A_i)$ and $P(A_i \cap A_j)$ for each i and j, we would be able to completely pin down the systemic probabilities P_1, P_2 and P_3. However, this is not the case: in general, an information set of order M cannot determine probabilities of order greater than M. A simple graphical example of this is reported in Figure 2, a Venn diagram in which areas represent probabilities. In that Figure, the area of each event is the same across the two panels, so that the marginal probabilities of defaults are the same. The same is true for the pairwise default probabilities. However, it is easy to see that P_3, the intersection of all three events, is positive in the left panel and zero in the right panel.

Knowledge of the low-order probabilities, however, allows us to put bounds on higher-order probabilities, and therefore on systemic default risk. While finding the upper bound for P_3 is immediate[19] ($P_3 \leq 0.01$), finding the other bounds is more complicated, and especially so when there are more than three banks in the financial sector. The exact way to obtain tightest bounds is the object of the rest of this section. When applied to this example, it yields the following bounds:

$$0.45 \leq P_1 \leq 0.46$$

$$0.13 \leq P_2 \leq 0.15$$

$$0 \leq P_3 \leq 0.01$$

This simple example already shows one of the main points of this analysis: the set of bonds and CDSs, when counterparty risk is taken into account, represents a rich information set that can be used to learn about systemic risk even when we make *no assumptions* on the way these low-order probabilities aggregate at the higher-order level.

This example can be used to illustrate two additional concepts related to the measurement of systemic risk. The first one is that simply averaging the CDS spreads of financial institutions can lead to an erroneous measure of systemic risk. Using the notation introduced above, such a calculation would measure systemic risk using an index constructed as:

$$\frac{1}{6}\sum_i \sum_{j \neq i} z_{ij} = \frac{(1-R)}{3}\left[\sum_i P(A_i) - (1-S)\sum_{i,j<i} P(A_i \cap A_j)\right] \tag{3}$$

Suppose that banks become more correlated (the joint default probabilities increase) while the marginal default probabilities remain the same or increase only slightly. Then, this index in fact decreases. The reason is that, because of counterparty risk, systemic risk reduces the quality of

[19]Because we know $P(A_1 \cap A_2 \cap A_3) \leq P(A_1 \cap A_3) = 0.01$, and it is easy to see using Venn diagrams how this bound can be attained.

insurance, and therefore the average cost of insurance decreases. At least in some cases, then, this simple index fails to capture systemic risk properly.

The second idea that emerges from this example is the importance of using *all* the different prices available to construct the bounds, rather than using only the information contained in the *average* bond and CDS spreads. While information on average default probabilities is enough to construct some bounds on the probability of systemic events, the additional restrictions given by the different prices can significantly tighten these bounds. Later in this section I prove that among all networks with the same average marginal and pairwise default probabilities, the widest, least informative bounds on systemic risk are obtained when the network is symmetric, i.e. when all institutions have the same marginal and pairwise default probabilities. Therefore, asymmetry always results in more informative bounds.

Following the example above, suppose that instead of observing all the marginal and pairwise default probabilities we only observe the average probabilities (this is a *partially aggregated* information set): $\frac{1}{3}\sum_i P(A_i) = 0.2$, and $\frac{1}{3}\sum_{i,j<i} P(A_i \cap A_j) = 0.05$. In this case, the upper bound on P_3 is $P(A_1 \cap A_2 \cap A_3) \leq 0.05$. In fact, this bound is attained by a symmetric network configuration in which $P(A_1 \cap A_2 \cap A_3) = P(A_i \cap A_j) = 0.05$ for all i and j: all joint default events (of two or three institutions) perfectly overlap at the upper bound. Using the *full information set* (that includes all bond and CDS prices) rather than the partially aggregated one allows us to reduce the upper bound on P_3 from 0.05 to 0.01. Similar improvements are seen in the other degrees of systemic risk.[20]

3.2 General theory of the probability bounds

In this section, I show how to represent and solve the problem of constructing tightest bounds for probabilities of high-order events given a low-order information set. For now, I assume that we have already extracted the probability information set from the observed prices of traded securities.

Consider a finite set of *basic events* $\mathcal{A} = \{A_1, ..., A_N\}$, which in this paper I interpret as the default events of a set of N financial intermediaries. The relation between low-order probabilities (probabilities of unions and intersections of a few events in $\mathcal{A}$) and higher-order ones (that involve many events in $\mathcal{A}$) has previously been explored in mathematics. Two famous results are Boole's and Bonferroni's inequalities, which state that:

$$P\left(\bigcup_i A_i\right) \leq \sum_i P(A_i) \tag{4}$$

and

$$P\left(\bigcap_i A_i\right) \geq \sum_i P(A_i) - (N-1) \tag{5}$$

These bounds are not tight, in the sense that tighter inequalities can be written based on the same information set (the set of all marginal probabilities of events in $\mathcal{A}$).

[20] In fact, we can improve on the probabilities of default event of *all* degrees $r = 1, 2, 3$. Under average information, the bounds are $0.45 \leq P_1 \leq 0.50$, $0.05 \leq P_2 \leq 0.15$, and $0 \leq P_3 \leq 0.05$.

As discussed in the introduction, I define a systemic event of *degree* r to be the default of at least r out of the N intermediaries. Systemic events defined in this way are *all of order* N, because they involve unions and intersections of all the events in $\mathcal{A}$. I employ an estimation method that allows me to obtain the tightest possible bounds for the probabilities of systemic events given the available information set, which includes bond prices and CDS spreads. The approach is based on linear programming (LP), and consists in writing the bounds as the solution to a LP problem.[21] While difficult to solve analytically, a LP problem is easy to solve numerically even as the scale of the problem gets large. Additionally, the linearity of the problem guarantees that the global optimum is always found when solving it numerically.

To see how the LP approach works, for a sample space Ω consider the finest partition of Ω created by unions and intersections of the basic events $A_1, ..., A_N$; call it V. Then, the probability of each union or intersection of the basic events can be expressed as the sum of the probabilities of some events in V. Since V contains exactly 2^N elements, it is possible to represent the probability space by a vector with 2^N elements, each corresponding to the probability of an elementary event in V . Formally, the following proposition holds (see Boros and Prekopa (1989)):

Proposition 1. *Call $\mathcal{F}$ the σ-algebra generated by the finite set of events $A_1, ..., A_N$ on a sample space Ω. Call V the finest partition of Ω that is included in $\mathcal{F}$. Then, V has 2^N elements, and any probability system on $(\Omega, \mathcal{F})$ can be represented by a vector $p \in \mathbb{R}^{2^N}$, in the sense that $\forall A \in \mathcal{F}$, $\exists I_A \subseteq \{1, 2, 3, ...2^N\}$ s.t. $P(A) = \sum_{i \in I_A} p_i$.*

In general, there are different vectors p that represent the same probability system. In this paper, I use the one constructed according to the following Proposition.

Proposition 2. *For a set B, call $\overline{B} \equiv \Omega \backslash B$, the complement of B. For every integer i between 0 and 2^{N-1}, consider its binary representation b_i, which consists of a vector of N numbers, each either 0 or 1. Construct p_{i+1} as follows:*

$$p_{i+1} = P(A_1^* \cap A_2^* \cap ... \cap A_N^*)$$

where $A_j^* = A_j$ if element j of b_i is 1, and $A_j^* = \overline{A}_j$ if element j of b_i is 0. Then, p represents a probability system on $\mathcal{F}$ in the sense of Proposition 1.

A simple example can help illustrate this Proposition. In the case of three banks, we have three basic events, A_1, A_2 and A_3, corresponding to the default of each bank. The finest partition of the sample space obtained from unions and intersection of these events will have $2^3 = 8$ elements. Figure 3 shows the 8 elements of this partition. From the Figure, it is evident how one can express the probability of any union or intersection of the A_i's as the sum of the probabilities of a subset of these 8 elements. These 8 probabilities can be collected in a vector p with 8 elements, and therefore the probability of any event A' in $\mathcal{F}$ can be represented as a product $a'p$ for a certain vector a. The consistency of the probability system p is assured by imposing $p \geq 0$ and $i'p = 1$, where i is a vector of ones.

[21] See Kwerel (1975). A LP problem is a constrained maximization problem in which both the objective and the constraints are linear in the maximization variable.

The ordering of the elements of p is arbitrary, and Proposition 2 shows a way to construct a vector p that leads to a unique choice for the order of its elements. The i^{th} element of the vector p is obtained as follows. First, obtain the binary representation of the number $i-1$, b_i. For example, $b_1 = [0\,0\,0]$, $b_2 = [0\,0\,1]$, $b_3 = [0\,1\,0]$ and so on up to $b_8 = [1\,1\,1]$.

Each of these vectors can be interpreted as a vector of indicators of one of the three basic events. For example, $[0\,1\,1]$ represent the event in which A_1 does not occur, A_2 and A_3 occur. The element i of p_i will then be the probability of the event represented in this way by b_i. For example, $p_3 = Pr\{\overline{A}_1 \cap A_2 \cap \overline{A}_3\}$. This is precisely the ordering represented in Figure 3.

Propositions 1 and 2 imply that bounds on the probability of a (systemic) event A' in $\mathcal{F}$ subject to constraints on low-order probabilities can be rewritten as a linear programming problem. In particular, the following Corollary holds:

Corollary. *The upper bound for the probability of $P(A')$, i.e. the solution to:*

$$max\, P(A') \tag{6}$$

s.t.

$$P(A_i) = a_i$$

$$...$$

$$P(A_i \cap A_j) = a_{ij}$$

can be found as the solution to the problem:

$$max_p\, c'p \tag{7}$$

s.t.

$$p \geq 0$$

$$i'p = 1$$

$$Ap = b$$

for c, A, b depending only on the available information. The lower bound is obtained by solving the corresponding minimization problem.

Proof. The corollary is an immediate consequence of the fact that the probability of every union or intersection of events in $\mathcal{A}$ can be expressed as a product $a'p$ for some a. □

As mentioned before, the definition of systemic event I employ is indexed by r: "at least r institutions default" within a specified time period. Since all these events are within the $\sigma-$algebra generated by the basic events $A_1, ..., A_N$, their probability can be represented by the product of p with some vector c_r: $c_r'p$.

For the simple case of three banks reported above, it is easy to verify looking at Figure 3 that the probabilities P_r of at least r institutions defaulting can be expressed as

$$P_1 = [0\,1\,1\,1\,1\,1\,1\,1] \cdot p$$

$$P_2 = [0\,0\,0\,1\,0\,1\,1\,1] \cdot p$$

$$P_3 = [0\,0\,0\,0\,0\,0\,0\,1] \cdot p$$

and that the constraints can be rewritten as:

$$P(A_1) = [0\,0\,0\,0\,1\,1\,1\,1] \cdot p = a_1$$

$$P(A_1 \cap A_2) = [0\,0\,0\,0\,0\,0\,1\,1] \cdot p = a_{12}$$

and so on. These constraints can then be collected in a matrix A and a vector b, obtaining the linear representation above. A detailed description of the setup of the LP problem is reported in Appendix B.

3.3 Properties of the bounds

In this section I discuss two important properties of the bounds. First, I prove in Proposition 3 that symmetric probability systems attain the widest bounds on systemic risk given average marginal and pairwise probabilities. Second, I study the uniqueness of the solution. Analytical results for the width of the bounds, in the special case of symmetric systems, are derived in Appendix B.

3.3.1 Symmetry of the probability system

Definition. Consider the vector $p \in \mathbb{R}^{2^N}$ representing a probability system on the σ-algebra generated by the basic events $A_1, ..., A_N$, as in Propositions 1 and 2. Consider a permutation J of the indices of the basic events: $A_{J_1}, ..., A_{J_N}$, and call $\mathcal{M}$ the set of permutations. Define $p_J \in \mathbb{R}^{2^N}$ the vector representing the probability system generated by $A_{J_1}, ..., A_{J_N}$ that corresponds to p, constructed as in Proposition 2.

For example, take two events A_1 and A_2. A vector p representing the probability system constructed as in Proposition 2 would have four elements: $p_1 = P(\overline{A}_1 \cap \overline{A}_2)$, $p_2 = P(\overline{A}_1 \cap A_2)$, $p_3 = P(A_1 \cap \overline{A}_2)$ and $p_4 = P(A_1 \cap A_2)$. In this case, only one additional permutation of the generating events is possible, $J = \{2, 1\}$, with $p_{J1} = p_1$, $p_{J2} = p_3$, $p_{J3} = p_2$, and $p_{J4} = p4$.

Definition. A linear combination of the elements of p defined by the vector c is *symmetric* with respect to the generating events $A_1, ..., A_N$ if $c'p = c'p_J \;\; \forall J \in \mathcal{M}$.

An example of a symmetric weighting vector c is the one corresponding to the probability of the union of the events, $c = [1\,1\,0\,1]'$, since $c'p = c'p_J = P(A_1 \cup A_2)$.

Definition. A probability system p is *symmetric* if every event in V (the finest partition of the sample space generated by the basic events) has the same probability in all permutations of the generating events.

For example, with three generating events ($N = 3$), the probability system is symmetric if $P(A_1) = P(A_2) = P(A_3)$ and $P(A_1 \cap A_2) = P(A_1 \cap A_3) = P(A_1 \cap A_3)$.

Definition. A linear programming problem

$$max\, c'p$$

$$s.t.\, Ap \leq b$$

is *symmetric* if c and all rows of A are *symmetric* with respect to the generating events $A_1, ..., A_N$.

We can now state the following proposition:

Proposition 3. *Suppose that the probability bounds correspond to a* symmetric *LP problem. Then, the bounds are attained by a* symmetric *probability system.*

Proof. See Appendix B. □

Corollary. *The bounds on systemic events of the type "at least r institutions default" given a symmetric constraint set (for example, constraints on the* average *marginal and pairwise default probabilities) are attained by a symmetric probability system.*

The bounds obtained in a symmetric network in which we observe all marginal and pairwise probabilities will always be at least as wide as those obtained in an asymmetric network with the same averages of the low-order probabilities. The difference between the bounds obtained in the two cases captures precisely the extent to which asymmetry in the shape of the network affects the probability of systemic events.

3.3.2 Uniqueness of the solution

At first sight, it might seem that if a financial network is asymmetric enough (relative to the observed marginal and pairwise default probabilities), there will be a unique probability system that attains the upper bound on systemic risk, and similarly a unique one that attains the lower bound. However, this is in general not the case: several probability systems exist that attain each bound. The existence of multiple solutions to the maximization problem plays an important role in section 5, where I study the contribution of different institutions to systemic risk.

If more than one probability system attains the upper (lower) bound, it is possible that some probability $P(A')$ of an event A' is not completely pinned down at the bound. In this case, we can characterize the range of possible values for this probability by solving a second maximization/minimization problem of the type:

$$max_p (min_p)\, v'p$$

s.t.

$$p \geq 0$$

$$i'p = 1$$

$$Ap = b$$

$$c_r'p = c_{bound}$$

where v is the vector that corresponds to the event A' (as described in Propositions 1 and 2), and c_{bound} is the value of the upper or the lower bound on systemic risk. In other words, we look for the probability system, among those which attain the upper or the lower bound, that maximizes (minimizes) the particular probability we are interested in, $P(A')$. Of course, if $max_p v'p = min_p v'p$, that particular probability is completely pinned down at that bound on systemic risk.

In addition, we can also try to find a "representative" probability system at the upper (lower) bound. One way to construct such a system is the "AFROS" procedure described in Appa (2002) and reported in Appendix B. The "representative" probability system is obtained averaging different probability systems, chosen from the space of solutions to the bounds that are as distant as possible from each other. In other words, then, it is an average of solutions that are as different from each other as possible within the space of solutions to the probability bounds.

4 Implementation

The method presented in section 3 requires conditioning on marginal and pairwise default probabilities for all banks and all pairs of banks. When working with actual data, however, estimating the bounds requires some additional steps. First, in order to extract marginal and pairwise default probabilities from observed prices, I need to specify a pricing model for bonds and CDSs that takes into account not only default risk, but also other important determinants of prices - in particular liquidity premia, shown to be an important factor by several studies.[22] Second, the implementation of the bounds is affected by the availability of CDS data. Because I only observe, for each bank, the average CDS spread quoted by its counterparties, the bounds I can estimate condition on a smaller information set, which captures average counterparty risk. I now discuss each of these issues in turn.

4.1 Pricing bonds

To price bonds, I use a simple pricing model of the reduced-form class with constant risk-neutral hazard rates of default, as in Lando (1997), Duffie and Singleton (1999), and Hull and White (2000, 2001).

In general, four elements are crucial in determining the price of a bond: credit risk, the recovery rate, the risk-free rate process and the liquidity premium. Assume that the recovery rate process is independent of all other processes and call the expected recovery rate R. Call T the maturity of the bond and short term r_t^F the riskless rate. The reduced-form approach specifies a risk-neutral default hazard process h_t, the risk-neutral probability of default between t and $t+1$ conditional on survival until t. Following Duffie (1999), I incorporate a liquidity process $\gamma_t > 0$. This process is modeled as a per-period proportional cost of holding the bond. Later I discuss how we can interpret this parameter in light on the theoretical literature on liquidity in asset markets.

I use a simplified pricing model[23] that assumes that, for any given firm, the prices of all of its bonds are determined independently at each time t, under the assumption that from time t onwards h_{t+s} and γ_{t+s} will be constant and equal to h_t and γ_t, respectively. Naturally, this is just an approximation, because prices do not take into account that at each future date these parameters

[22]For example, see Bao et al. (2010), Chen, Lesmond and Wei (2007), Collin-Dufresne, Goldstein and Martin (2007), Huang and Huang (2003), or Longstaff et al. (2005).

[23]I report the general discrete-time pricing model in Appendix C.

are going to be revised, since at every future date $t+r$ prices will be recomputed assuming a constant hazard rate and liquidity process from $t+r$ on, at new levels h_{t+r} and γ_{t+r}. I discretize the model to a monthly horizon, and I assume that coupons are paid monthly. The choice of a month is motivated by the relative reference period for the CDS spreads discussed in section 2.

Theoretically it would be possible to specify and estimate more sophisticated models for h_t^i and γ_t^i: at each time t we observe several bonds outstanding with different maturities, and therefore we could extract information about the term structure of h_t^i at each point t looking forward. The reason why I choose a much simpler model as a baseline case comes from limitations in the CDS data. Because CDS quotes obtained from the main data vendors are reliable only for the contracts of 5 year maturity, I do not have enough information to identify a corresponding term structure of joint hazard rates at every time t for CDSs. The assumption of constant hazard rates allows me to identify the marginal and joint hazard rates directly combining bond and CDS spreads.[24]

Calling $\delta(t,T)$ the risk-free discount rate between times t and T (the price at time t of a risk-free zero-coupon bond with maturity T), the price at time t of a senior unsecured bond j issued by firm i, with coupon c^{ij} and recovery equal to a fraction R of the face value of the bond is:

$$B^{ij}(t,T^{ij}) = c^{ij}\left(\sum_{s=t+1}^{T^{ij}} \delta(t,s)(1-h_t^i)^{s-t}(1-\gamma_t^i)^{s-t}\right) +$$

$$+\delta(t,T^{ij})(1-h_t^i)^{T^{ij}-t}(1-\gamma_t^i)^{T^{ij}-t} + R\left(\sum_{s=t+1}^{T^{ij}} \delta(t,s)(1-h_t^i)^{s-t-1}(1-\gamma_t^i)^{s-t-1}h_t^i\right) \tag{8}$$

Before tackling the problem of calibrating the process γ_t^i (later in this section), it can be useful to discuss a possible interpretation for this variable. The role of γ_t^i in the bond pricing formula is to capture in a reduced-form way all the different elements that result in a liquidity discount in bonds, which can arise for a variety of reasons (such as funding costs, search costs and other transaction costs, or asymmetric information). Funding costs, in particular, are an appealing motivation for modeling the liquidity process because they are known to have been especially relevant during the recent crisis.

Garleanu and Pedersen (2010) present a model in which liquidity discounts arise because some market participants are required to use part of their own equity in purchasing the assets. In particular, they are required to post a margin m_t^{ij} on security j of firm i at time t. Because this uses up part of their own capital, they require an additional return that is proportional to the product of m_t^{ij} and ψ_t, the shadow cost of funds at time t. This, in the presence of a group of traders for which the financing constraint is binding, leads to an adjusted CCAPM of the form:

$$E_t[R_{t+1}^{ij} - R_{t+1}^f] = -\frac{Cov_t(M_{t+1}, R_{t+1}^{ij} - R_{t+1}^f)}{E_t[M_{t+1}]} + m_t^{ij} x_t \psi_t \tag{9}$$

[24] A way to tackle this problem under additional assumptions (essentially imposing that the term structure of joint default follows the term structure of marginal default hazard estimated from bond prices) is explored in section 6 and Appendix D.

where M_{t+1} is the CCAPM stochastic discount factor and R^{ij}_{t+1} is the return on the bond, which includes the potential liquidity discount that might arise in the future. The liquidity discount also depends on x_t, the proportion of liquidity-constrained agents in the economy. Note that since all bonds I consider have the same seniority, we can assume that $m^{ij}_t = m^i_t$ for all the bonds issued by the same firm.

In the simple pricing model presented here, the term γ^i_t approximately corresponds to $m^{ij}_t x_t \psi_t$.[25] In light of this, the liquidity component γ^i_t can be interpreted as a time-varying shadow cost of the capital that needs to be put as margin on the bond issued by firm i. Variation of γ^i_t over time can be attributed to changes in the margin requirements specific to firm i or to variations in the economy-wide weighted shadow cost of capital, $x_t \psi_t$.

4.2 Pricing CDSs

As explained in section 2, a potentially important component of the spread of a CDS is counterparty risk.[26] Counterparty risk arises because in some states of the world the protection seller cannot pay the buyer the full amount owed. A fraction of that amount can still be recovered thanks to collateralization and to the seniority of CDS claims in bankruptcy relative to junior claims.

As with bonds, I discretize the model to one month intervals, and I assume that both marginal and joint hazard rates are constant from the perspective of the time of the pricing until the maturity of the CDS (5 years). I assume that the payoff of the CDS for the following month is as follows. If the seller does not default within the month but the reference entity defaults, the payment is made in full. Therefore, a month is considered an amount of time sufficient for the seller to establish whether to default or not on the CDS obligation, given that the reference entity defaulted. If the seller defaults within the month but the reference entity does not, the contract terminates with either a positive or a negative value (depending on whether the default probability of the reference entity increased or decreased relative to when the contract was written). Here I assume that the expected value of the contract conditional on the reference entity surviving until the next month is zero.[27]

If both the seller and the reference entity default in the same month, I assume that the two defaults happen in a connected way and only an amount S of the full payment is recovered - this case corresponds to the *double default* case, in which counterparty losses are important. The pricing formula remains approximately the same if, conditional of both firms defaulting in the same month, the default of the seller induces a jump in the order of magnitude of the probability of default of the other institution for a certain amount of time; as explained in section 2, double default does

[25] Apart from a term $\frac{E_s[M_{s+1}]-E[M_{s+1}]}{E[M_{s+1}]}$ which captures fluctuations in the short-term risk-free rate and should be relatively small and not very volatile.

[26] The literature on pricing credit derivatives is very large, and examples include Das and Sundaram (2000) and Duffie (1999). Hull and White (2001) and Jarrow and Turnbull (1995) in particular consider models of credit derivatives subject to counterparty risk.

[27] This is consistent with the assumption of constant hazard rates, in which the credit risk of the reference entity is constant over time as long as it does has not defaulted. The pricing formula remains approximately the same as long as the hazard rate of default of the reference entity, conditional on surviving until the next month, remains of the same order of magnitude as it was before the default of the counterparty. In this case, the effect on the price of the CDS is of the order of magnitude of the square of the CDS spread, which is very small.

not need exact simultaneity of the defaults. Note that different assumptions about the jump in probability when the seller defaults conditional on both banks defaulting in the same month can be mapped into different recovery rates in case of double default, S. Robustness to assumptions about S is explored in section 6.

Calling $P(A_i)$ the (constant) monthly default probability of institution i, $P(A_i \cap A_j)$ the probability of joint default during each month, and under the assumptions that these hazard rates are constant over time and independent of the risk-free rate process, the discretized CDS pricing equation can be written as:

$$\sum_{s=t}^{T-1} \delta(t,s)(1 - P(A_i \cup A_j))^{s-t} z_{ji} = \tag{10}$$

$$= \sum_{s=t+1}^{T} \delta(t,s)(1 - P(A_i \cup A_j))^{s-t-1} \left\{ \left[P(A_i) - P(A_i \cap A_j) \right] (1-R) + S \left[P(A_i \cap A_j) \right] (1-R) \right\}$$

where z_{ji} is the spread of the CDS written by j to insure against i's default.

The left-hand side of the formula represents the present value of payments to the protection seller; they only occur as long as neither a credit event occurred nor the counterparty defaulted. The right-hand side represents the expected payment in case of default. In each period, conditional on both firms surviving until then, there is a probability $P(A_i) - P(A_i \cap A_j)$ that the reference entity defaults while the counterparty has not defaulted, so that the payment of $(1-R)$ is made in full. With probability $P(A_i \cap A_j)$, there is a double-default event, and thus only a fraction S of that payment is recovered. Note that if only the counterparty defaults the contract ends with zero value due to the assumption of constant hazard rates.

Using a linear approximation derived and discussed in Appendix C, it is possible to rewrite the spread as:

$$z_{ji,t} = (P(A_i) - (1-S)P(A_i \cap A_j)) \frac{\left[\sum_{s=t+1}^{T} \delta(t,s) \right] (1-R)}{\left[\sum_{s=t}^{T-1} \delta(t,s) \right]} \tag{11}$$

This representation is linear in the event probabilities and can then be imposed directly as a constraint in the LP problem.

While the model I use in this paper uses a simple discrete-time framework, it would be possible to build continuous-time models of CDS prices that take into account the exact dynamics of events and the timing of defaults. Several models of joint default that appear in the literature assume that defaults in each (short) period Δt are independent conditional on the realization of a state vector. In these models, counterparty risk at short horizons is very small by construction, which makes these models less suitable to model the case in which joint default is relevant even at short horizons. A valid alternative are models of correlated default intensities, as introduced by Jarrow and Yu (2001). In these models, the default of one institution increases immediately the default intensity of the others. Because of limited availability of CDS pricing data, I do not have enough flexibility to estimate these models directly. However, my discrete-time pricing formulation is compatible with a model in which defaults within a month are correlated due to spillovers from one institution to the other, but, conditional on one institution surviving until the next month, the spillover effect on the default intensity of other banks is small or negligible.

4.3 Implementation of the bounds: liquidity assumptions

Once the pricing models for bonds and CDSs are specified, the implementation of the bounds requires dealing with the presence of the bond liquidity process γ_i^t, which is unobservable but a crucial determinant of bond prices. Because γ_i^t is unobservable, and because there is no unique way of interpreting it, estimating the liquidity process directly is extremely difficult. However, it is at least possible to obtain plausible *lower bounds* for it, which translate into upper bounds for the marginal probabilities $P(A_i)$.

I start by assuming that γ_t^i can be decomposed into a fixed firm-specific component α_i, and a time-varying component, λ_t, common to all senior unsecured bonds issued by large financial institutions:

$$\gamma_t^i = \alpha_i \lambda_t \tag{12}$$

where λ_t is a latent variable normalized to be 1 on average during 2004 (the beginning of the sample), so that α_i captures the average liquidity component of each bank in 2004. This formulation is flexible enough to capture constant differences among firms in my sample, as well as changes in margins and other liquidity-related costs that are common to the firms in the sample even though they are different from the rest of the economy.

This decomposition, together with the interpretation of γ_t^i discussed above, suggests three possible ways to impose a plausible lower bound on γ_t^i for the financial institutions in my sample. The first approach just requires that the liquidity premium for bonds should not be negative: $\gamma_t^i \geq 0$.

The idea behind the second approach is that we can use the early part of the sample (2004) to identify α^i, because in this period counterparty risk was considered to be essentially zero. Then, using bond prices together with CDS prices allows us to perfectly identify γ_t^i during this period: the average basis in 2004 pins down α^i for each i. If we are willing to assume that liquidity premia were no lower during the crisis than they were in 2004, we can impose an alternative constraint $\gamma_t^i \geq \alpha^i$.

A third, more sophisticated approach tries to obtain a time-varying lower bound for the liquidity process, γ_t^i, by comparing the financial institutions in the sample to other non-financial institutions with high credit ratings and therefore likely similar margins and cost of funding. A CDS written by a financial institution on a safe non-financial firm is much less likely to be affected by the *risk of double default*. Under this assumption, I proceed as follows. For a set J of nonfinancial firms with high credit rating, I estimate γ_t^j using bond yield spreads and CDS spreads together and assuming no counterparty risk. Note that assuming independence yields very similar results, since the order of magnitude is the square of the CDS spread and therefore extremely small. I then decompose it as

$$\gamma_t^j = \alpha_j \lambda_t^* \tag{13}$$

therefore allowing the component common to nonfinancial firms (λ_t^*) to be different from the common component of financial firms, λ_t. Assuming that γ_t^j is observed with independent proportional noise

ϵ_t^j , i.e. we observe:

$$\tilde{\gamma}_t^j = \gamma_t^j \epsilon_t^j \tag{14}$$

we can then estimate the series λ_t^* for each t using OLS (again, normalizing λ_t^* to be 1 on average in 2004). Since this series captures the cost of funds as well as the margin requirement of the bonds of these non-financial institutions (relative to the pre-crisis level), for a group J of high credit rating nonfinancial firms it is reasonable to assume that $\lambda_t \geq \lambda_t^*$. In other words, the liquidity component common to financial firms was, during the crisis, at least as high as the component common to nonfinancial firms. We then obtain a third possible constraint on the liquidity process: $\gamma_t^i \geq \alpha^i \lambda_t^*$.

Once a lower bound $\underline{\gamma}_t^i$ is obtained in this way for each i, from the bond pricing equation we immediately obtain an upper bound on $P(A_i)$, which is the value of $P(A_i)$ that is estimated from bond prices when $\gamma_t^i = \underline{\gamma}_t^i$. This defines an upper bound function $h_i(\underline{\gamma}_t^i)$. We can then modify the maximization problem to find the bounds on systemic risk by replacing at each t the constraints $P(A_i) = \alpha_i$ with the inequality constraints (one for each i)

$$P(A_i)_t \leq h_i(\underline{\gamma}_t^i) \tag{15}$$

This allows to preserve the LP formulation in computing the bounds. On the other hand, it will result in wider bounds, since equalities have been replaced by inequalities.

A final caveat with the use of these liquidity assumptions is that in some periods, and for some banks, the upper bound on the marginal default probability $h_i(\underline{\gamma}_t^i)$ might be lower than the CDS-implied default probability (the one obtained under the assumption that counterparty risk is nonnegative). In other words, after removing the component of the yield spread attributed to liquidity, the remaining part of the yield spread can be lower than the CDS spread. For example, when the liquidity process is set so that the average basis in 2004 is zero, in 2004 half of the banks (on average) will have a positive liquidity adjusted basis. In these cases I reduce the effect of liquidity to the point where bond-implied probabilities are as high as the CDS-implied probabilities. This in turn means that the default correlation of such banks with the rest of the financial system is zero: the whole basis (in fact, even slightly more) is explained by liquidity premia, and therefore there is no room for counterparty risk. This phenomenon occurs less and less frequently as the financial crisis unfolds and the basis widens for more banks. Under the calibration of liquidity to the level of 2004, half of the banks have a zero or positive liquidity-adjusted basis on average between 2007 and Bear's collapse, a fifth of the banks between Bear Stearns' and Lehman's default, and about 2 banks on average after that. Under the calibration of liquidity to the basis of nonfinancials, which is higher, we have a zero or positive liquidity-adjusted basis for two thirds of the banks between 2007 and Bear, about half of the banks between Bear's and Lehman's collapse, and again 2 after that.

4.4 Availability of CDS data and choice of the set of intermediaries

An important factor to take into account in implementing the LP problem is the fact that I do not observe the various spreads written on a given bond i by every other institution j, but only an average of the quotes provided by the $N-1$ counterparties:

$$\bar{z}_i = \frac{1}{N-1} \sum_{j \neq i} z_{ji} \tag{16}$$

This means that when I compute the bounds, instead of the set of constraints generated from individual counterparties shown in (11), I can only impose the constraint

$$\bar{z}_i = \left[P(A_i) - (1-S) \left(\frac{1}{N-1} \sum_{i \neq j} P(A_i \cap A_j) \right) \right] \frac{\left[\sum_{s=t+1}^{T} \delta(t,s) \right] (1-R)}{\left[\sum_{s=t}^{T-1} \delta(t,s) \right]} \tag{17}$$

for each i.

It is important to note that this constraint assumes that the spread $\bar{z}_i$ is obtained by averaging across the spreads quoted by all other $N-1$ institutions in the group considered. The CDS spreads I use (obtained from Markit Group) are constructed as an equal-weighted average of the quotes reported by a set of dealers. Unfortunately I do not observe exactly which dealers contributed quotes at each point in time. Because of this, I consider a group of dealers that are most likely to represent the sample from which the quotes come from. Since this market is very concentrated, and the top 10 firms alone account for about 90% of the protection sold by volume, and for at least 2/3 by trade count, including all the largest dealers according to such measures of activity should ensure that the average spread reflects the average counterparty risk of these financial institutions. Since all of these institutions are very active in the CDS market, and therefore approximately equally likely to contribute quotes, equal weighting seems a reasonable assumption even if not all firms contribute quotes at all times.

Of course, it is possible that the spread partly reflects quotes obtained from financial institutions outside the group I consider, or that some of the dealers in the group do not post quotes at all times. In both cases, I would likely underestimate counterparty risk: in the former case, because quotes may be obtained from smaller institutions for which the recovery rate of the CDS in case of double default could be lower; in the latter case, because if the institutions that are not posting a quote are the riskier ones, the average spread observed would be biased upwards. However, as long as these problems affect only a few institutions at a time, the effect on the average spread should be small.

To find the most active dealers during the crisis, I employ a list of the Top 15 dealers by activity in July 2008 provided by Credit Derivatives Research. While Bear Stearns could not be a part of that list (it had already been bought by JP Morgan), a report by Fitch Ratings[28] shows that it was an important player in the CDS market in 2006, and therefore I include it in the sample. I drop HSBC for lack of enough bond data, so that in the end my sample includes 15 banks, 9 American and 6 European.[29] Note that after March 15th 2008 Bear Stearns disappears, and after September

[28] Fitch Ratings, 2006, Global Credit Derivatives Survey.

[29] The banks are: Bank of America, Bear Stearns, Citigroup, Goldman Sachs, Lehman Brothers, JP Morgan, Merrill Lynch, Morgan Stanley, Wachovia, Abn Amro, Bnp Paribas, Barclays, Credit Suisse, Deutsche Bank, UBS. Note that AIG does not appear because it was holding large net positions, but was not one of the main dealers in the CDS market by volume or trade count.

12th 2008 both Lehman Brothers and Merrill Lynch drop out of the group.

The assumption of equal weighting can be relaxed, by allowing the top institutions to be over-represented in the CDS spread. In particular, because we know the ranking of the top 5 institutions by number of contracts written for each year between 2006 and 2010, we can compute bounds in which these institution's weight in the average CDS spreads is higher than the other banks. Section 6 shows that results are robust to this change.

4.5 Feasible bounds

The analysis presented above allows me to reformulate the maximization problem to obtain bounds on systemic risk that take into account all these factors. The optimal bounds that can be computed given the assumptions discussed above are as follows:

$$max\, P_r$$

s.t.

$$P(A_i) \leq h_i(\underline{\gamma}_i) \;\; \forall i \tag{18}$$

$$\left[P(A_i) - (1-S) \left(\frac{1}{N-1} \sum_{i \neq j} P(A_i \cap A_j) \right) \right] \frac{\left[\sum_{s=t+1}^{T} \delta(t,s) \right] (1-R)}{\left[\sum_{s=t}^{T-1} \delta(t,s) \right]} = \overline{z}_i \;\; \forall i \tag{19}$$

which can be represented in linear form as:

$$max_p\, c_r' p$$

s.t.

$$p \geq 0$$

$$i' p = 1$$

$$Cp \leq d \tag{20}$$

$$Ep = f \tag{21}$$

where the constraint (20) corresponds to the set of constraints (18), and the constraint (21) corresponds to (19). These bounds can then be computed separately at each t using the cross-section of bond and CDS prices, as well as the series for $\underline{\gamma}_i^t$.

Note that the bond-implied constraint $h_i(\underline{\gamma}_i)$ can be computed beforehand using the cross section of bonds, as explained in the next section. The bounds can then be calculated taking the upper bound on marginal probabilities, $h_i(\underline{\gamma}_i)$, as given. Instead, the equations derived from CDS spreads, equations (19), do not involve a separate estimation step, and they are directly imposed as a constraint in the maximization (minimization) problem.

4.6 Data

Before turning to the presentation of the results, I present in this section some information about the data used and the estimation method for the marginal default probabilities $h_i(\underline{\gamma}_i)$. The data cover, with daily frequency, the period from January 2004 to June 2010.

For each of the 15 institutions considered, I obtain clean closing prices from Bloomberg for senior unsecured zero and fixed coupon bonds with maturity less than 10 years[30]. These are indicative quotes; however, if the bond is TRACE-eligible, Bloomberg reports the closing price from TRACE, which corresponds to an actual trade[31]. I exclude callable, putable, sinkable, and structured bonds, since their prices reflect the value of the embedded options. I remove all bonds for which I have price information for less than 5 trading days. I consider bonds denominated in five main currencies: USD, Euro, GBP, Yen, CHF. Since Bloomberg data on European bonds is fairly limited, I integrate it with bond pricing data from Markit, whenever it adds at least 5 observations to the price series of each bond.

As the reference risk-free rate, I use government zero-coupon yields, obtained from Bloomberg.[32] An alternative would be to use swap rates (see for example Houweling and Vorst (2005)). However, swap rates contain counterparty risk, and they are indexed to LIBOR (see Sundaresan (1991) and Duffie and Huang (1996)). Because LIBOR is the rate on unsecured loans between banks, it cannot be considered risk-free, especially in the context of systemic risk in the financial system. The empirical results are robust to the use of swap rates, as discussed in section 6.

Using these data, together with a calibration of the lower bound for the liquidity process $\underline{\gamma}_t^i$, for every trading day t I estimate the risk-neutral default probability $h_t^i(\underline{\gamma}_t^i)$ separately for each firm, using the cross section of bonds issued by firm i which are still outstanding at time t. I employ the constant hazard rate model described above, and estimate the hazard rate using least absolute deviations to reduce the impact of outliers. All the results are robust to the use of OLS.

Table 1 reports some statistics on the availability of bond data. In the first column, for each institution, we can see the average daily number of valid bond prices available for the estimation. For example, the default probability for Bank of America h_t^i is estimated using on average 32 bonds each day. The next columns break down this number by year. The Table shows that for some European dealers, bond data is scarce especially in the early part of the sample.

Turning to CDS data, since credit default swaps are derivatives traded over the counter, they are not always standardized. However, the 5-year CDS based on the ISDA format (which standardizes credit event definitions and collateral requirements) has emerged as the reference contract, which brings high transparency and liquidity. For this CDS contract, Markit reports quotes that are obtained by averaging the quotes reported by different dealers, after removing stale prices and outliers.[33] The series contain a few missing values, which are filled by interpolation. The CDSs are euro-denominated for European banks and dollar-denominated for US banks. All the results are robust to the use of CDS data from Bloomberg - CMA New York.

[30] Given that the maturity of CDSs is 5 years, it would make sense to use outstanding bonds of remaining maturity close to 5 years when comparing bonds and CDSs. The results are robust to other maturity brackets around 5 years - for example (2,8) and (4,6), even though this would reduce the number of institutions I can use for the estimation.

[31] Section 6 reports robustness result based only on TRACE prices.

[32] I instead bootstrap the Swiss yield curve from the government coupon yield curve, using linear interpolation and assuming bonds trade at par.

[33] Removing outliers, while helpful to remove noise coming from erroneous prices, can potentially bias the reported CDS spread away from the average spread if the distribution of quotes is skewed. This in particular can be a concern if the distribution of quotes is left-skewed, i.e. most dealers have low counterparty risk but a few dealers have higher counterparty risk, as the observed CDS spread would be biased upwards.

Table 2 reports summary statistics on CDS spreads. While CDS spreads between 2004 and 2010 are usually quite low, on the order of 50bp, they reach levels higher than 1000bp in some periods. On the right side of Table 2, I report statistics for the basis $z_i - (y_i - r^F)$, computed using the interpolated 5-year bond yield and the 5-year Treasury rate. As the Table shows, the basis is usually negative, because the CDS spread is lower than the corresponding bond yield spread. In a few cases, however, the basis becomes positive. This can be due to noise in bond pricing data. It can also occur when, in response to bad news, CDS spreads jump upwards quickly while bond yields react with a lag, because of the lower liquidity of the bond market. Of course, as discussed above, one limitation of the methodology presented in this paper is that all such events will be interpreted as a drop in counterparty risk to zero, at least until the bond yields adjust. However, since this phenomenon typically occurs for only a few days, it will not affect the general behavior of the bounds. Smoothing the bounds will further reduce the problem.

5 Empirical Bounds on Systemic Risk

5.1 Bounds on systemic risk with nonnegative liquidity

In this section, I present results under the assumption that the per-period cost of holding bonds (liquidity process) is non-negative: $\gamma_t^i \geq 0$. Note that because I only impose a lower bound on γ_t^i, but not an upper bound, I allow the bond/CDS basis to be explained by any combination of liquidity and counterparty risk. The bounds are obtained under the assumption that the recovery rate on bonds (R) and that on CDSs in case of double default (S) are both equal to 30%. In section 6 I explore robustness to different assumptions on recovery rates.

Figure 4 plots the bounds for the average monthly probability of default of at least r institutions, P_r, for r between 1 and 4. To make the graphs more readable, I plot a 3-day moving average of each bound. It is useful to compare these bounds, that use all available information, to partial-information bounds obtained using only bond prices or only CDS spreads. This allows one to see the gain from using the information from all the traded securities.

In Figure 4, the thin lines represent the upper and lower bounds obtained using only bond prices. These bounds impose only constraints (18) but not constraints (19) in the maximization and minimization problems. The dotted lines, on the contrary, use only CDS data; they impose only constraints (19) but not constraints (18). Both these sets of bounds ignore the information on counterparty risk contained in the bond/CDS basis. The thick lines represent the *full-information bounds*. They are the optimal bounds conditional on the prices of bonds and CDSs that we observe. Note that the CDS-only bounds (dotted line) often coincide with the full-information bounds, and that all lower bounds for $r > 1$ are zero.

Even under these minimal assumptions on the liquidity process, the bounds are tight enough to be informative about systemic risk. All the bounds in Figure 4 suggest an increase in the maximum possible amount of systemic risk during the financial crisis, up to early 2009, followed by a decrease at the end of 2009. As we can see in the first panel, the probability that at least one bank defaults (P_1) clearly increased during the crisis. However, the other panels do not clearly pin down whether a similar increase occurred for the default risk of more than one institution. All bounds on P_2 to P_4 start to widen in late 2007, which is consistent with an increase in systemic risk, but also with an increase in uncertainty about systemic risk.

Comparing the bounds obtained using different information sets, an interesting pattern emerges. The full-information bounds, that use both bond and CDS spreads, are noticeably tighter than the bounds that only use bond prices, but, in general, do not seem to improve much over the CDS-only bounds. The reason for this is that the upper bound on marginal default probability derived from bond prices ($h_i(0)$) is often so high that constraint (18) is not binding at the maximum of systemic risk. This result arises from the restrictions imposed by the internal consistency of the probability system. In these cases counterparty risk cannot be so high as to explain the whole bond/CDS basis. This shows why often the full-information cannot improve over the CDS-only bounds.

When the bond and CDS spreads increase, but the difference between the two does not increase as much, the constraints imposed by bond prices start to bind. In such situations the full-information bounds become significantly more informative. In other words, the full-information bounds have more bite when the basis is small relative to the spreads. This is precisely what happened during the peak periods of the crisis (Bear Stearns' and Lehman's failures), as can be seen in Figure 4.

The optimal bounds also allow an interesting decomposition of the movements in bond yields and CDS spreads into idiosyncratic and systemic risk. This decomposition emerges clearly if we compare the top panel to the bottom three panels of Figure 4. The top panel indicates that the probability that at least one bank would default (P_1) spiked during these episodes. However, P_2, P_3 and P_4 do not show a similar spike (Bear's case) or show a smaller one (as in the month after Lehman's default). This tells us that during these episodes what spiked was idiosyncratic risk rather than systemic risk.

Note that we cannot reach similar conclusions if we look at the bond-only or CDS-only bounds. These bounds do not capture the difference between the different degrees of systemic risk, and suggest a story in line with the naive measures of Figure 1.

Looking at Figure 4, one might be tempted to conclude that systemic risk *decreased* in the Bear Stearns and Lehman episodes: the lower bound for systemic risk stays at zero while the upper bound tightens. This, however, is not the correct interpretation. The bounds *reveal* that systemic risk was low during these episodes, but they do not establish that it was higher before or after, because at those times the bond/CDS basis was too wide to be informative about systemic risk. In other words, a similar basis is *more informative* about systemic risk at times when CDS spreads and yield spreads are high than in times when they are low.

These results were obtained under minimal assumptions on liquidity. In the next section, I show that imposing stricter assumptions on liquidity makes the basis more informative even in the periods before and after these peaks, when the levels of the spreads are lower.

5.2 Bounds with other calibrations of liquidity

As explained in section 4, if we assume that the liquidity cost of holding a bond is at least $\underline{\gamma}_t^i$, the upper bound on each of the marginal default probabilities decreases to $h_i(\underline{\gamma}_t^i)$. In other words, the part of the bond/CDS basis that can be attributed to counterparty risk is reduced by a corresponding amount. In this section I show how the bounds on systemic risk vary as we make different assumptions about $\underline{\gamma}_t^i$. Note that lowering the upper bound on the marginal probabilities of default has three effects on the measures of systemic risk. First, it directly rules out high values for the individual default probabilities of each bank. Second, it indirectly lowers the maximum amount of counterparty risk present in CDS contracts, by lowering the basis for all banks. Third, to the extent that the effect of liquidity on the basis is different across dealers, it increases the asymmetry of the network and therefore leads to a greater benefit from using all the information available. The bounds constructed using bond prices but not CDS spreads reflect only the first effect. The bounds that use only CDS data are not affected at all by bond liquidity assumptions. The optimal bounds reflect all three effects, and therefore become significantly tighter than both the other bounds.[34]

The first panel of Figure 5 reports the bounds on the monthly probability of at least four institutions defaulting (P_4) when $\underline{\gamma}_t^i = 0$ (which corresponds to the last panel of Figure 4). The second panel plots bounds obtained when $\underline{\gamma}_t^i$ is calibrated to match the bond/CDS basis of each bank in 2004. This calibration implies a liquidity component of the yield spread of about 50-60% in 2004 - in line with the calibrations of Huang and Huang (2003) and Longstaff, Mithal and Neis (2005).

The third panel plots the bounds obtained by calibrating the liquidity process to the one that fits the entire basis of nonfinancial firms, as explained in section 4. In particular, I look at the nonfinancial firms that compose the CDX IG index (a standard index of CDS spreads on investment-grade bonds), restricting to those with credit rating of A1 or higher (according to Moody's). For 8 of these firms[35] I have enough data to include them in the estimation of λ_t^*. This time-varying component stays approximately constant until 2007, and during the crisis it increases by up to five times.

Once stronger liquidity assumptions are taken into account, we discover that systemic risk was particularly low before Bear Stearns' collapse. In particular, the upper bound on P_4 does not reflect a sharp increase in systemic risk at the beginning of 2008, contrary to the other measures of systemic risk discussed in the introduction. The third panel in Figure 5 shows how that increase was due to idiosyncratic, not systemic risk. After jumping in March 2008, systemic risk kept increasing in a relatively smooth way up to April 2009, showing only a small spike immediately following Lehman's default. The large spike that during that period we observe in CDS spreads and bond yields, reaching the highest levels observed in the sample, corresponds to a spike in idiosyncratic default risk rather than in systemic default risk. This decomposition is only possible using all the information available, and is not captured when bounds are constructed using only bond prices or only CDS spreads.

[34] Note that in order to calibrate the liquidity process, we use some information obtained from both bonds and CDS spreads, because we use the bond/CDS basis in 2004 to calibrate the liquidity level α^i, and because we adjust the liquidity process so that the bond/CDS basis of each bank never becomes positive (as discussed in section 4). So the bond-only bounds and the CDS-only bounds, that use either constraints (18) or constraints (19) after such adjustments for liquidity, will indirectly reflect some information contained in the bond/CDS basis. The bounds look very similar if the bond/CDS basis is used only to calibrate the liquidity process in 2004 but not used in the following periods.

[35] The eight firms are Boeing, Caterpillar, John Deere, Disney, Honeywell, IBM, Pfizer, Walmart.

The reason for this decomposition between idiosyncratic and default risk during these key episodes of the crisis is simple. High systemic risk in the financial sector requires low CDS spreads, relative to bond yields, for contracts written by large dealers against the default of other dealers. Since the counterparties to those contracts are in fact these same dealers, a high risk of systemic default ought to reduce the value of insurance contracts bought from them. However, until March 2008, and again during September 2008, the bond/CDS basis was not wide enough to correspond to such spikes in systemic risk. This effect is evident especially when liquidity premia explicitly account for part of the basis.

A final remark concerns the interpretation of this decomposition in terms of objective, rather than risk-neutral, probabilities. We can reasonably assume marginal utility to be higher, in expectation, in systemic default events than in idiosyncratic default events. Therefore, a period in which P_1 is particularly high, but P_4 is low, is likely to be due to high objective idiosyncratic default risk and low objective systemic risk. If the large value of P_1 was due to high expected marginal utility in the corresponding idiosyncratic default events, we would expect to find P_4 to be high as well, in contradiction with the empirical findings.

5.3 Full-information bounds and average-information bounds

An interesting question is to what extent one could achieve similar results on systemic risk by looking only at *average* bond and CDS spreads. As shown in section 3, bounds obtained using only average information coincide with bounds obtained in a fully symmetric system. Therefore, this exercise also reveals the gain in terms of information about systemic risk due to the asymmetry of the financial network.

The last panel of Figure 5 plots the bounds obtained under the full versus the average information set. These are derived under the calibration of liquidity to nonfinancial firms, as in the third panel. The average information set is obtained by replacing the two *sets* of constraints (18) and (19) with one constraint each, obtained by averaging each of the two sets across dealers. The Figure shows that looking at average spreads allows us to distinguish between idiosyncratic risk and systemic risk in some cases (like the weeks before Bear Stearns' failure) but not in others (Lehman's default). The reason is that in the first episode the spreads and bases of most banks moved in a similar way, but in the second episode the risks were concentrated in a few banks.

5.4 Individual contributions to systemic risk

The method described in this paper also allows one to study the evolution of the default risk of each bank and its relation with the rest of the network. In particular, I solve for the probability systems that *attain* the upper bound and study the configuration of the financial network in the scenarios of highest systemic risk. In this section I look at the bounds for P_4 obtained by calibrating the liquidity process to that of nonfinancial firms.

In general, several probability systems attain the upper (and lower) bounds on systemic risk. This means that there will not be a unique configuration of the network at the bounds in terms of default probabilities. In this section I follow the procedure described in section 3 and Appendix B to numerically solve for the "representative" probability system in the space of solutions. This gives us an idea of the typical configuration of the network at the bounds. I also report the maximum and minimum values for all the magnitudes of interest (for example, the pairwise probabilities of default) that can be achieved within the space of solutions whenever these are not completely pinned down at the bounds.

I focus on the marginal and pairwise probabilities of default at the upper bound for P_4. As an example, in Figure 6 I plot a partial snapshot of the network as of August 6th 2008, five weeks before Lehman's collapse. The nodes of the diagram are associated with the individual banks and present monthly marginal probabilities of default. The segments that connect the nodes report the joint default probability of the two intermediaries. Note that the marginal probabilities of default are completely pinned down at the upper bound for P_4. The same is not true for pairwise probabilities. This should not be surprising, since from CDSs we obtain constraints only on average pairwise probabilities. Figure 6 reports the estimate of pairwise default probability in the "representative" solution. In parentheses, I report the range of maximum and minimum possible values within the space of solutions. Even though joint default probabilities are not completely determined, in most cases the range is relatively tight.

In decreasing order of individual default risk we find Lehman Brothers, Merrill Lynch, Morgan Stanley, Citigroup, and JP Morgan. In addition, we can see that the pair at highest risk of joint default is Merrill Lynch with Lehman Brothers, followed by the pair Lehman Brothers and Citigroup. It is evident from this graph that the prices of bonds and CDSs were consistent with a high joint default risk of Lehman and Merrill even 5 weeks before the weekend in which both went under (September 13-14th). The rest of the graph shows considerable heterogeneity in the marginal probabilities of default, but especially so in the pairwise probabilities, which are particularly informative about the structure of the network. For example, from the graph we learn that even though Citigroup has a much smaller individual default probability than Morgan Stanley, the two appear equally connected to Lehman Brothers and Merrill Lynch.

Not plotted in the graph are several banks for which the joint default risk with other banks is zero or approximately zero (in this day: the European banks and Goldman Sachs). For these banks, the basis becomes essentially zero after adjusting for liquidity, and therefore they appear disconnected from the network. Note that the cases of zero joint default probability and independence are difficult to distinguish because the order of magnitude of the latter is the square of the default probabilities, which is extremely small. However, this clearly indicates that there is ample heterogeneity across banks in their maximum possible contribution to systemic risk.

Using a similar approach, for each pair of banks i and j we can track the evolution of $P(A_i), P(A_j)$ and $P(A_i \cap A_j)$ over time. Figure 7 plots a 3 day moving average of these probabilities for three different pairs (all combinations of Lehman, Merrill Lynch, and Citigroup) - in the "representative" solution. The upper panel reports the marginal probabilities, and the lower panel reports the joint probabilities. These graphs confirm the relatively high degree of heterogeneity and variability in marginal default probabilities across banks, but even more so for joint probabilities. It is particularly interesting to note that the markets anticipated the joint collapse of Lehman Brothers and Merrill

Lynch for the two months prior to that event.

We can now turn to study how each institution contributed to systemic risk. One way to capture this is to compute the probability that institution i is involved in a multiple default event:

$$Pr\{at\, least\, 4\, default\, \cap\, i\, defaults\}$$

By applying the techniques described in section 3, we can verify that this probability is uniquely identified at the upper bound on systemic risk. Figure 8 plots this contribution for four banks (Citigroup, Lehman Brothers, Merrill Lynch and Bank of America) as well as the average across the other banks.[36]

As the Figure shows, there is large heterogeneity across dealers, both in the levels and in the changes. While the contribution to systemic risk increases for all banks after August 2007, the growth is faster for Lehman and Merrill Lynch than for the other banks. Markets seem to have anticipated the increased systemic importance of these two institutions before the weekend of September 15th. After September 2008, Citigroup displayed the largest spike, followed immediately by Bank of America, that ended with the release of the results of the U.S. government stress test on banks' capital in May 2009.

6 Robustness

In this section I study the robustness of the main results of the paper to different assumptions. While the main results are easy to read in graphs like Figures 4 and 5, for reasons of space I cannot report all the graphs for each robustness test. Instead, I report the average value of the bounds for different subperiods, chosen to reflect the main events identified in the Figures. For each robustness test (all performed under the calibration of the liquidity to the basis of nonfinancial institutions, $\underline{\gamma}_t^i = \alpha_t^i \lambda_t^*$) I report in Table 3 the average value of the bounds, in basis points per month, during different periods: January to December 2007, January 2008 to March 15 2008 (the run-up to Bear Stearns' collapse), from Bear's episode to Lehman's default (on September 15th 2008), the month after Lehman's default (in which CDS spreads and bond yields spiked), the period between September 2008 and April 2009 (the latest peak of the crisis, just before the stress test results were released) and finally from May 2009 to June 2010. In the three panels, I show values for the lower and upper bound on P_1, and the upper bound on P_4 (the lower bound on P_4 is always 0).

Besides showing the level and the time series of the bounds, this Table allows us to check that the main results reported in the paper hold under different assumptions. The bold line in each panel of Table 3 reports the baseline case presented in the paper. We can confirm the result, presented in Figure 5 (third panel), that systemic risk was low in the months preceding Bear Stearns' collapse, while idiosyncratic risk was already high. Besides, we can see that during the month after Lehman's default, idiosyncratic risk spiked (it increased sharply and then decreased as sharply). The upper bound on P_4 increases as well, but does not even reach the level observed in the following 6 months: systemic risk keeps increasing until March 2009.

[36] To improve readability, I plot a two week moving average.

6.1 Pricing assumptions

6.1.1 Value of the constant expected recovery rates

Let us start with robustness with respect to the assumed recovery rate of CDSs when double default occurs, $S \in [R, 1]$. The effect of changes in this assumption depends crucially on the liquidity-adjusted bond/CDS basis of each bank. For some banks, the basis is small enough that can be completely explained by counterparty risk. For these banks, an increase in S means that the same basis can account for higher counterparty risk. For other banks, instead, the basis is large enough that, due to internal constraints of the probability system, it cannot be completely explained by counterparty risk: even at the upper bound for systemic risk, a part of the basis must be explained by liquidity. For these banks, an increase in S means that the same amount of counterparty risk - which was already at the maximum possible - will explain an even smaller fraction of the basis. This means that the marginal probability of default, $P(A_i)$, has to decrease. In turn, this directly reduces the maximum possible amount of counterparty risk for contracts written by i against other banks, since for each j we must have $P(A_i \cap A_j) \leq P(A_i)$.

An increase in S then has a different effect on banks with a relatively small basis and banks with a large basis. The two effects are also at play for each bank individually, for different starting levels of S: when S is low enough counterparty risk has a large effect on CDS spreads, and therefore the basis will be relatively small - it can be completely explained by counterparty risk. When S is large enough, not all basis can be explained by counterparty risk, and the second mechanism operates. In Appendix B, I show these results formally in the symmetric case, where all banks have the same basis. In the asymmetric case, the two opposite forces counterbalance each other in such a way that the bounds on systemic risk are not very sensitive to assumptions on S. The reason is that in an asymmetric network different banks will have different bases, so that for most values of S the two effects described above will operate for some banks in one direction and for other banks in the opposite direction. This analysis explains why we see the bounds on systemic risk being very robust to changes in S (at least up to a recovery rate of 90%), as shown in Table 3.

The case for the recovery rate of bonds R is different. As shown in equations (8) and (11), R affects the prices of *both* bonds and CDSs. A higher expected recovery rate in case of default increases the value of a bond, and at the same time decreases the value of CDS insurance written on that bond, since the payment from the CDS seller covers only the amount of bond value not recovered in default. Because this recovery rate multiplies the marginal and joint default probabilities in the pricing formulas, when R changes all probabilities implied in bonds and CDSs are scaled up or down by approximately the same amount.[37] Therefore, the bounds on systemic risk will scale in a similar way. However, the main results on the time series of the bounds will not change, as shown by Table 3.

[37]The difference between the two comes from differences in the cash flow timing of bonds and CDSs. They are scaled by exactly the same amount in the simple two-period example of section 2.

6.1.2 Time varying recovery rates

Above I have studied robustness to different assumptions about S and R, when these are assumed to be constant during the whole sample period. In theory, it is possible that these recovery rates vary over time in a way that affects the results on the time-series of systemic risk presented in section 5. Suppose that at every time t bonds and CDSs are still priced assuming that at all future periods $t+s$ the recovery rates are constant and equal to S_t and R_t; however, let now S_t and R_t vary over time. How will this affect the bounds on systemic risk?

The tests presented above show that the bounds on systemic risk scale in the same direction as R. If we believe that, during peak episodes like the one following Lehman's default, recovery rates R might have dropped, this would in fact strengthen the result that the spike in systemic risk was then relatively low, because it would further reduce the bound on P_4 during that month.

Another possibility is a reduction in the recovery rate of CDSs, S, in times when systemic risk increases. However, it is easy to see that this case actually reinforces the main empirical results. If the recovery rate S becomes smaller during the key episodes of the crisis, then joint default risk has to be smaller as well. This stems once more from the fact that during these episodes the bond/CDS basis is small relative to CDS and yield spreads. When S is reduced, the probability of joint default has a *greater* effect on the basis. To still match the basis even if S is higher, joint default risk has to decrease. Therefore, the main results in the paper will be robust to a decrease in the recovery rate S in times of crisis.

6.1.3 Stochastic recovery rate on bonds R

Another possibility is that when pricing bonds and CDSs, agents incorporate the possibility that recovery rates might be stochastic and correlated with the default events in the financial sector. In particular, one could think that recovery rates of both bonds and CDSs might deteriorate the more defaults happen in the financial system.

Because of the limited data available, it is difficult to solve explicitly for the case of stochastic recovery rates. However, it is possible to gain some intuition on the effect of this assumption under simple modeling assumptions. Suppose that the recovery rate on bonds is R_H whenever one bank defaults alone, and $R_L < R_H$ whenever two or more banks default. Then, I show in Appendix D that we can decompose as follows the change in the bounds on systemic risk, going from a non-stochastic recovery rate R to the stochastic recovery process described above. First, we can shift the (constant) recovery rate downwards for both bonds and CDSs to R_L. This component scales down the bond-implied and the CDS-implied probabilities by a similar amount, as discussed above. This would scale the bounds on systemic risk downwards. Second, we increase the present value of bonds by an amount Y_{bond}, and decrease the present value of payments of the CDS contract by an amount $Y_{CDS} \approx Y_{bond}$ (in a first-order approximation with small probabilities of default). This second effect shifts the CDS spread and the yield spread in the same direction by a similar amount, with minimal effect on the basis and hence on counterparty risk. We then expect the bounds on systemic risk to become lower if we introduce a stochastic recovery rate with $R_L < R$. The reason is that for the purpose of systemic risk, the relevant recovery rate is the one that obtains in states of multiple defaults, or R_L in this case. However, as long as the recovery rates R_L and R_H themselves do not vary over time, the time series of the bounds should still look as in Figures 4 and 5.

6.1.4 Assumptions about the hazard rate

In this section I allow for a more flexible form for the hazard rate process. In particular, for each institution i, from the perspective of an agent pricing bonds at time t, the hazard rate at time $t+s$ follows the (deterministic but time-varying) process

$$h_{t+s} = (1-\rho_t)\overline{h}_t + \rho_t h_{t+s-1} \tag{22}$$

where parameters h_t, $\overline{h}_t$ and ρ_t are determined at time t. As before, all bonds are priced at every time t assuming that the hazard rate process is known for all future dates.

At each time t, I can use bond prices to estimate h_t, ρ_t and $\overline{h}_t$. This representation allows to capture cases in which the hazard rate is higher at shorter horizons and then reverts to a lower long-term value. I can therefore construct the bounds on systemic risk using h_t, the probability of default in the month after t.

As discussed in section 4, the main problem with this approach is that while it is easy to estimate a more flexible function for the marginal hazard rate of default using bond prices, CDS data do not contain enough information to estimate a similarly flexible process for joint default risk (because at each time t we only observe the spread of one CDS, with maturity of 5 years). To tackle this limitation, I assume that the joint hazard process replicates the shape of the marginal hazard process of the reference entity: the process decays at the same rate (ρ_t) and displays the same ratio between short-term and long-term default hazards ($h_t/\overline{h}_t$). The details of the estimation method and of the additional assumptions involved are reported in Appendix D. Table 3 shows that the main empirical results are confirmed under these assumptions.

6.2 Other robustness tests

6.2.1 Using interest rate swaps as the risk-free rate

While swap rates may not be the appropriate rate to discount cash flows under risk neutral probabilities (because they are indexed to a risky reference, LIBOR, and because they contain counterparty risk), it is interesting to check how the results change if we use them in place of Treasury rates.[38] Because these rates are higher than the Treasury rates, and therefore result in a lower basis for all banks, we would expect the upper bounds on systemic risk to decrease noticeably. At the same time, remember that we are calibrating the time variation in the liquidity process to the basis of non-financial firms, and the level of the liquidity process to the basis of each bank in 2004. Therefore, the change in the risk-free rate will be offset by a corresponding decrease in the liquidity process (even though the offset is not exactly one to one). Table 3 shows that the change in the bounds is very small.

[38] I bootstrap the zero-coupon yield curve from the par swap rate curve of the different currencies using linear interpolation.

6.2.2 Assumptions about the weighting of contributors in CDS contracts

As discussed in section 4, the bounds are computed under the assumption that the CDS spreads are obtained by averaging quotes obtained from all the other dealers in the sample. If some dealers do not post quotes at all times, the average spread observed will, in expectation, overweight dealers which send quotes in more frequently. In turn, this is most likely related to how active the dealer is in the CDS market.

While we cannot obtain directly estimates of the activity of the dealers (in terms of number of contracts written and volume of CDS protection sold), Fitch Ratings[39] reports a ranking of the top 5 counterparties by trade count (which in turn is very correlated with gross positions sold), for each year between 2006 and 2010. We might then think that because these dealers are more active, quotes are more likely to be obtained from them, and therefore the average CDS spread observed will in expectation reflect more their contribution. Given this, as a robustness test I compute bounds that overweight the top 5 institutions in the formula for CDS contracts. I consider two relatively extreme weighting schemes. In all of them, institutions ranked below 5 have the same weight (I do not have information about the relative ranking of these dealers). In the first weighting scheme, I compute the bounds assuming that the top 5 institutions are 5 times more likely than the other 10 to contribute quotes, and therefore their contribution is weighted 5 times more than the other institutions in the sample. The second weighting scheme again assumes that all institutions ranked 6-15 have the same weight, and the top dealer has 10 times their weight, the second dealer 8 times, and so on up to the 5th largest dealer (with a weight twice that of the smaller dealers).

The effect of this overweighting on the bounds of systemic risk is not immediate. Suppose, for example, that in the bounds computed under equal weighting, systemic risk comes from the joint default risk among top-5 banks. Then, increasing the weight on these banks will have the effect, everything else constant, of lowering the weighted CDS spreads. But this is not possible because the CDS spreads were chosen to match the observed ones. Therefore, the joint default risk among these banks will have to decrease. At the same time, joint default risk with smaller banks can increase. But if these smaller banks were contributing little to default risk before the change in weights, an increase in the possibility of joint default risk with them might not make up for the reduction in maximum systemic risk coming from the top-5 dealers. In this example, systemic risk will likely decrease when we overweight top-5 dealers. It is easy to see that the opposite is true if systemic risk mainly comes from non top-5 dealers.

If instead in both groups (top-5 and non-top-5 dealers) we find dealers with large contribution to systemic risk as well as dealers with small contribution to systemic risk, under equal weighting, the bounds will be relatively robust to changes in the weights. In fact, this is the case. The top 5 banks include both banks with high contribution to systemic risk as well as banks with low contribution to systemic risk, such as one or two European banks. Table 3 shows that under both weighting schemes the main results still hold.

This robustness test also allows us to say something about heterogeneity in collateral agreements across counterparties. All the results in the paper have been derived assuming that the recovery

[39]Fitch Ratings, 2008, Global Credit Derivatives Survey.

rate in case of double default, S, is the same for all banks. How do the main results change if instead (because of different collateral agreements and exposure to other shocks) the recovery rate is different across institutions? While we have no direct information about the expected recovery rates of each counterparty, it is easy to show that if the recovery rates S_j are different across counterparties j, the average quote reflects not an equally weighted average across j's of the joint default probabilities $P(A_i \cap A_j)$, but rather a weighted average $\sum w_j P(A_i \cap A_j)$, where $w_j = \frac{(1-S_j)}{(N-1)(1-\overline{S})}$, and $\overline{S} = \frac{1}{N-1}\sum_j S_j$. Therefore, given a certain average recovery rate $\overline{S}$, the joint default risk with counterparty j will be weighted more in the observed quote if j's recovery rate is lower. Now, it is reasonable to assume that more important counterparties (that have a larger volume of the business) are also the counterparties that are able to obtain less stringent collateral agreements – and therefore buyers of CDSs from them might obtain a lower recovery rate in case of double default. As a consequence, the robustness test presented in this Section can also be interpreted as robustness to this case of heterogeneity in recovery rates.

6.2.3 Assumptions about the exchange rate

The construction of the bounds on systemic risk involves the estimation of risk-neutral probabilities from bond prices and of joint default probabilities from CDS spreads. Using probabilities obtained from different securities to obtain risk-neutral probabilities of joint default requires additional assumptions if the securities are denominated in different currencies. In particular, while most bonds issued by American firms and the CDSs written on them are denominated in dollars, European firms issue several bonds in Euros and in other currencies, and the CDSs written on them are denominated in Euros.

To simplify the discussion, consider one-period bonds and CDSs written by banks i and j. Call m_{se} the stochastic discount factor of a US investor in state (s, e). Here, s indicates the default state of the banks i and j, so that it can take values i (only i defaults), j (only j defaults), ij (both default), and 0 (none defaults). e indicates the exchange rate with a foreign currency. Call π_s the probability of s occurring, and note that $\pi_s E[m_{se}|s]$ is the price of a security that pays 1 if default state s happens. The price of a state-contingent security that pays a unit of foreign currency if default state s happens is then $\pi_s E[e m_{se}|s]$.

Appendix D shows that a sufficient condition for correctly estimating risk-neutral default probabilities using bonds and CDSs denominated in different currencies (using the risk-free rates denominated in the respective currencies to discount cash flows) is:

$$\frac{\pi_s E[e \cdot m_{se}|s]}{\pi_s E[m_{se}|s]} = \frac{E[e \cdot m_{se}]}{E[m_{se}]} \tag{23}$$

which requires that the relative price of domestic and foreign risk-free securities is the same as the relative price of domestic and foreign state-contingent securities that pay off in the various default states. Of course, it is reasonable to assume that the relative price of dollar-denominated and foreign currency-denominated default-contingent securities might be different depending on the default state (think for example of a flight-to-quality to US securities if several banks default). As a robustness test for the validity of the bounds in case these conditions are violated, I perform the estimation exercise including only American firms, for which all bonds and CDSs are dollar-denominated. Table 3 shows that the results still hold for this subset of banks.

6.2.4 Using only larger transactions from TRACE

A concern with using bond prices from Bloomberg is that they might incorporate stale information (for European bonds, for which I use quoted prices), or they might depend on very small trades, which might be less reflective of credit risk (see for example Dick-Nielsen et al. (2010)). To make sure results are robust to these problems, I compute the bounds for the subset of US firms using only transaction data from TRACE, and ignoring all trades with nominal amounts of less than $100,000. Of course, this will exclude several bonds for several days. Table 3 reports that the bounds change very little.

7 Conclusion

This paper shows that bond prices and CDS spreads represent a rich information set for learning about joint default probabilities of intermediaries in the financial network. This information set can be used to construct bounds on the probability of systemic events, defined as the probability of several institutions defaulting within a short time horizon. These bounds are the best possible ones given the available information, and they are obtained as the solution to a linear programming problem.

Even under minimal assumptions on the liquidity premia in bond markets, the bounds obtained in the paper are significantly tighter than the ones obtained from bond prices alone. They are also tighter than bounds that only use CDS spreads, particularly when CDS and bond spreads increase while the bond/CDS basis remains relatively small (such as early 2008 and October 2008). Adding relatively mild assumptions on the bond liquidity process further tightens the bounds. Using this methodology to estimate stress in the financial network helps us to better understand the recent financial crisis. Systemic risk started to increase in August 2007 but remained relatively low until after Bear Stearns' collapse. After that, the bounds widen considerably, which would be consistent both with an increase in systemic risk and with an increase in uncertainty about the severity of such risks. Finally, the bounds on systemic risk converge to lower values starting in the second half of 2009.

This paper shows how it is possible to decompose the changes in bond and CDS spreads across the network into an idiosyncratic and a systemic component of default risk. By using the full information set available from observed prices, the bounds show that some of the spikes in bonds yields and CDS spreads of financial institutions during the crisis correspond to increases in *idiosyncratic* default risk rather than systemic risk. This is the case for the months preceding March 15th 2008 (Bear Stearns episode) and following Lehman's default. This decomposition is important to understand the market's perceptions of risk during the financial crisis.

The same methodology can also be used to track the contribution to systemic risk of the institutions comprising the financial network, as well as to obtain a full representation of the default probabilities in the financial network in the scenario of highest systemic risk. An interesting result that emerges from this analysis is that markets seem to have anticipated the possible systemic nature of the risk exposure of some banks (Lehman and Merrill Lynch), several weeks before these institutions faced a severe crisis in September 2008: the upper bound on the joint default probability and the contributions to systemic risk of these two banks increased dramatically already by July 2008.

It is worth noting some limitations of the approach. The bounds capture market perceptions about probabilities rather than the true probabilities. This means that just as securities can be mispriced, the bounds can reflect various imperfections and mispricing that happen in financial markets, including slow incorporation of information and underestimation of risks. In addition, because they are reduced-form, they do not allow us to distinguish whether the low systemic risk estimated, for example, around March 2008 was due to the configuration of structural links between the banks or because of the expectation of government intervention in case multiple banks defaulted (which is what happened in September 2008). This also limits their use for policy action. Bond, Goldstein and Prescott (2010) warn us of the danger of taking actions based on market prices when these prices incorporate anticipation of the policy response. On the other hand, the bounds rule out some of the accounts about perceived risks, such as the fear of the government leaving several banks to default in the case that shocks propagated throughout the system. Also, while the approach is built to minimize the number of assumptions about the correlation structure of the network, especially regarding high-order joint risks, some assumptions are still necessary to obtain the bounds. For example, estimating risk-neutral marginal and pairwise default probabilities from prices requires imposing a pricing model and taking a stand on the effects of liquidity in bond markets.

The method developed in this paper can easily be extended to incorporate information from additional securities. The price of any security whose payoff depends on individual or joint defaults can be added as a constraint in the maximization and minimization problems, thus tightening the bounds. For example, the prices of CDS baskets and equity index options may enrich the information set and improve the measurement of systemic risk. In the future, as better data on these and other similar securities become more easily available, the measurement of systemic risk will become even more precise.

Because this method allows us to measure not only the aggregate level of systemic risk, but also the linkages between institutions, it opens the way to new studies of the structure of the financial system. For example, one can explore which network structures are more robust to financial distress, or the effect of individual banks' decisions on the aggregate level of systemic risk.

Since the bounds presented in this paper can be constructed in real time, they can be used to complement other measures in monitoring the market's perceptions of systemic risk. Not only do they track the effect of macroeconomic and financial shocks on perceived systemic risk, but they also reflect the market's expectations regarding government intervention in financial markets. Finally, the possibility of measuring in real time the structure of the financial network can be a valuable tool for identifying the sources of distress among banks at the core of the financial system.

References

Acharya, Viral V., Lasse H. Pedersen, Thomas Philippon and Matthew Richardson, 2010, "Measuring systemic risk", NYU Working Paper.

Adrian, Tobias and Markus K. Brunnermeier, 2009, "CoVaR," Staff Reports 348, Federal Reserve Bank of New York.

Anderson, Ron, 2009, "What Accounts for Time Variation in the Price of Default Risk?", LSE Working Paper.

Appa, Gautam, 2002, "On the uniqueness of solutions to linear programs", Journal of the Operational Research Society, 53, no. 10, 1127-1132.

Arora, Navneet, Priyank Gandhi and Francis Longstaff, 2009, "Counterparty Credit Risk and the Credit Default Swap Market", UCLA Working Paper.

Avesani, Renzo, Antonio Pascual and Jing Li, 2006, "A New Risk Indicator and Stress Testing Tool: A Multifactor Nth-to-Default CDS Basket," IMF Working Papers 06/105.

Bao, Jack, Jun Pan and Jiang Wang, 2008, "Liquidity of Corporate Bonds", MIT Sloan Working Paper.

Barro, Robert, 2006, "Credit-Default Swaps and Default Probabilities", unpublished, Harvard University.

Berndt, Antje, Robert Jarrow, and ChoongOh Kang, 2007, "Restructuring Risk in credit default swaps: An empirical analysis", Stochastic Processes and their applications, Vol 117 n. 11.

Blanco, Roberto, Simon Brennan, and Ian W. Marsh, 2003, An empirical analysis of the dynamic relationship between investment grade bonds and credit default swaps, The Journal of Finance, Vol. 60, No. 5 (Oct., 2005), pp. 2255-2281.

Bond, Philip, Itay Goldstein and Ed Prescott, 2010, "Market-Based Corrective Actions", Review of Financial Studies 23, 781-820.

Boros, E. and A. Prékopa, 1989, "Closed Form Two-Sided Bounds for Probabilities that at Least r and Exactly r out of n Events Occur", Mathematics of Operations Research 14, 317-342.

Chen, Long, David Lesmond and Jason Wei, 2007, "Corporate Yield Spreads and Bond Liquidity", Journal of Finance 42, 119-149.

Cihak, Martin and Klaus Shaeck, 2007, "How Well Do Aggregate Bank Ratios Identify Banking Problems?", IMF Working Paper 07/275.

Collin-Dufresne, Pierre, Robert S. Goldstein, and J. Spencer Martin, 2001, "The determinants of credit spread changes", Journal of Finance 56, 2177–2207.

Das, Sanjiv R., and Rangarajan K. Sundaram, 2000, "A discrete-time approach to arbitrage-free pricing of credit derivatives", Management Science 46, 46–62.

Demirguc, Asli and Enrica Detragiache, 1998, “The Determinants of Banking Crises in Developing and Developed Countries”, IMF Staff Paper 45/1.

Demirguc, Asli and Enrica Detragiache, 1999, “Monitoring Banking Sector Fragility: A Multivariate Logit Approach”, IMF Working Paper 99/147.

De Wit, Jan, 2006, “Exploring the CDS-bond basis”, National Bank of Belgium Working Paper.

Dick-Nielsen, Jens, Peter Feldhuetter, and David Lando, 2010, “Corporate bond liquidity before and after the onset of the subprime crisis”, Working Paper.

Duffie, Darrell, 1999, “Credit swap valuation”, Financial Analysts Journal 55, 73–87.

Duffie, Darrell, and Ming Huang, 1996, “Swap rates and credit quality”, Journal of Finance 51, 921-949.

Duffie, Darrell, and Kenneth J. Singleton, 1999, “Modeling term structures of defaultable bonds”, Review of Financial Studies 12, 687–720.

Duffie, Darrell, and Haoxiang Zhu, “Does a Central Clearing Counterparty Reduce Counterparty Risk?”, Stanford GSB Research Paper.

Garleanu, Nicolae G. and Lasse H. Pedersen, 2010, “Margin-Based Asset Pricing and Deviations from the Law of One Price”, NYU Working Paper.

Gonzales-Hermosillo, Brenda, 1999, “Determinants of ex-ante Banking System Distress: A Macro-Micro Empirical Exploration of Some Recent Episodes”, IMF Working Paper 99/23.

Houweling, Patrick, and Ton Vorst, 2004, “Pricing default swaps: Empirical evidence”, Journal of International Money and Finance, 29 1331–1358.

Huang, Jing-zhi, and Ming Huang, 2003, “How much of the corporate-Treasury yield spread is due to credit risk?”, Working paper, Penn State University.

Huang, Xin, Hao Zhou, and Haibin Zhu, 2009, “A Framework for Assessing the Systemic Risk of Major Financial Institutions”, Journal of Banking and Finance, vol 33 no. 11.

Hull, John C., and Alan D. White, 2000, “Valuing credit default swaps I: No counterparty default risk”, The Journal of Derivatives 8, 29–40.

Hull, John C., and Alan D. White, 2001, “Valuing credit default swaps II: Modeling default correlations”, The Journal of Derivatives 8, 12–22.

Jarrow, Robert A., and Stuart M. Turnbull, 1995, “Pricing derivatives on financial securities subject to credit risk”, Journal of Finance 50, 53–85.

Jarrow, Robert A., and Fan Yu, 2001, “Counterparty risk and the pricing of defaultable securities”, Journal of Finance 56, 1765-1799.

Kritzman, Mark, Yuanzhen Li, Sebastien Page and Roberto Rigobon, 2010, “Principal Components as a Measure of Systemic Risk”, Revere Street Working Paper 272-28.

Kwerel, S.M., 1975, "Most stringent bounds on aggregated probabilities of partially specified dependent probability systems," Journal of the American Statistics Association 70, 472-479.

Lando, David, 1997, "Modelling bonds and derivatives with credit risk", in M. Dempster and S. Pliska (eds.), Mathematics of Financial Derivatives, pp. 369–393, Cambridge University Press.

Lehar, Alfred, 2005, "Measuring systemic risk: A risk management approach", Journal of Banking and Finance, vol. 29 no. 10.

Longstaff, Francis, Sanjay Mithal and Eric Neis, 2005, "Corporate Yield Spreads: Default Risk or Liquidity? New Evidence from the Credit-Default Swap Market", Journal of Finance, Vol. 60, No. 5.

Merton, Robert C., 1974, "On the Pricing of Corporate Debt: The Risk Structure of Interest Rates", Journal of Finance, Vol. 29, No. 2.

Poghoshan, Tigran and Martin Cihak, 2009, "Distress in European Banks: An Analysis Based on a New Data Set", IMF Working Paper 09/09.

Segoviano, Miguel, 2008, "The CIMDO Copula. Modeling of a Non-Parametric Copula", IMF Working Paper

Segoviano, Miguel and Charles Goodhart, 2009. "Banking Stability Measures," IMF Working Papers 09/4, International Monetary Fund.

Singh, Manmohan and James Aitken, 2009, "Counterparty Risk, Impact on Collateral Flows, and Role for Central Counterparties", IMF working paper 09/173.

Singh, Manmohan, 2010, "Collateral, Netting and Systemic Risk in the OTC Derivatives Market", IMF working paper 10/99.

Sundaresan, S. ,1991, "Valuation of Swaps.", in Recent Developments in International Banking and Finance, S. Khoury, ed. Amsterdam: North Holland (Vols. IV and V, 1991).

Appendix

Appendix A - Collateral agreements and the pricing of counterparty risk

In the text, I argue that the collateral agreements used for CDS contracts during the financial crisis were unlikely to eliminate counterparty risk. Buyers of CDS protection were aware of this and possibly priced it into the spreads. Here I report some evidence for the main points of the argument.

An initial question is whether counterparty risk was perceived at all by market participants. The growth of the percentage of OTC derivative contracts covered by some form of collateral confirms this indirectly: for credit derivatives, the volume-weighted percentage of collateralized contracts went from 39% in 2004 to 58% in 2005, to 66% in 2007 and 2008 (ISDA Margin Survey 2006, 2008). Besides, documents and interviews from practitioners directly confirm that the issue was taken into account by financial participants throughout the crisis. Robert McWilliam, head of Counterparty Risk management at ABN Amro, reports in January 2008[40]*: "The golden rule is to start early. If you start worrying about the counterparty when they are under duress your options are fairly limited"*. A document from Barclays dated February 2008[41] states:*"While the maximum potential loss to the seller of protection is the contract spread for the rest of the contract duration, the buyer of protection could arguably lose the full notional of the contract (in case of simultaneous defaults by counterparty and the reference credit and zero recovery). Thus, counterparty risk is evidently more of a concern for buyers of protection."*

Even if agents were aware of counterparty risk, it was standard practice to ask for relatively little collateral, especially from the largest counterparties. ISDA reports that only about 2/3 of the contracts were covered by a collateral agreement, up to 2009. Besides, calculations by Singh and Aitken (2009) and Singh (2010) show that, even at the end of 2009, large financial institutions still carried large under-collateralized derivative liabilities. In particular, they compute the total value of "residual derivative payables" - liabilities from derivative positions after netting under master netting agreements and in excess of the collateral posted. For the 5 largest US dealers this amount was more than $250bn. Even though these numbers include all derivative contracts, and not only CDSs, they suggest a general under-collateralization of derivative positions from these counterparties. As an example of this, in 2008 Goldman Sachs had received collateral for 45% of the value of its receivable OTC derivatives, but posted only 18% of the value of payables. Similarly, JP Morgan in the same year had received collateral for 47% of receivables, but posted only 37% of the payables[42]. Finally, as reported in the main text, even the most active dealer in counterparty risk management, Goldman Sachs, failed to cover the full value of exposure on its CDS position with AIG.

Even when a collateral agreement is in place and actively managed, residual counterparty risk cannot be eliminated when the value of the derivative is subject to jumps. While during the crisis we did see gradual increases in CDS spreads of banks, a crucial episode - the Lehman bankruptcy - shows that correlated jumps in credit risk (and defaults) are indeed possible. Just before the weekend of the 13th and 14th of September 2008, many institutions were considered at risk, but neither the credit ratings nor the CDS spreads indicated an extremely high likelihood of immediate default. For example, the Lehman 5 year CDS was trading at around 700bp per year, Merrill's at 400bp, and the credit ratings of their debt were still as high as 4 months before, with an implied default probability of less than 0.25% per annum. A buyer who bought a Lehman or a Merrill CDS at 350bp per year a month before the default would have seen the value of the contract (the present discounted value of the difference in spreads) grow to 15 cents and 5 cents on the dollar respectively on Friday September 12th. Therefore, even if the buyers had called for enough collateral to cover the current

[40]Reuters,"Banks move to guard against counterparty failures", Jan 24, 2008.
[41]Barclays Capital, 2008, "Counterparty Risk in Credit Markets".
[42]Bloomberg, "Goldman Sachs Demands Collateral It Won't Dish Out", March 15, 2010.

value of such contracts, they would have improved their recovery rate by only 5% to 15%.

For the reasons explained above, buyers were generally aware that the collateral agreements in place (if any) would have left them exposed to the risk of double default. In fact, several sources document that in early 2008 buyers of CDS contracts were buying additional CDS contracts *against their counterparties* to hedge the residual counterparty risk. For example, from the documents on the AIG bailout (Maiden Lane III) from the Financial Crisis Inquiry Commission, we see that starting November 2007, Goldman Sachs - which had bought $22bn of CDS on a super-senior tranche of a CDO from AIG - was adjusting the amount of CDS protection *against* AIG together with their margin calls to AIG (which were caused by increases in the default probability of the underlying asset). Up to June 2008, the nominal amount of protection bought against AIG was of the same order of magnitude as the total amount of collateral called by Goldman.

In a document issued by Goldman Sachs in 2009 regarding the AIG bailout[43], the firm declares: *"In mid-September 2008, prior to the government's action to save AIG, a majority of Goldman Sachs' exposure* [current market value] *to AIG was collateralized and the rest was covered through various risk mitigants. Our total exposure on the securities on which we bought protection was roughly $10 billion. Against this, we held roughly $7.5 billion in collateral. The remainder was fully covered through hedges we purchased, primarily through CDS for which we received collateral from our market counterparties. Thus, if AIG had failed, we would have had the collateral from AIG and the proceeds from the CDS protection we purchased."*. Similarly, in an interview with ABN Amro, Reuter reports[44]: *When counterparties* [to OTC derivatives] *are large corporations, which do not usually put up collateral, ABN buys protection in the CDS market against the default of the counterparty itself. ABN's trading desk must go into the market constantly to rebalance those CDS holdings so that its protection equals its counterparty risk profile."*.

This evidence indicates that buyers were understanding the direct and indirect costs of the residual counterparty risk. Note that the fact that collateral was not enough to eliminate counterparty risk does not mean that buyers were making a bad deal on their contracts. Simply, they would have been compensated by paying a lower spread for the contracts when the counterparty was at higher risk of double default. In fact, the 2008 Barclays report titles a section: *How much should I pay for a higher-rated counterparty?* (The analysis then quantifies this number for generic corporate reference entities of different credit rating).

Appendix B - Implementation of the Linear Programming Problem

This appendix describes in detail the algorithm employed to transform the probability bounds problem into a linear programming problem. It also describes the bond pricing formula and the linear approximation to the CDS pricing formula that allows to write the CDS constraints as linear constraints.

B.1 - Linear programming representation in the general case

This section describes the algorithm used to transform the probability problem

$$max\, P_r$$

s.t.

$$P(A_i) = a_i$$

$$...$$

[43]Goldman Sachs, "Overview of Goldman Sachs' Interaction with AIG and Goldman Sachs' Approach to Risk Management", March 2009.

[44]Reuters, "Banks move to guard against counterparty failures", Jan 24, 2008.

$$P(A_i \cap A_j) = a_{ij}$$

into the LP representation

$$max_p \, c_r' p$$

s.t.

$$p \geq 0$$

$$i'p = 1$$

$$Ap = b$$

for the general case of N banks.

Start with a matrix B of size $(2^N, N)$ whose rows contain the binary representation of all numbers between 0 and $2^N - 1$. For example, with N=4:

$$B = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 1 \\ & ... & & \\ 1 & 1 & 1 & 1 \end{bmatrix}$$

Each row of this matrix corresponds to a particular element of the partition of the sample space described in Proposition 2: the event

$$A_1^* \cap A_2^* \cap ... \cap A_N^*$$

where $A_j^* = A_j$ if element j of the row is 1, and $A_j^* = \overline{A}_j$ if element j of the row is 0. The probability system p will then be determined as a vector of 2^N elements containing the probability of each of the elements of the partition represented by the 2^N rows of the matrix B. For example p_1 will be the probability that none of the A_i events occur, p_2 will represent the probability that event A_N occurs but none of the other events does, and so on. Finally, the element p_{2^N} will represent the probability that all events occur.

The maximization problem presented above tries to find the vector p that maximizes the probability of systemic event of degree r (P_r) while satisfying constraints on marginal and pairwise default probabilities, as well as the constraints implied by the consistency of the probability measure. The latter are immediate: because the events represented by the rows of B are a partition of the sample space, and p is a probability measure on these events, all elements of p need to be nonnegative and sum to one:

$$p \geq 0$$

$$p'i = 1$$

To obtain in LP form the inequalities and equalities that involve marginal and pairwise default probabilities, note first that because the elements of the partition are disjoint events, the probability of any union of them is equal to the sum of their probabilities. Therefore, to find the probability of an event A_i, $P(A_i)$, in terms of p, one needs to sum the probabilities of all the elements of the partition in which event A_i occurs. But this is immediate given the representation in B:

$$P(A_i) = \sum_{j:B(j,i)=1} p_j$$

or:

$$P(A_i) = a^{i'} p$$

for a vector a_i of size $(2^N, 1)$ s.t.:

$$a^i_j = B(j, i)$$

In other words, to find which elementary events form event A_i one needs to find all the rows of B in which element i is equal to 1. The union of these events will coincide with A_i, and therefore the sum of their probabilities will be $P(A_i)$. Given the linearity, this sum is equivalent to the product of the vector p with a vector a_i, whose elements are ones whenever the corresponding elementary event is a subset of A_i.

Similarly, the probability of a joint event:

$$P(A_i \cap A_k) = \sum_{j:B(j,i)=1 \, and \, B(j,k)=1} p_j$$

or:

$$P(A_i \cap A_k) = b^{ik'} p$$

for a vector b_{ik} of size $(2^N, 1)$ s.t.:

$$b^{ik}_j = B(j, i)B(j, k)$$

i.e., the probability of the joint default is obtained summing the elements of p s.t. the corresponding element of the partition involves *both* the occurrence of A_j and of A_k. All these constraint can then be collected in the matrix form $Ap = b$.

Finally, the probability that at least r events occur can be found as follows:

$$P_r = \sum_{j:\left(\sum_{h=1}^{N} B(j,h)\right) \geq r} p_j$$

or:

$$P_r = c^{r'} p$$

for a vector c^r of size $(2^N, 1)$ s.t.:

$$c^r_j = I \left[\sum_{h=1}^{N} B(j, h) \geq r \right]$$

where $I[]$ is the indicator function.

Given this decomposition, the LP representation follows immediately.

B.2 - Proof of Proposition 3

Start from a symmetric LP problem

$$max \, c'p$$

$$s.t. \, Ap \leq b$$

Suppose that p^* is a solution to the problem. Given the definition of symmetry presented in the text, it is clear that p^*_J is also a solution to the problem: $c'p^* = c'p^*_J$ and similarly hold for every row of the constraints, for every J.

Now, construct p^{**} as follows:

$$p^{**} = \frac{1}{\#J} \sum_J p^*_J$$

where the first J correspond to no permutation, and J cycles across all permutations of indices $A_1, ..., A_N$.

Note that it is also possible to construct p^{**} in the following way, considering the binary representation introduced in Proposition 2. Every b_i vector has O_i ones and $N - O_i$ zeros. Call H_i the set of all vectors of size N that have O_i ones and $N - O_i$ zeroes in different positions. Call b_{ih} the vector corresponding to element h from H_i. Then, for every i, construct p^{**} as:

$$p_i^{**} = \left(\begin{array}{c} O_i \\ N \end{array} \right)^{-1} \sum_{h \in H_i} b_{ih}$$

From the first construction, it is clear why p^{**} is a solution to the maximization problem, being just an average of solutions. Additionally, p^{**} is symmetric, which proves the statement of the Proposition.

An example with $N = 3$. As explained in Proposition 2, we can construct the probability system p^* as follows:

$$p_1^* = Pr\{\overline{A}_1 \cap \overline{A}_2 \cap \overline{A}_3\}$$

$$p_2^* = Pr\{\overline{A}_1 \cap \overline{A}_2 \cap A_3\}$$

$$p_3^* = Pr\{\overline{A}_1 \cap A_2 \cap \overline{A}_3\}$$

$$p_4^* = Pr\{\overline{A}_1 \cap A_2 \cap A_3\}$$

$$p_5^* = Pr\{A_1 \cap \overline{A}_2 \cap \overline{A}_3\}$$

$$p_6^* = Pr\{A_1 \cap \overline{A}_2 \cap A_3\}$$

$$p_7^* = Pr\{A_1 \cap A_2 \cap \overline{A}_3\}$$

$$p_8^* = Pr\{A_1 \cap A_2 \cap A_3\}$$

Suppose p^* solves the maximization problem, and construct p^{**} as:

$$p_1^{**} = p_1^*$$

$$p_2^{**} = p_3^{**} = p_5^{**} = \frac{p_2^* + p_3^* + p_5^*}{3}$$

$$p_4^{**} = p_6^{**} = p_7^{**} = \frac{p_4^* + p_6^* + p_7^*}{3}$$

$$p_8^{**} = p_8^*$$

p^{**} solves the maximization problem and is symmetric.

B.3 - Uniqueness of the solution and representative solution

As described in section 3, it is possible to describe the range of possible values for a particular probability $P(A') = v'p$ among the probability systems p that attain the solution to the probability bounds problem:

$$c_{max} = max_p \, c_r' p$$

s.t.

$$p \geq 0$$

$$i'p = 1$$

$$Ap = b$$

(and similarly for the lower bound). To find the maximum (minimum) value of $v'p$ at the upper bound on systemic risk, simply solve the max (or min) problem

$$max_p(min_p)\, v'p$$

s.t.

$$p \geq 0$$

$$i'p = 1$$

$$Ap = b$$

$$c_r'p = c_{max}$$

To find the maximum and minimum value of $v'p$ at the lower bound, replace the last constraint with $c_r'p = c_{min}$.

Appa (2002) describes an algorithm (called AFROS) to find a "representative" solution to the bounds. Call $Ep = d$ the set of all equality constraints in the problem, and note that the original maximization problem can always be rewritten with a series of equality constraints and a nonnegativity constraint for an expanded vector p. The algorithm proceeds as follows.

1. Solve P^0 : max $c_r'p$ s.t. $Ep = d, p \geq 0$. Let $p = p^0$ be the solution to P^0.

2. Solve P^1: max d^1p s.t $Ep = d, p \geq 0, c_r'p = c_r'p^0$, where d^1 is a row vector with $d_j^1 = 1$ if $p_j^0 = 0$, otherwise $d_j^1 = 0$. If maximum value of $d^1p = 0$, stop. There is no alternative solution. Otherwise, let $p = p^1$ be the optimal solution to P^1. Set counter $q = 1$

3. $q = q + 1$. While $q < s$, solve P^q: max d^qp s.t. $Ep = d, p \geq 0, c_r'p = c_r'p^0, d^rp \geq \alpha^r$ for $r = 1, ..., q-1$. $d_j^q = 1$ if $p_j^{q-1} = 0$, 0 otherwise, and α^r are positive numbers.

4. After S steps, the "representative solution" is the average of the p from p^0 to p^S.

The algorithm starts from a solution p^0 and looks for another solution that maximizes the sum of all and only those elements of p that are hitting the constraint $p \geq 0$ at p^0. This is the sense in which p^1 will be as dissimilar as possible from p^0. From p^1, it will repeat the same procedure to find another solution p^2, and so on, with additional constraints that ensure that the algorithm never converges back to a previous solution. The average of these solutions, itself a probability system, is the "representative" solution.

B.4 Width of the bounds

Given Proposition 3, we can derive the analytical formulation of the bounds derived as a special case of the bounds presented in Boros and Prekopa (1989), section 8. Here, $S_1 = Nq_1$ and $S_2 = \frac{N(N-1)}{2}q_2$. The bounds are nonlinear functions of S_1 and S_2 and take the following form, for $3 \leq r \leq N-1$. Square brackets indicate integer part.

Lower bound in the symmetric case

Call

$$A_1 = (r-1)N - (r+N-2)Nq_1 + N(N-1)q_2$$

$$A_2 = -(r-2)Nq_1 + N(N-1)q_2$$

and remember that $q_2 \leq q_1$. Then,

$$A_2 \leq 0 \Longleftrightarrow q_2 \leq \frac{(r-2)}{(N-1)}q_1$$

$$A_1 \leq 0 \Longleftrightarrow q_2 \leq \frac{(r+N-2)q_1 - r + 1}{N-1} = \frac{(r-2)q_1}{N-1} + \frac{Nq_1 - r + 1}{N-1}$$

The bounds then can be written as follows.

CASE 1: If $q_1 < \frac{r-1}{N}$

- If $q_2 \leq \frac{(r-2)}{(N-1)}q_1$

$$p_r \geq 0$$

- If $\frac{(r-2)}{(N-1)}q_1 \leq q_2 \leq q_1$

$$p_r \geq \frac{-(r-2)Nq_1 + N(N-1)q_2}{(N+1-r)N}$$

CASE 2: If $q_1 > \frac{r-1}{N}$

- If $q_2 \leq \frac{(r-2)q_1}{N-1} + \frac{Nq_1-r+1}{N-1}$

$$i = \left\lceil \frac{N(N-1)q_2 - (r-2)Nq_1}{Nq_1 - (r-1)} \right\rceil$$

$$p_r \geq \frac{(r-1)(r-2i-2) + 2iNq_1 - N(N-1)q_2}{(i-r+2)(i-r+1)} \tag{24}$$

- If $q_2 \geq \frac{(r-2)}{(N-1)}q_1 + \frac{Nq_1-r+1}{N-1}$

$$p_r \geq \frac{-(r-2)Nq_1 + N(N-1)q_2}{(N+1-r)N}$$

Upper bound in the symmetric case

Call

$$B_1 = rN - (r+N-1)Nq_1 + N(N-1)q_2$$

$$B_2 = -(r-1)Nq_1 + N(N-2)q_2$$

The key cases are distinguished by whether

$$\frac{(r+N-1)}{N-1}q_1 - \frac{r}{N-1} > \frac{(r-1)}{(N-2)}q_1$$

or

$$\frac{(N-2)(r+N-1) - (r-1)(N-1)}{(N-2)}q_1 > r$$

$$q_1 > \frac{N-2}{N(N-2) - r + 1}r$$

Then we have:
CASE 1: If $q_1 > \frac{N-2}{N(N-2)-r+1}r$

- If $q_2 < \frac{(r+N-1)}{N-1}q_1 - \frac{r}{N-1}$
$$p_r \leq 1$$

- If $q_2 \geq \frac{(r+N-1)}{N-1}q_1 - \frac{r}{N-1}$
$$p_r \leq \frac{(r+N-1)Nq_1 - N(N-1)q_2}{rN}$$

CASE 2: If $q_1 < \frac{N-2}{N(N-2)-r+1}r$

- If $q_2 < \frac{(r-1)}{(N-2)}q_1$
$$p_r \leq \frac{i(i+1) - 2iNq_1 + N(N-1)q_2}{(r-i-1)(r-i)} \tag{25}$$
where
$$i = \left\lceil \frac{(r-1)Nq_1 - N(N-1)q_2}{r - Nq_1} \right\rceil$$

- If $q_2 \geq \frac{(r-1)}{(N-2)}q_1$
$$p_r \leq \frac{(r+N-1)Nq_1 - N(N-1)q_2}{rN}$$

Width of the bounds

The formulation presented above allows us to state the following result. Suppose that the probability system is symmetric with respect to marginal and pairwise default probabilities, i.e. $\forall i\ P(A_i) = q_1$ and $\forall i, j$ $P(A_i \cap A_j) = q_2$. Then, for given q_1, the upper and lower bounds for P_r, $r \geq 3$, have the following properties:

- If $q_1 < \frac{r-1}{N}$, the *lower bound* is 0 for low q_2 and is increasing in q_2 for higher q_2. If $q_1 > \frac{r-1}{N}$, the *lower bound* is first decreasing and then increasing in q_2.

- If $q_1 < \frac{N-2}{N(N-2)-r+1}r$, the *upper bound* is first increasing and then decreasing in q_2. If $q_1 > \frac{N-2}{N(N-2)-r+1}r$, the *upper bound* is first 1 and then decreasing in q_2.

Similar results hold for q_2 fixed, when q_1 varies. The upper bound is concave and the lower bound is convex. Both bounds are continuous. The width of the bounds is continuous and concave, and has a maximum in the interior of the parameter space.

Proof. Because we are maximizing or minimizing a continuous linear function over a convex set which is a continuous function of q_1, and q_2, the upper bound will be concave in q_1 and q_2 and the lower bound will be convex, and the bounds will be continuous functions of q_1 and q_2. As shown above, they are piecewise linear. Note that we can check that equation (24) is decreasing in q_2 because each of the linear parts are decreasing and the function is continuous. Similarly we can show that equation (25) is increasing in q_2. Therefore, by looking at the different combinations of values for r and N, we can check that for given q_1, the width of the bounds is first increasing and then decreasing in q_2 The same is true for fixed q_2. This also implies that the maximum width, a positive concave function, is attained in the interior of the parameter space (q_1, q_2).

This result implies that the tightness of the bounds on systemic risk varies in a very precise way with changes in the low-order risks. When pairwise default probabilities q_2 are either very high or very low compared to the marginal default probabilities q_1, the structure of the network is pinned down very precisely and there is little uncertainty about systemic default risk, based on bond and CDS prices alone; different parametric models that aggregate in different ways this low-order information will agree on an estimate of systemic risk. For intermediate values of q_1 and q_2, however, low-order probabilities are less informative about systemic events. The scope for modeling assumptions in aggregating low-order probabilities is greater.

Appendix C: additional pricing details

C.1 Bond pricing

In this section I show how to obtain the bond pricing formula in the text starting from the general discrete-time formulation of bond pricing in the reduced form model. Call h_t^i the hazard rate process for i: the probability of default in month t conditional on survival until then. Call r_t^F the monthly risk-free process. Call the $r_{t,T}^F$ the realized return of the short-term risk-free security between t and T, s.t.

$$(1 + r_{t,T}^F) = \prod_{s=t+1}^{T} (1 + r_s^F)$$

The time t price of a risk-free zero-coupon bond of face value \$1 at time T is:

$$\delta(t, T) = E_t^Q \left[\frac{1}{1 + r_{t,T}^F} \right]$$

Call $G^i(t, s)$ the probability of survival of firm i up to s under a certain realization of hazard rates of default, i.e.:

$$G^i(t, s) = \prod_{r=t}^{s-1} (1 - h_r^i)$$

where E^Q indicates the expectation taken under the risk neutral probability measure.

The price of a *liquid* bond j issued by firm i of face value \$1, maturity T^{ij}, coupon rate c^{ij} and recovery equal to a fraction R^{ij} of the value of a Treasury zero-coupon bond of comparable maturity is:

$$B^{ij}(t, T^{ij}) = E_t^Q \left[\sum_{s=t+1}^{T^{ij}} \frac{G^i(t, s)}{1 + r_{t,s}^F} c^{ij} + \frac{G^i(t, T^{ij})}{1 + r_{t,T}^F} + \sum_{s=t+1}^{T^{ij}} \frac{G^i(t, s-1) h_{s-1}^i}{1 + r_{t,s}^F} R^{ij} \right]$$

In this paper I consider only senior unsecured bonds of different coupons and maturities, so I assume that the recovery rate is the same for all bonds and that it is also the same for similar bonds of other firms in the financial industry, i.e. $R^{ij} = R$.

We can then add a liquidity process γ_t^i , assumed to be the same for all bonds of equal seniority of firm i. Following Duffie (1999), this liquidity cost will appear in the bond pricing equation as a per-period proportional cost incurred while holding the bond.

In theory, it is possible to write down a parametric version of the (generally not independent) processes that govern the evolution of r_t^F,h_t, γ_t, and additionally a time-varying recovery rate. Examples of this can be found in Duffie and Singleton (1997) and Longstaff, Mithal and Neis (2005). In this paper, I use a simplified pricing model that assumes that, for any given firm, the prices of all of its bonds are determined independently at each time t, under the assumption that from time t onwards h_{t+s} and γ_{t+s} will be constant and equal to h_t and γ_t, respectively. This is just an approximation, because prices do not take into account that at each future date these parameters are going to be revised, since at every future date $t + r$ prices will be recomputed assuming a constant hazard rate and liquidity process from $t+r$ on at new levels, h_{t+r} and γ_{t+r}. I discretize the model to a monthly horizon, and I assume that coupons are paid monthly. Further assuming independence of the risk-free rate process from all other processes (under Q) we obtain:

$$B^{ij}(t,T^{ij}) = c^{ij}\left(\sum_{s=t+1}^{T^{ij}} \delta(t,s)(1-h_t^i)^{s-t}(1-\gamma_t^i)^{s-t}\right)+$$

$$+\delta(t,T^{ij})(1-h_t^i)^{T^{ij}-t}(1-\gamma_t^i)^{T^{ij}-t} + R\left(\sum_{s=t+1}^{T^{ij}} \delta(t,s)(1-h_t^i)^{s-t-1}(1-\gamma_t^i)^{s-t-1}h_t^i\right)$$

C.2 - Additional details of CDS contracts

Besides those considered explicitly in this paper, there are other elements of CDS contracts that potentially affect their spreads.

First, liquidity of the CDS market could influence the CDS spreads, just as bond liqudiity is known to affect bond prices. In this paper, I explicitly take into account liquidity premia in bond prices, but not in CDS spreads. For the case of CDSs, liquidity is much less likely to be an issue, especially because they require much less capital at origination and they are not in fixed supply.[45]

Also, I abstract from restructuring clauses and the cheapest-to-deliver option sometimes present in CDS contracts. A restructuring clause (under which payment is triggered for simple debt restructuring, in addition to bankruptcy) is more frequent for European bonds, and this results in the contract being triggered in cases close to the Chapter 11 for the US. Berndt, Jarrow and Kang (2007) estimate that the presence of such clause increases the value of the CDSs by 6-8%, and all the results in this paper are robust to an adjustment of CDS spreads of that magnitude. The value of the cheapest-to-deliver option (which allows the buyer to deliver to the seller the cheapest of the defaulted bonds of the same seniority as the reference bond) will be small relative to the CDS spread as long as in default all senior unsecured bonds have similar recovery rates. Additionally, as observed in the Delphi and Calpine defaults in 2005, the high demand for the cheapest bonds might determine shortages of such securities and therefore, anticipating this, a reduction in the ex-ante value of the option.[46]

C.3 - CDS pricing approximation

Start from the discretized pricing equation with constant hazard and risk-free rates, starting from period 0 for notational simplicity.

$$\sum_{s=1}^{T} \delta(0,s-1)(1-P(A_i \cup A_j))^{s-1} z_{ji} =$$

$$= \left[\sum_{s=1}^{T}(1-P(A_i \cup A_j))^{s-1}(P(A_i)-(1-S)P(A_i \cap A_j))\delta(0,s)(1-R)\right]$$

We can rewrite the equation as:

$$\frac{z_{ij}}{(1-R)} = \frac{\left[\sum_{s=1}^{T} \delta(0,s)(1-P(A_i \cup A_j))^{s-1}(P(A_i)-(1-S)P(A_i \cap A_j))\right]}{\sum_{s=1}^{T} \delta(0,s-1)(1-P(A_i \cup A_j))^{s-1}}$$

[45] For an additional discussion of this and on the supporting evidence, see Blanco, Brennan and Marsh (2003,2005).

[46] De Wit (2006).

and then approximate the right hand side around $P(A_i)=0, P(A_j)=0, P(A_i\cap A_j)=0$, remembering that

$$P(A_i\cup A_j)=P(A_i)+P(A_j)-P(A_i\cap A_j)$$

To obtain the approximation, note that while we cannot vary $P(A_i)$ and $P(A_i\cap A_j)$ independently around 0, we can rewrite the expression in terms of:

$$\pi_i = P(A_i\cap \overline{A}_j)$$

$$\pi_j = P(\overline{A}_i\cap A_j)$$

$$\pi_{ij} = P(A_i\cap A_j)$$

which can be varied independently of each other. The right-hand side (call it G, and call $G(0)$ the function at the approximation point) can then be written as:

$$G=\frac{\left[\sum_{s=1}^{T}\delta(0,s)(1-\pi_i-\pi_j-\pi_{ij})^{s-1}(\pi_i+\pi_{ij}-(1-S)\pi_{ij})\right]}{\sum_{s=1}^{T}\delta(0,s-1)(1-\pi_i-\pi_j-\pi_{ij})^{s-1}}$$

First, note that $G(0)=0$. Second, take $G_{\pi_i}(0)$:

$$G_{\pi_i}(0)=\frac{\frac{d}{d\pi_i}\left[\sum_{s=1}^{T}\delta(0,s)(1-\pi_i-\pi_j-\pi_{ij})^{s-1}(\pi_i+\pi_{ij}-(1-S)\pi_{ij})\right]\left[\sum_{s=1}^{T}\delta(0,s-1)(1-\pi_i-\pi_j-\pi_{ij})^{s-1}\right]+}{\left[\sum_{s=1}^{T}\delta(0,s-1)(1-\pi_i-\pi_j-\pi_{ij})^{s-1}\right]^2}$$

$$\frac{\left[\sum_{s=1}^{T}\delta(0,s)(1-\pi_i-\pi_j-\pi_{ij})^{s-1}(\pi_i+\pi_{ij}-(1-S)\pi_{ij})\right]\frac{d}{d\pi_i}\left[\sum_{s=1}^{T}\delta(0,s-1)(1-\pi_i-\pi_j-\pi_{ij})^{s-1}\right]}{\left[\sum_{s=1}^{T}\delta(0,s-1)(1-\pi_i-\pi_j-\pi_{ij})^{s-1}\right]^2}$$

Note that the second part is 0 and at the approximation point, and at that point we have:

$$\frac{d}{d\pi_i}\left[(1-\pi_i-\pi_j-\pi_{ij})^{s-1}(\pi_i+\pi_{ij}-(1-S)\pi_{ij})\right]=$$

$$\frac{d}{d\pi_i}(1-\pi_i-\pi_j-\pi_{ij})^{s-1}\left[(\pi_i+\pi_{ij}-(1-S)\pi_{ij})\right]+(1-\pi_i-\pi_j-\pi_{ij})^{s-1}\frac{d}{d\pi_i}\left[(\pi_i+\pi_{ij}-(1-S)\pi_{ij})\right]=1$$

So we have:

$$G_{\pi_i}(0)=\frac{\sum_{s=1}^{T}\delta(0,s)}{\sum_{s=1}^{T}\delta(0,s-1)}$$

Similarly:

$$G_{\pi_j}(0)=\frac{\sum_{s=1}^{T}\delta(0,s)\frac{d}{d\pi_j}\left[(1-\pi_i-\pi_j-\pi_{ij})^{s-1}(\pi_i+\pi_{ij}-(1-S)\pi_{ij})\right]\left[\sum_{s=1}^{T}\delta(0,s-1)\right]}{\left[\sum_{s=1}^{T}\delta(0,s-1)\right]^2}$$

with

$$\frac{d}{d\pi_j}\left[(1-\pi_i-\pi_j-\pi_{ij})^{s-1}(\pi_i+\pi_{ij}-(1-S)\pi_{ij})\right]=$$

$$\left[\frac{d}{d\pi_j}(1-\pi_i-\pi_j-\pi_{ij})^{s-1}\right]\left[(\pi_i+\pi_{ij}-(1-S)\pi_{ij})\right]=0$$

so that

$$G_{\pi_j}(0) = 0$$

Finally we get:

$$G_{\pi_{ij}}(0) = \frac{\sum_{s=1}^{T} \delta(0,s)}{\sum_{s=1}^{T} \delta(0,s-1)} S$$

So that:

$$G \simeq \frac{\sum_{s=1}^{T} \delta(0,s)}{\sum_{s=1}^{T} \delta(0,s-1)} [\pi_i + S\pi_{ij}] = \frac{\sum_{s=1}^{T} \delta(0,s)}{\sum_{s=1}^{T} \delta(0,s-1)} [\pi_i + \pi_{ij} - (1-S)\pi_{ij}]$$

$$= \frac{\sum_{s=1}^{T} \delta(0,s)}{\sum_{s=1}^{T} \delta(0,s-1)} [P(A_i) - (1-S)P(A_i \cap A_j)]$$

The result is:

$$\frac{z_{ij}}{(1-R)} \simeq \frac{\left[\sum_{s=1}^{T} \delta(0,s)\right]}{\left[\sum_{s=1}^{T} \delta(0,s-1)\right]} (P(A_j) - (1-S)P(A_i \cap A_j))$$

It is important to check the accuracy of the approximation for a realistic range of parameters. For several different points in time (every 50 days) between 1/1/2007 and 3/31/2009, I compare the correct spread and the approximated spread, computed using the US yield curve at that time, considering:

- different values of $P(A_j)$: between 0 and the maximum probability implied by bond data under no liquidity assumptions ($max_j\{h_j(0)\}$).
- different values of $P(A_i \cap A_j)$: between 0 and $P(A_j)$
- different values of R and S: between 0.1 and 0.4.

In all these simulations, the approximation error is between 0.2% and 0.3% of the true value of the CDS spread.

Appendix D - Details on the Robustness Tests

D.1 - Assumptions on S

The recovery rate of a CDS in case of double default, S, affects positively the value of that security. In addition, a higher recovery rate S implies a higher sensitivity of the CDS spreads to changes in the joint default risk. A recovery rate of zero ($S = 0$) means that counterparty risk has the highest impact on CDS spreads, while a recovery rate of 1 means that joint default risk has no effect on the spread.

To understand the effect of assumptions about S on the measure of systemic risk presented in this paper, remember that the upper bound on systemic risk is attained by the most correlated probability system that satisfies the constraints:

$$P(A_i) \leq h_i(\underline{\gamma}_t^i)$$

$$P(A_i) - (1-S)\left(\frac{1}{N-1}\sum_{j \neq i} P(A_i \cap A_j)\right) = b_i$$

where $b_i = \overline{z}_i \frac{\left[\sum_{s=1}^{T} \delta(0,s-1)\right]}{\left[\sum_{s=1}^{T} \delta(0,s)\right](1-R)}$ [47].

Intuitively, for a given S, one can obtain the most correlated probability system by setting $P(A_i)$ as high as possible (up to the constraint $h_i(\underline{\gamma}_t^i)$) for all banks and then increasing the term $\frac{1}{N-1}\sum_{i\neq j} P(A_i \cap A_j)$ to match the CDS spreads (b_i). Counterparty risk would explain the whole bond/CDS basis, and a higher recovery rate S would imply that a higher joint default probability is needed to match it, increasing the upper bound on systemic risk. This intuitive reasoning, however, does not take into account the internal restrictions of consistency of the probability system. These can be best understood in the context of a symmetric network.

For a symmetric network, call the marginal probability of default of each bank q_1 and the pairwise joint probabilities of default of each pair q_2. The previous constraints become:

$$q_1 \leq h$$

$$q_1 - (1-S)q_2 = b$$

where h is the (common) upper bound on the marginal probability of default and b is the (common) b_i.

To maximize systemic risk, we would intuitively set $q_1 = h$, and then q_2 will be set to match CDS spreads:

$$q_2 = \frac{q_1 - b}{1-S} \tag{26}$$

For given q_1, q_2 is increasing in S, as is systemic risk. This captures the intuition that a higher recovery rate of CDSs implies that higher counterparty risk is needed to explain the same bond/CDS basis.

In fact, this effect is at play only when S is small enough. As S grows, q_2 keeps increasing, and at some point it will reach the level $q_2 = q_1$. At that point, the internal consistency of the probability system kicks in, preventing further increases: it would violate the implicit constraint that $q_2 \leq q_1$.

What happens then if S increases further? The only way to satisfy the constraints is to *lower* q_1 *below* h: for $q_1 = h$ there might exist no probability systems able to satisfy both constraints: matching the CDS spread and satisfying internal consistency. Instead, with a lower q_1, it is possible to set q_2 to be equal to q_1 and satisfy the CDS constraint, so that:

$$q_2 = q_1 = \frac{b}{S}$$

which is *decreasing* in S. This means that for large enough values of S, the bond/CDS basis is too large to be explained by counterparty risk. Even at the upper bound on systemic risk, liquidity has to explain part of the basis. In a symmetric system, then, the bounds on systemic risk first increase and then decrease with S.

These forces play out in similar but nonlinear ways for asymmetric networks. In that case, the asymmetry in the bond/CDS basis across banks means that the upper bounds on marginal probabilities (that are obtained from bond prices) will bind for some banks and not for others. The overall effect on the bounds is difficult to describe analytically but can be tested numerically. The results, described in the text, are reported in Table 3.

[47] I focus on the upper bound for the probability of at least $r > 1$ events occurring. Following the analysis reported in section 3, the same argument holds for the *lower* bound for the probability that at least 1 institution defaults, since that is achieved for a very correlated system. It is easy to see why the results for the lower bound for $r > 1$ and the upper bound for $r = 1$ do not depend on S: these bounds look for the least correlated system, which can always be obtained by setting the marginal default probabilities at the levels implied by the CDS spreads and attributing the bond/CDS basis entirely to liquidity.

D.2 - Stochastic recovery rate

This section shows that if recovery rates are stochastic, the change in bond and CDS prices can be decomposed as follows. Starting from the prices obtained in the case of a constant recovery rate R (section 4), the prices are first adjusted by changing the recovery rate to a lower value R_L. In addition, the price of the bond and the present value of CDS payments are shifted in opposite directions by an amount Y (to a first order approximation).

To see this, assume that the recovery rate on bonds is R_H if one bank defaults alone but is R_L if more than one bank defaults at the same time. Call $X_i = \cup_{k \neq i} A_k$ the event of at least one default among the banks different from i, and similarly $X_{ij} = \cup_{k \neq i,j} A_k$. Call $B_R(0,T)$ the price of a bond under the assumption of constant recovery rate R and $B_{R_L,R_H}(0,T)$ the price of a bond with stochastic recovery rate described above, and similarly for the CDS spreads. Then, it is easy to see that (setting liquidity to 0 for simplicity)

$$B_{R_L,R_H}(0,T) = c\left(\sum_{t=1}^{T}\delta(0,t)(1-P(A_i))^t\right) + \delta(0,T)(1-P(A_i))^T +$$

$$+R_H\left(\sum_{t=1}^{T}\delta(0,t)(1-P(A_i))^{t-1}P(A_i\cap\overline{X}_i))\right) + R_L\left(\sum_{t=1}^{T}\delta(0,t)(1-P(A_i))^{t-1}P(A_i\cap X_i)\right)$$

while the CDS spread solves:

$$\sum_{s=1}^{T}\delta(0,s-1)(1-P(A_i\cup A_j))^{s-1}z_{R_L,R_H,ji} =$$

$$=\sum_{s=1}^{T}(1-P(A_i\cup A_j))^{s-1}\delta(0,s)[P(A_i\cap\overline{A}_j\cap\overline{X}_{ij})(1-R_H)$$

$$+P(A_i\cap\overline{A_j}\cap X_{ij})(1-R_L) + P(A_i\cap A_j)(1-R_L)S]$$

Now, rewrite bond prices as:

$$B_{R_L,R_H}(0,T) = c\left(\sum_{t=1}^{T}\delta(0,t)(1-P(A_i))^t\right) + \delta(0,T)(1-P(A_i))^T +$$

$$+(R_H - R_L)\left(\sum_{t=1}^{T}\delta(0,t)(1-P(A_i))^{t-1}P(A_i\cap\overline{X}_i)\right) + R_L\left(\sum_{t=1}^{T}\delta(0,t)(1-P(A_i))^{t-1}\left[P(A_i\cap X_i)+P(A_i\cap\overline{X}_i)\right]\right)$$

Noting that the term $P(A_i\cap X_i)+P(A_i\cap\overline{X}_i) = P(A_i)$, we can rewrite

$$B_{R_L,R_H}(0,T) = c\left(\sum_{t=1}^{T}\delta(0,t)(1-P(A_i))^t\right) +$$

$$+\delta(0,T)(1-P(A_i))^T + R_L\left(\sum_{t=1}^{T}\delta(0,t)(1-P(A_i))^{t-1}P(A_i)\right) + Y_{bond}$$

or:

$$B_{R_L,R_H}(0,T) = B_{R_L}(0,T) + Y_{bond}$$

where

$$Y_{bond} = \sum_{t=1}^{T}\delta(0,t)(1-P(A_i))^{t-1}P(A_i\cap\overline{X}_i)(R_H - R_L)$$

The price of the bond is now equal to the price of the bond in case that the recovery rate is constant and equal to R_L plus the last term, which is the present value of the additional recoveries in case i defaults alone.

A similar fromula holds for CDS spreads. Since $P(A_i \cap \overline{A}_j \cap \overline{X}_{ij}) = P(A_i \cap \overline{X}_i)$, we can rewrite the CDS spread as

$$\sum_{s=1}^{T} \delta(0, s-1)(1 - P(A_i \cup A_j))^{s-1} z_{R_L,R_H,ji} =$$

$$= \sum_{s=1}^{T} (1 - P(A_i \cup A_j))^{s-1} \delta(0,s)[P(A_i \cap \overline{A}_j \cap \overline{X}_{ij})(1 - R_H) - P(A_i \cap \overline{A}_j \cap \overline{X}_{ij})(1 - R_L)$$

$$+ P(A_i \cap \overline{A}_j \cap \overline{X}_{ij})(1 - R_L) + P(A_i \cap \overline{A_j} \cap X_{ij})(1 - R_L) + P(A_i \cap A_j)(1 - R_L)S]$$

$$= \sum_{s=1}^{T} (1 - P(A_i \cup A_j))^{s-1} \delta(0,s)[P(A_i \cap \overline{A}_j \cap \overline{X}_{ij})(R_L - R_H)$$

$$+ (1 - R_L)\left(P(A_i \cap \overline{A}_j \cap \overline{X}_{ij}) + P(A_i \cap \overline{A_j} \cap X_{ij}) + P(A_i \cap A_j)S\right)]$$

Since

$$P(A_i \cap \overline{A}_j \cap \overline{X}_{ij}) + P(A_i \cap \overline{A_j} \cap X_{ij}) + P(A_i \cap A_j)S =$$

$$= P(A_i \cap \overline{A}_j) + P(A_i \cap A_j)S = P(A_i) - (1 - S)P(A_i \cap A_j)$$

we can write:

$$\sum_{s=1}^{T} \delta(0, s-1)(1 - P(A_i \cup A_j))^{s-1} z_{R_L,R_H,ji} =$$

$$= \sum_{s=1}^{T} (1 - P(A_i \cup A_j))^{s-1} \delta(0,s)[P(A_i \cap \overline{A}_j \cap \overline{X}_{ij})(R_L - R_H)$$

$$+ (1 - R_L)(P(A_i) - (1 - S)P(A_i \cap A_j))]$$

$$= \left\{ \sum_{s=1}^{T} (1 - P(A_i \cup A_j))^{s-1} \delta(0,s) \left(P(A_i) - (1 - S)P(A_i \cap A_j)\right)(1 - R_L) \right\} - Y_{CDS}$$

or:

$$\sum_{s=1}^{T} \delta(0, s-1)(1 - P(A_i \cup A_j))^{s-1} z_{R_L,R_H,ji} = \sum_{s=1}^{T} \delta(0, s-1)(1 - P(A_i \cup A_j))^{s-1} z_{R_L,ji} - Y_{CDS}$$

where

$$Y_{CDS} = \sum_{s=1}^{T} (1 - P(A_i \cup A_j))^{s-1} \delta(0,s) P(A_i \cap \overline{X}_i)(R_H - R_L)$$

Now we can show that when probabilities are small, Y_{CDS} and Y_{bond} are approximately the same. To do this, rewrite the two formulas in terms of the following probabilities: $\pi_i = P(A_i \cap \overline{A}_j \cap \overline{X}_{ij}) = P(A_i \cap \overline{X}_i)$, $\pi_j = P(A_j \cap \overline{A}_i)$, $\pi_{ij} = P(A_i \cap A_j)$, $\pi_k = P(A_i \cap \overline{A}_j \cap X_{ij})$. Note that these three sets are disjoint and therefore if the probabilities of the other events are small we can vary them independently of each other.Then,

$$Y_{CDS} = \sum_{s=1}^{T} (1 - \pi_i - \pi_j - \pi_{ij} - \pi_k)^{s-1} \delta(0,s) \pi_i (R_H - R_L)$$

$$Y_{bond} = \sum_{s=1}^{T} (1 - \pi_i - \pi_{ij} - \pi_k)^{s-1} \delta(0,s) \pi_i (R_H - R_L)$$

Approximating each of these equations around $\pi_i = \pi_j = \pi_{ij} = \pi_k = 0$, we note that:

$$Y_{CDS}(0) = 0$$

$$\frac{d}{d\pi_i} Y_{CDS}(0) = \sum_{s=1}^{T} \delta(0,s)(R_H - R_L) \left[\left(\frac{d}{d\pi_i} (1 - \pi_i - \pi_j - \pi_{ij} - \pi_k)^{s-1} \right) \pi_i + \left((1 - \pi_i - \pi_j - \pi_{ij} - \pi_k)^{s-1} \right) \frac{d}{d\pi_i} \pi_i \right]_0$$

$$= \sum_{s=1}^{T} \delta(0,s)(R_H - R_L)$$

and

$$\frac{d}{d\pi_j} Y_{CDS}(0) = \frac{d}{d\pi_{ij}} Y_{CDS}(0) = \frac{d}{d\pi_k} Y_{CDS}(0) = 0$$

And similarly for Y_{bond}: so, to a first approximation, $Y_{CDS} = Y_{bond}$.

D.3 - Alternative pricing model

Suppose that at every time t agents prices bonds assuming for the life of the bond a deterministic hazard rate of the form:

$$h_{t+s} = (1 - \rho_t) \overline{h}_t + \rho_t h_{t+s-1}$$

with a certain $\overline{h}_t, \rho_t$, and h_t determined at time t.

Note that

$$h_{t+s} = (1 - \rho_t)(1 + ... + \rho_t^{s-1}) \overline{h}_t + \rho_t^s h_t$$

$$= (1 - \rho_t) \frac{1 - \rho_t^s}{1 - \rho_t} \overline{h}_t + \rho_t^s h_t$$

$$= (1 - \rho_t^s) \overline{h}_t + \rho_t^s h_t$$

for $s \geq t$.

The probability of surviving until $t + r$ is

$$H_t(t + r; h_t, \rho_t) = (1 - h_t)...(1 - h_{t+r})$$

From the cross section of outstanding bonds, we can then estimate at each t the three parameters $h_t, \rho_t, \overline{h}_t$.

Since CDS spreads depend on the process of joint hazard rate of default, but we only observe the price of the 5-year CDS, I assume that the shape of the joint default hazard rate h_s^{ij} is similar to that of the marginal hazard rate of the reference entity i (similar results hold if we assume that it inherits the shape estimated from the bond prices of the seller, or a combination of both). In particular, after having estimated the three parameters for the hazard rate of bank i, I define

$$\alpha_t^i = \frac{\overline{h}_t^i}{h_t^i}$$

so that:

$$h_{t+s} = (1 - \rho_t^s) \overline{h}_t + \rho_t^s h_t = (1 - \rho_t^s) \alpha_t h_t + \rho_t^s h_t = (\alpha_t - \rho_t^s \alpha_t + \rho_t^s) h_t$$

Call h_t^i the probability that i defaults, similarly for j and finally h_s^{ij} the probability of joint default.

Assume that:

$$h^i_{t+s} = (\alpha_{i,t} - \rho^s_{i,t}\alpha_{i,t} + \rho^s_{i,t})h^i_t$$

$$h^j_{t+s} = (\alpha_{j,t} - \rho^s_{j,t}\alpha_{j,t} + \rho^s_{j,t})h^j_t$$

$$h^{ij}_{t+s} = (\alpha_{ij,t} - \rho^s_{ij,t}\alpha_{ij,t} + \rho^s_{ij,t})h^{ij}_t$$

Note that this requires $h^i_t > 0$ for $\alpha_{i,t}$ to be defined. Therefore, I impose a lower bound on h^i_t of 10^{-6}, or a hundredth of a basis point.

Define

$$H^{ij}_t(t+s; h_t, \rho_t) = (1 - h^i_t - h^j_t + h^{ij}_t)(1 - h^i_{t+1} - h^j_{t+1} + h^{ij}_{t+1})...(1 - h^i_{t+s} - h^j_{t+s} + h^{ij}_{t+s})$$

which is the probability of having no credit events until time $t+s$, and

$$H^{ij}_t(t; \rho_t, h_t) = 1$$

The CDS spread at time t satisfies:

$$\frac{z^{ij}_t}{(1-R)} = \frac{\left[\sum_{s=1}^T \delta(t,t+s)H^{ij}_t(t+s-1;\rho_t,h_t)(h^i_{t+s} - (1-S)h^{ij}_{t+s})\right]}{\sum_{s=1}^T \delta(t,t+s-1)H^{ij}_t(t+s-1;\rho_t,h_t)}$$

We now approximate this formula around $h^i_t = h^j_t = h^{ij}_t = 0$. To do this, we rewrite the formula in terms of probabilities of disjoint events:

$$\pi_i = h^i_t - h^{ij}_t$$

$$\pi_j = h^j_t - h^{ij}_t$$

$$\pi_{ij} = h^{ij}_t$$

So we have:

$$h^i_{t+s} = (\alpha_{i,t} - \rho^s_{i,t}\alpha_{i,t} + \rho^s_{i,t})(\pi_i + \pi_{ij})$$

$$h^j_{t+s} = (\alpha_{j,t} - \rho^s_{j,t}\alpha_{j,t} + \rho^s_{j,t})(\pi_j + \pi_{ij})$$

$$h^{ij}_{t+s} = (\alpha_{ij,t} - \rho^s_{ij,t}\alpha_{ij,t} + \rho^s_{ij,t})\pi_{ij}$$

Call

$$G = \frac{\left[\sum_{s=1}^T \delta(t,t+s)H^{ij}_t(t+s-1;\rho_t,h_t)(h^i_{t+s} - (1-S)h^{ij}_{t+s})\right]}{\sum_{s=1}^T \delta(t,t+s-1)H^{ij}_t(t+s-1;\rho_t,h_t)}$$

We can start by noting that at the approximation point, $\pi_i = \pi_j = \pi_{ij} = 0$, we have:

$$(h^i_{t+s} - (1-S)h^{ij}_{t+s})|_0 = 0$$

$$H^{ij}_t(t+s-1;\rho_t,h_t)|_0 = 1$$

We then have, for each of the π

$$\frac{d}{d\pi}G|_0 = \frac{\frac{d}{d\pi}\left[\sum_{s=1}^T \delta(t,t+s)H^{ij}_t(t+s-1;\rho_t,h_t)(h^i_{t+s} - (1-S)h^{ij}_{t+s})\right]\left[\sum_{s=1}^T \delta(t,t+s-1)H^{ij}_t(t+s-1;\rho_t,h_t)\right]}{\left[\sum_{s=1}^T \delta(t,t+s-1)H^{ij}_t(t+s-1;\rho_t,h_t)\right]^2}$$

$$-\frac{\left[\sum_{s=1}^{T}\delta(t,t+s)H_t^{ij}(t+s-1;\rho_t,h_t)(h_{t+s}^i-(1-S)h_{t+s}^{ij})\right]\frac{d}{d\pi}\left[\sum_{s=1}^{T}\delta(t,t+s-1)H_t^{ij}(t+s-1;\rho_t,h_t)\right]}{\left[\sum_{s=1}^{T}\delta(t,t+s-1)H_t^{ij}(t+s-1;\rho_t,h_t)\right]^2}|_0$$

The second term, at the approximation point, is always 0. We can reduce the derivative to:

$$\frac{d}{d\pi}G|_0=\frac{\sum_{s=1}^{T}\delta(t,t+s)\left\{\left[\frac{d}{d\pi}H_t^{ij}(t+s-1;\rho_t,h_t)\right](h_{t+s}^i-(1-S)h_{t+s}^{ij})+H_t^{ij}(t+s-1;\rho_t,h_t)\frac{d}{d\pi}(h_{t+s}^i-(1-S)h_{t+s}^{ij})\right\}}{\sum_{s=1}^{T}\delta(t,t+s-1)}|_0$$

which in turn becomes

$$\frac{d}{d\pi}G|_0=\frac{\sum_{s=1}^{T}\delta(t,t+s)\frac{d}{d\pi}(h_{t+s}^i-(1-S)h_{t+s}^{ij})}{\sum_{s=1}^{T}\delta(t,t+s-1)}|_0$$

Now,

$$h_{t+s}^i=(\alpha_{i,t}-\rho_{i,t}^s\alpha_{i,t}+\rho_{i,t}^s)(\pi_i+\pi_{ij})$$

$$h_{t+s}^j=(\alpha_{j,t}-\rho_{j,t}^s\alpha_{j,t}+\rho_{j,t}^s)(\pi_j+\pi_{ij})$$

$$h_{t+s}^{ij}=(\alpha_{ij,t}-\rho_{ij,t}^s\alpha_{ij,t}+\rho_{ij,t}^s)\pi_{ij}$$

$$\frac{d}{d\pi}(h_{t+s}^i-(1-S)h_{t+s}^{ij})=$$

$$=\frac{d}{d\pi}((\alpha_{i,t}-\rho_{i,t}^s\alpha_{i,t}+\rho_{i,t}^s)(\pi_i+\pi_{ij})-(1-S)(\alpha_{ij,t}-\rho_{ij,t}^s\alpha_{ij,t}+\rho_{ij,t}^s)\pi_{ij})$$

So that:

$$\frac{d}{d\pi_i}(h_{t+s}^i-(1-S)h_{t+s}^{ij})=(\alpha_{i,t}-\rho_{i,t}^s\alpha_{i,t}+\rho_{i,t}^s)$$

and

$$\frac{d}{d\pi_{ij}}(h_{t+s}^i-(1-S)h_{t+s}^{ij})=(\alpha_{i,t}-\rho_{i,t}^s\alpha_{i,t}+\rho_{i,t}^s)-(1-S)(\alpha_{ij,t}-\rho_{ij,t}^s\alpha_{ij,t}+\rho_{ij,t}^s)$$

$$\frac{d}{d\pi_j}(h_{t+s}^i-(1-S)h_{t+s}^{ij})=0$$

Therefore we have:

$$\frac{d}{d\pi}G|_0=\frac{\sum_{s=1}^{T}\delta(t,t+s)\frac{d}{d\pi}(h_{t+s}^i-(1-S)h_{t+s}^{ij})}{\sum_{s=1}^{T}\delta(t,t+s-1)}|_0$$

$$G\simeq\frac{\sum_{s=1}^{T}\delta(t,t+s)\left[(\alpha_{i,t}-\rho_{i,t}^s\alpha_{i,t}+\rho_{i,t}^s)\right]\pi_i+\left[(\alpha_{i,t}-\rho_{i,t}^s\alpha_{i,t}+\rho_{i,t}^s)-(1-S)(\alpha_{ij,t}-\rho_{ij,t}^s\alpha_{ij,t}+\rho_{ij,t}^s)\right]\pi_{ij}}{\sum_{s=1}^{T}\delta(t,t+s-1)}$$

$$=\frac{\sum_{s=1}^{T}\delta(t,t+s)(\alpha_{i,t}-\rho_{i,t}^s\alpha_{i,t}+\rho_{i,t}^s)h_t^i-(1-S)(\alpha_{ij,t}-\rho_{ij,t}^s\alpha_{ij,t}+\rho_{ij,t}^s)h_t^{ij}}{\sum_{s=1}^{T}\delta(t,t+s-1)}$$

Calling

$$G_i=\frac{\sum_{s=1}^{T}\delta(t,t+s)(\alpha_{i,t}-\rho_{i,t}^s\alpha_{i,t}+\rho_{i,t}^s)}{\sum_{s=1}^{T}\delta(t,t+s-1)}$$

and

$$G_{ij}=\frac{\sum_{s=1}^{T}\delta(t,t+s)(\alpha_{ij,t}-\rho_{ij,t}^s\alpha_{ij,t}+\rho_{ij,t}^s)}{\sum_{s=1}^{T}\delta(t,t+s-1)}$$

we obtain:

$$\frac{z_t^{ij}}{(1-R)}\simeq G_ih_t^i-(1-S)G_{ij}h_t^{ij}$$

Summing over counterparties, we have:

$$\frac{(N-1)\bar{z}_t^i}{G_i(1-R)} \simeq (N-1)h_t^i - (1-S)\sum_{j\neq i}\frac{G_{ij}}{G_i}h_t^{ij}$$

which is again linear in the marginal and pairwise default probabilities and can be used in the LP formulation. To impose that CDS spreads inherit the shape of the hazard rates estimated from bonds, we can assume $\alpha_{ij,t} = \alpha_{i,t}$ and $\rho_{ij,t} = \rho_{i,t}$. Note that this reduces the CDS spread to a function of only one parameter: h_t^{ij}.

D.4 Currency assumptions

To simplify, consider one-period bonds and CDSs denominated in different currencies. Call m_{se} the stochastic discount factor of a US investor in state (s, e) where s indicates whether default of firm i ($s = i$ or 0) occurs or not and e is the exchange rate. Assume that international bond and CDS markets are not segmented, so that m prices all assets. Call $f(s, e)$ the joint density function s and e. To simplify notation, rewrite $f(s, e)$ as:

$$f(s, e) = \pi_s f_s(e)$$

Note that

$$\pi_0 E[m_{se}|s = 0] + \pi_i E[m_{se}|s = i] = E[m_{se}]$$

For a dollar-denominated risky bond (R is the recovery rate), the dollar price is:

$$p_i^{\$} = \pi_0 E[m_{se}|s = 0] + R\pi_i E[m_{se}|s = i] = E[m_{se}] - (1-R)\pi_i E[m_{se}|s = i]$$

Now consider a euro-denominated bond issued by the same firm, and of equal seniority. Calling e_0 the time-0 exchange rate, we obtain:

$$p_i^E e_0 = \pi_0 E[e \cdot m_{0e}|s = 0] + R\pi_i E[e \cdot m_{ie}|s = i]$$

$$= E[e \cdot m_{se}] - (1-R)\pi_i E[e \cdot m_{se}|s = i]$$

The prices of the respective risk-free securities are:

$$t^{\$} = E[m_{se}]$$

$$t^E e_0 = E[e \cdot m_{se}]$$

Combining defaultable and risk-free bonds we get:

$$p_i^{\$} = t^{\$}\left(1 - (1-R)\pi_i \frac{E[m_{se}|s = i]}{E[m]}\right)$$

$$p_i^E e_0 = t^E e_0 \left(1 - (1-R)\pi_i \frac{E[e \cdot m_{se}|s = i]}{E[e \cdot m_{se}]}\right)$$

We can then use either bond to estimate the risk-neutral probability of default of firm i

$$P(A_i) = \pi_i \frac{E[m_{se}|s = i]}{E[m_{se}]}$$

discounting cash flows by the appropriate risk-free rate as long as the following condition holds:

$$\frac{E[e \cdot m_{se}|s=i]}{E[m_{se}|s=i]} = \frac{E[e \cdot m_{se}]}{E[m_{se}]}$$

which requires that the relative price of domestic and foreign risk-free securities is the same as the relative price of domestic and foreign state-contingent securities that pay off if i defaults.

Now, consider the case of a CDS written by one bank on the default of another bank. The CDS is written on a european bank (i) but the counterparty (j) is american. The contract is denominated in euros. In this case, s represents all the combinations of default of the two banks, and can be i (only i defaults), j, ij (both default), and 0 (none defaults).

The CDS contract costs z_{ji} euros. So we must have

$$z_{ji}e_0 = E[e \cdot m_{se}] \left((1-R)\pi_i \frac{E\left[e \cdot m_{se}|s=i\right]}{E[e \cdot m_{se}]} + (1-R)S\pi_{ij} \frac{E\left[e \cdot m_{se}|s=ij\right]}{E[e \cdot m_{se}]} \right)$$

Therefore, as long as the european yield curve is used to discount cash flows for euro-denominated CDSs the sufficient condition is:

$$\frac{E[e \cdot m_{se}|s]}{E[m_{se}|s]} = \frac{E[e \cdot m_{se}]}{E[m_{se}]}$$

for every default event in s.

Tables

Table 1

	Avg valid bonds	2004	2005	2006	2007	2008	2009	2010
Abn Amro	3.3	1.8	2.2	4.0	4.4	3.0	3.5	5.3
Bank of America	32.3	17.5	25.4	29.3	32.9	35.3	41.8	55.8
Barclays	14.8	3.1	3.0	2.4	2.5	9.1	38.5	78.0
Bear Stearns	11.4	7.2	9.8	12.6	15.3	15.4	-	-
Bnp Paribas	7.0	0.5	2.0	3.0	3.9	6.7	18.0	22.6
Citigroup	36.5	21.6	24.3	31.7	40.0	43.2	49.5	54.5
Credit Suisse	5.4	1.9	2.3	2.8	2.7	5.0	11.6	17.4
Deutsche Bank	42.1	5.3	10.4	42.3	68.9	54.4	58.8	67.9
Goldman Sachs	39.6	19.3	26.1	34.3	40.4	49.0	57.0	63.0
JP Morgan	17.3	6.6	11.1	14.0	17.4	22.2	27.4	27.9
Lehman Brothers	20.1	10.5	15.2	20.5	26.5	31.4	-	-
Merrill Lynch	35.7	22.7	33.0	38.4	43.0	44.0	-	-
Morgan Stanley	25.5	12.5	14.6	17.5	22.2	30.0	45.0	49.0
UBS	8.2	0.3	0.7	1.0	3.1	8.3	22.4	36.6
Wachovia	6.1	2.9	3.5	5.7	7.4	9.1	7.7	7.3

Note: first column reports average number of bonds for each institution that are used for the estimation of marginal default probabilities. Columns 2-8 break this number down by year.

Table 2

	Avg CDS spread	Std CDS spread	Min spread	Max spread	Avg basis	Std basis	Min basis	Max basis
Abn Amro	45.8	46.1	5.0	190.5	-46.2	44.4	-248.2	34.6
Bank of America	66.5	71.7	7.4	390.7	-71.9	64.8	-412.5	217.5
Barclays	54.3	60.0	5.5	261.9	-41.2	60.1	-324.8	111.6
Bear Stearns	54.2	69.7	18.0	736.9	-53.6	24.6	-298.0	40.6
Bnp Paribas	33.8	32.2	5.4	163.9	-53.9	49.7	-321.8	86.8
Citigroup	100.4	129.7	6.5	638.3	-76.5	88.1	-804.6	59.2
Credit Suisse	53.0	51.3	9.0	261.4	-50.5	44.0	-276.6	52.6
Deutsche Bank	49.8	45.1	8.9	190.0	-24.4	27.6	-174.2	65.2
Goldman Sachs	84.2	86.4	17.2	579.3	-79.0	91.3	-502.4	75.1
JP Morgan	53.1	42.8	10.9	227.3	-76.5	57.9	-322.2	32.1
Lehman Brothers	70.7	86.9	18.0	701.7	-61.9	44.5	-540.0	10.9
Merrill Lynch	59.9	71.9	14.4	447.7	-51.7	40.4	-200.0	26.2
Morgan Stanley	112.5	144.2	16.6	1385.6	-82.6	107.3	-1256.5	223.2
UBS	59.3	72.4	4.2	357.2	-65.4	58.5	-343.5	34.2
Wachovia	73.9	93.5	9.3	1487.7	-87.2	120.1	-2509.8	88.0

Note: table reports descriptive characteristics on the CDS spread and the yield spread for the 15 institutions in the sample, in basis points per year. The yield spread is computed as the linearly interpolated yield for a 5-year maturity bond in excess of the corresponding Treasury rate.

Table 3a: max P1

Model			Average level of the bounds (bp per month)					
			2007	Jan 2008 to Bear	Bear to Lehman	Month after Lehman	Oct 2008 to April 2009	After April 2009
		Start	1/1/07	1/1/08	3/16/08	9/15/08	10/16/08	5/1/09
R	S	End	12/31/07	3/15/08	9/12/08	10/15/08	4/30/09	6/10/10
0.10	0.10		50.4	178.0	168.8	298.1	221.7	133.0
0.10	0.30		50.4	178.0	168.8	298.1	221.7	133.0
0.10	0.40		50.4	178.0	168.8	298.1	221.7	133.0
0.10	0.70		50.4	178.0	168.8	298.1	221.7	133.0
0.10	0.90		50.4	178.0	168.8	298.1	221.7	133.0
0.10	1.00		50.4	178.0	168.8	298.1	221.7	133.0
0.30	**0.30**		**64.8**	**228.9**	**217.0**	**383.3**	**285.0**	**171.1**
0.30	0.40		64.8	228.9	217.0	383.3	285.0	171.1
0.30	0.70		64.8	228.9	217.0	383.3	285.0	171.1
0.30	0.90		64.8	228.9	217.0	383.3	285.0	171.1
0.30	1.00		64.8	228.9	217.0	383.3	285.0	171.1
0.40	0.40		75.6	267.1	253.2	447.1	332.5	199.6
0.40	0.70		75.6	267.1	253.2	447.1	332.5	199.6
0.40	0.90		75.6	267.1	253.2	447.1	332.5	199.6
0.40	1.00		75.6	267.1	253.2	447.1	332.5	199.6
Using swap rates			64.8	229.0	217.1	383.6	285.2	171.1
US banks			49.4	166.7	156.7	278.7	182.7	102.2
US banks, larger trans			46.5	166.4	156.6	278.7	183.3	96.5
Reweight top 5 banks			65.0	228.9	217.0	383.3	285.8	171.3
Reweight, decreasing			65.0	228.9	217.0	383.3	285.8	171.3
Alternative bond model			35.2	112.0	162.2	579.0	216.9	51.9

Table 3b: max P4

Model			Average level of the bounds (bp per month)					
			2007	Jan 2008 to Bear	Bear to Lehman	Month after Lehman	Oct 2008 to April 2009	After April 2009
		Start	1/1/07	1/1/08	3/16/08	9/15/08	10/16/08	5/1/09
R	S	End	12/31/07	3/15/08	9/12/08	10/15/08	4/30/09	6/10/10
0.10	0.10		2.2	2.3	15.7	24.3	49.1	30.9
0.10	0.30		2.4	2.5	17.3	25.7	48.8	31.2
0.10	0.40		2.5	2.6	18.2	26.2	48.6	31.2
0.10	0.70		2.8	2.9	20.0	27.7	47.2	30.6
0.10	0.90		3.3	3.3	21.5	28.4	46.0	29.7
0.10	1.00		12.6	44.5	42.2	70.0	55.4	33.3
0.30	**0.30**		**3.3**	**3.5**	**24.7**	**42.1**	**64.8**	**40.6**
0.30	0.40		3.4	3.7	25.6	42.5	64.2	40.5
0.30	0.70		3.7	4.1	27.9	43.7	62.0	39.6
0.30	0.90		4.1	4.6	29.5	44.3	60.6	38.3
0.30	1.00		16.2	57.2	54.3	90.0	71.2	42.8
0.40	0.40		4.0	4.4	32.0	54.0	76.8	47.5
0.40	0.70		4.3	5.0	34.7	54.0	74.0	46.3
0.40	0.90		4.8	5.7	36.1	53.3	72.0	44.8
0.40	1.00		18.9	66.8	63.3	104.9	83.1	49.9
Using swap rates			2.3	2.6	19.3	42.5	58.1	28.0
US banks			1.3	0.8	11.3	36.9	36.0	17.0
US banks, larger trans			1.4	1.0	16.6	39.6	42.1	16.6
Reweight top 5 banks			5.3	5.0	32.4	50.3	74.5	44.7
Reweight, decreasing			5.2	5.1	32.9	50.1	75.6	44.8
Alternative bond model			2.6	4.7	18.7	49.5	36.7	10.1

Table 3c: min P1

Model			Average level of the bounds (bp per month)					
			2007	Jan 2008 to Bear	Bear to Lehman	Month after Lehman	Oct 2008 to April 2009	After April 2009
		Start	1/1/07	1/1/08	3/16/08	9/15/08	10/16/08	5/1/09
R	S	End	12/31/07	3/15/08	9/12/08	10/15/08	4/30/09	6/10/10
0.10	0.10		42.6	164.9	121.7	221.9	109.8	62.3
0.10	0.30		42.0	164.3	117.4	217.3	101.3	56.1
0.10	0.40		41.8	163.8	115.0	214.7	97.0	53.0
0.10	0.70		40.9	162.2	107.2	205.8	85.8	44.4
0.10	0.90		39.1	160.6	101.1	202.8	80.6	40.3
0.10	1.00		7.4	27.3	25.1	84.7	41.3	22.6
0.30	**0.30**		**53.9**	**210.3**	**146.4**	**252.2**	**125.9**	**71.3**
0.30	0.40		53.6	209.5	143.1	248.0	120.4	67.3
0.30	0.70		52.5	206.6	132.9	237.4	106.5	56.6
0.30	0.90		50.9	204.8	125.3	232.8	99.4	51.6
0.30	1.00		9.5	35.1	32.2	108.9	53.1	29.0
0.40	0.40		62.4	243.7	161.9	281.8	135.8	77.9
0.40	0.70		61.2	241.3	149.8	273.5	120.1	65.6
0.40	0.90		59.7	238.9	142.3	268.4	111.1	59.7
0.40	1.00		11.1	40.9	37.6	127.0	61.9	33.8
Using swap rates			56.6	217.0	156.1	258.2	141.5	97.8
US banks			43.3	157.9	117.0	187.0	110.6	62.3
US banks, larger trans			39.8	154.1	109.8	168.8	100.6	57.4
Reweight top 5 banks			50.4	204.8	128.2	243.5	121.5	67.5
Reweight, decreasing			50.9	204.4	128.2	249.7	121.1	64.6
Alternative bond model			25.2	92.4	107.5	397.8	121.3	26.6

Note: Table reports the average value of the bounds on monthly P1 (probability that at least one bank defaults) and P4 (probability that at least four banks default) for different nonoverlapping periods, under different assumptions discussed in the text. The lower bound for P4 is 0 throughout.

Figures

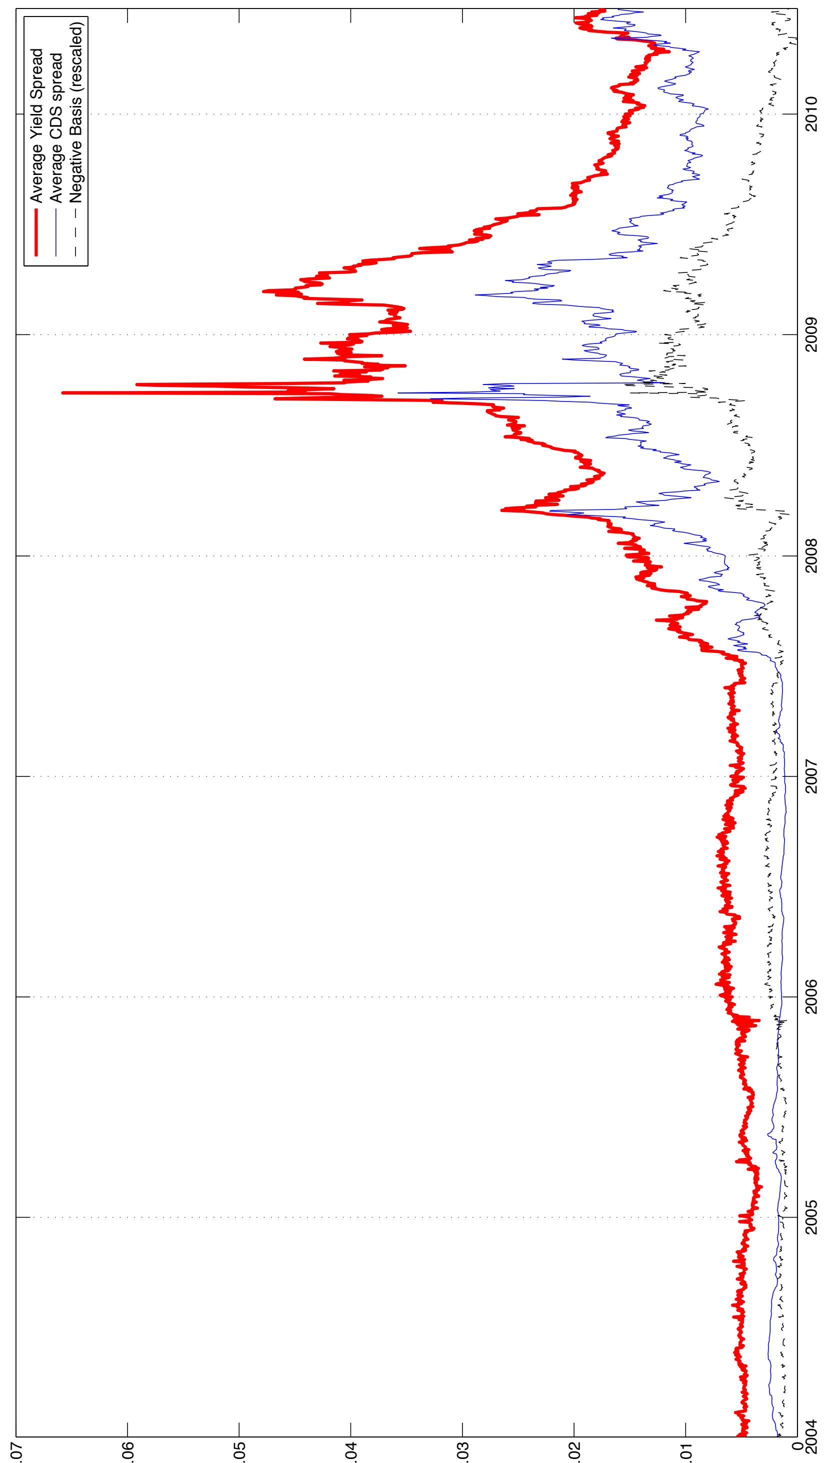

Figure 1: Average 5-year bond yield and CDS spread for the 15 financial intermediaries most active in the CDS market, together with the (rescaled) difference between the two (the negative bond/CDS basis). The bond yield is obtained by interpolating bond yields at the 5 year maturity.

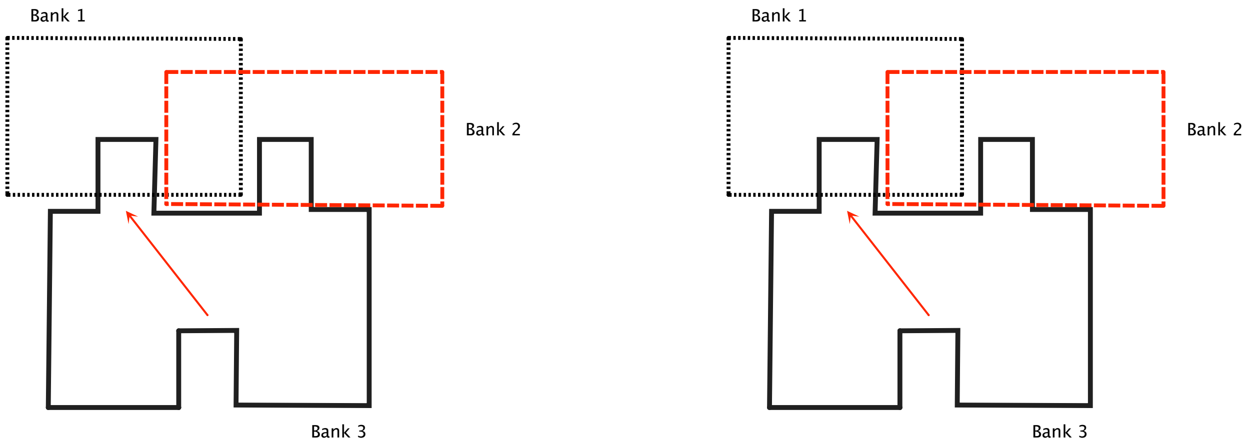

Figure 2: Example of the relation between low-order and high-order probabilities. The sets in the Venn diagram represent default events, and their areas represents the default probabilities.

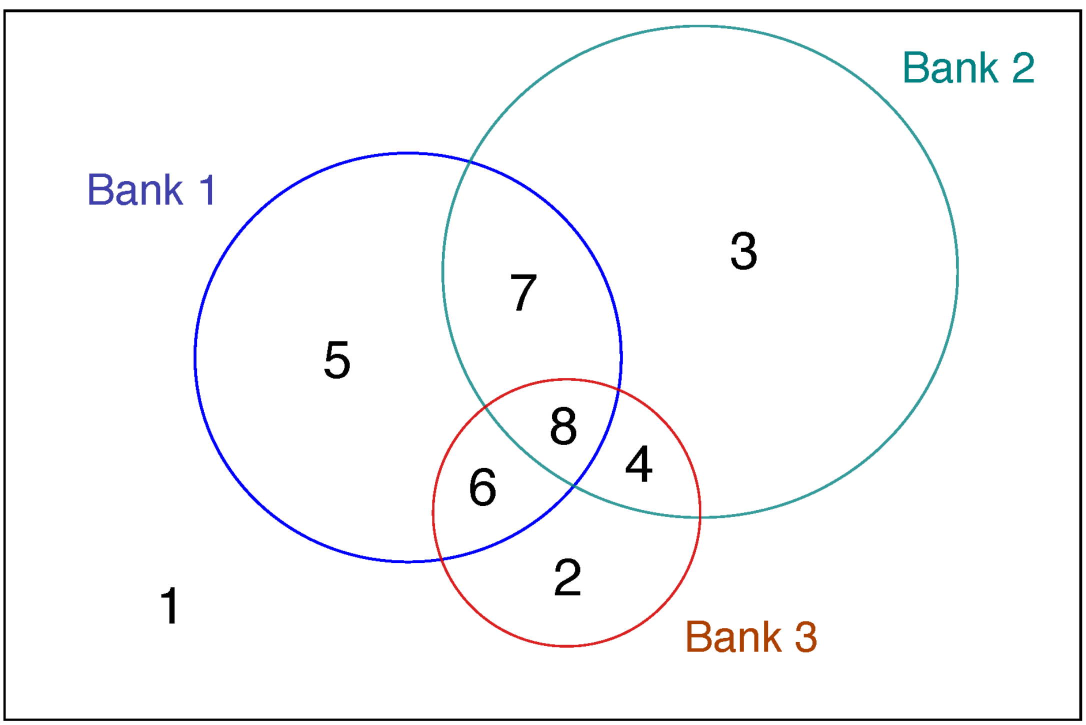

Figure 3: Construction of the Linear Programming representation according to Propositions 1 and 2.

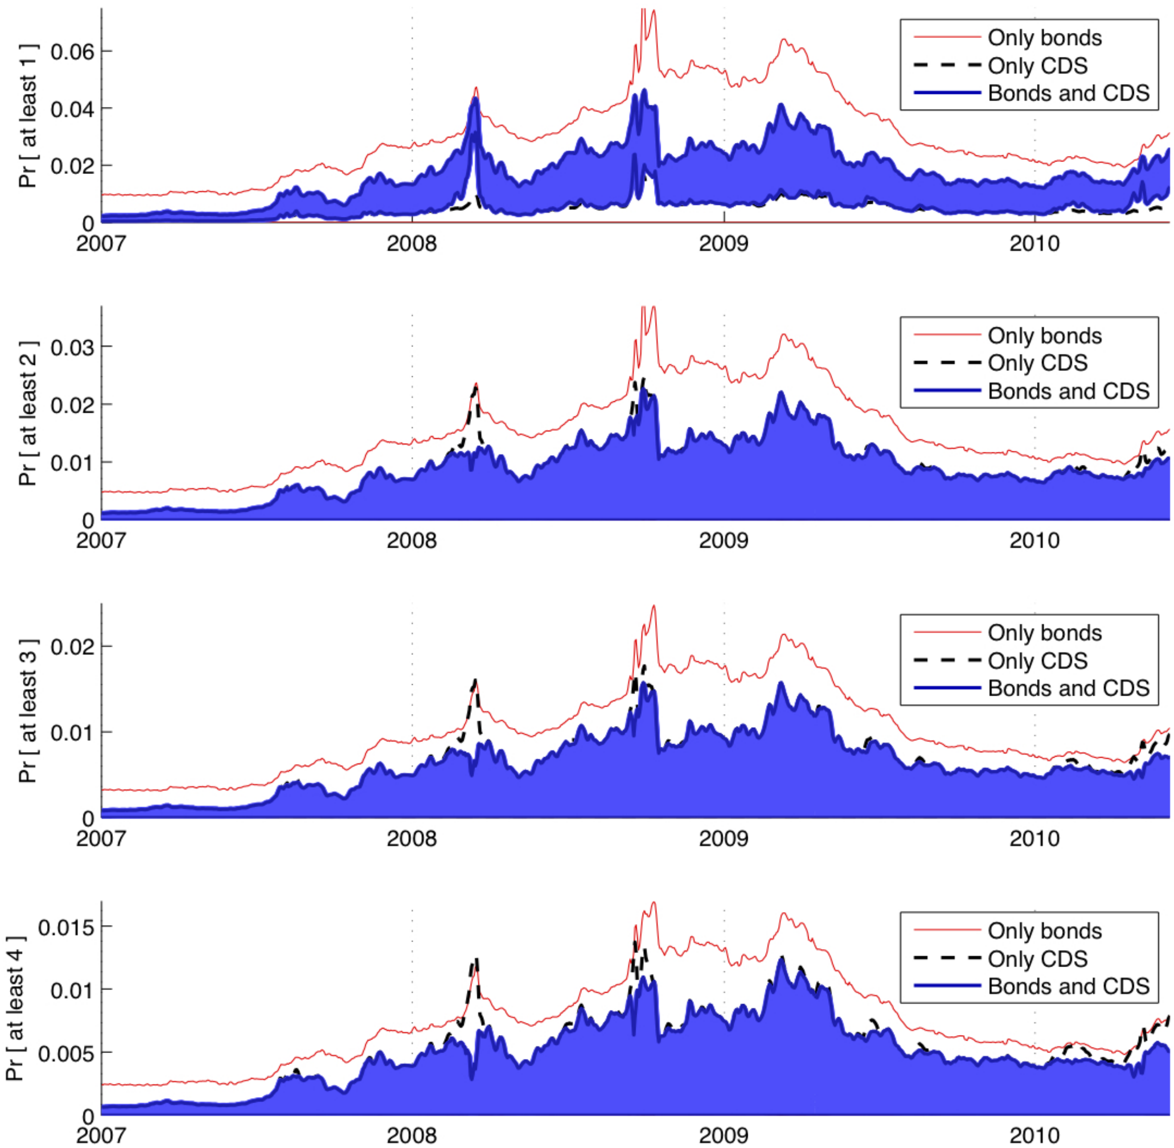

Figure 4: Upper and lower bounds on different systemic events under the assumption of nonnegative liquidity premia: average monthly probability of at least r banks defaulting, for $r = 1, 2, 3, 4$. The Figure reports bounds obtained using only information in bond prices (thin lines), only CDS spreads (dotted lines), and all information (shaded area). To reduce noise, I plot a 3 day moving average of each bound. Note that all lower bounds are 0 for $r > 1$, and that the CDS-only bounds often coincide with the full information bounds (that use both bonds and CDSs).

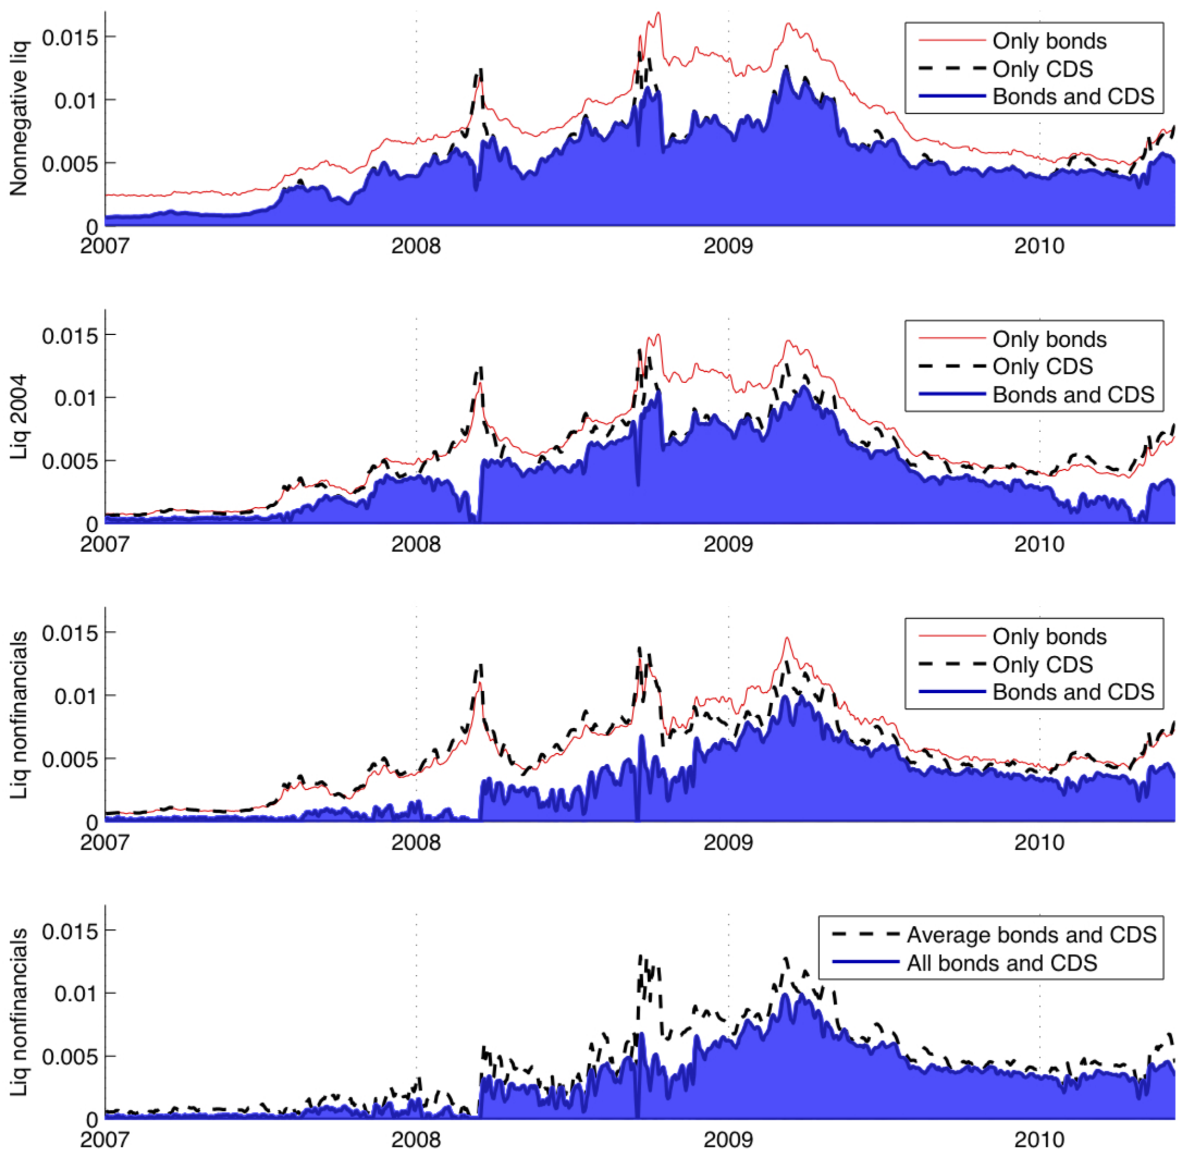

Figure 5: Bounds on the average monthly probability of at least 4 banks defaulting under nonnegative liquidity premia (first panel), liquidity premia at least as high as in 2004 (second panel), liquidity premia calibrated to the ones of nonfinancial firms (third panel). In the top three panels, I report bounds obtained using only information in bond prices (thin lines), only CDS spreads (dotted lines), and all information (shaded area). The last panel reports bounds that use all information (shaded area) and bounds that use only average information (dotted lines). To reduce noise, I plot a 3 day moving average of each bound.

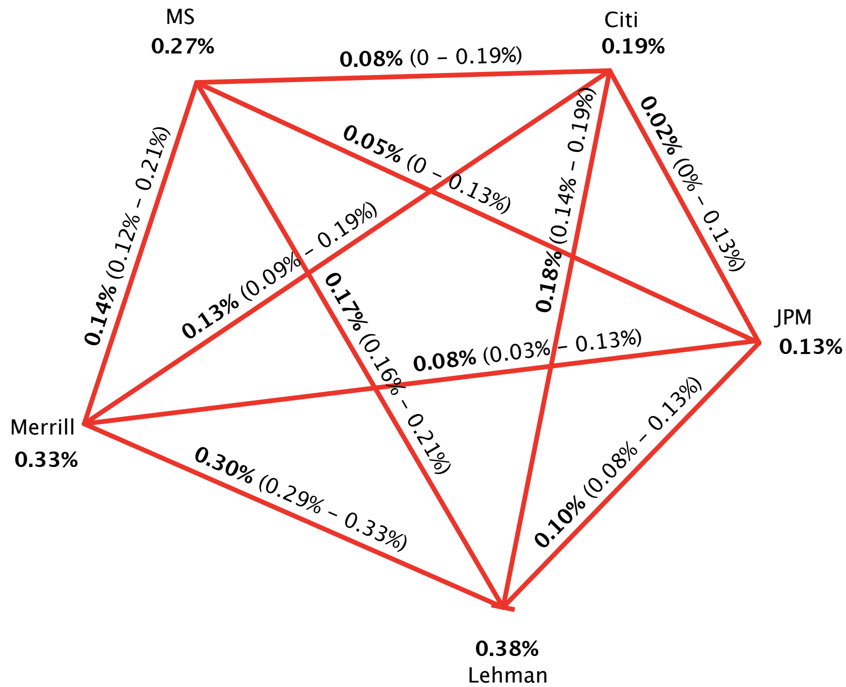

Figure 6: Marginal and pairwise average monthly default probabilities for part of the network in the high systemic risk scenario (max P_4) as of 08/06/2008, with the liquidity process calibrated to that of nonfinancial firms. Nodes report the marginal default probabilities of each banks, edges report pairwise joint default probabilities. The numbers in bold represent the probabilities in the "representative" solution to the bounds, as described in section 3. The numbers in parentheses report the possible range of probabilities across all probability systems that attain the upper bound on P_4.

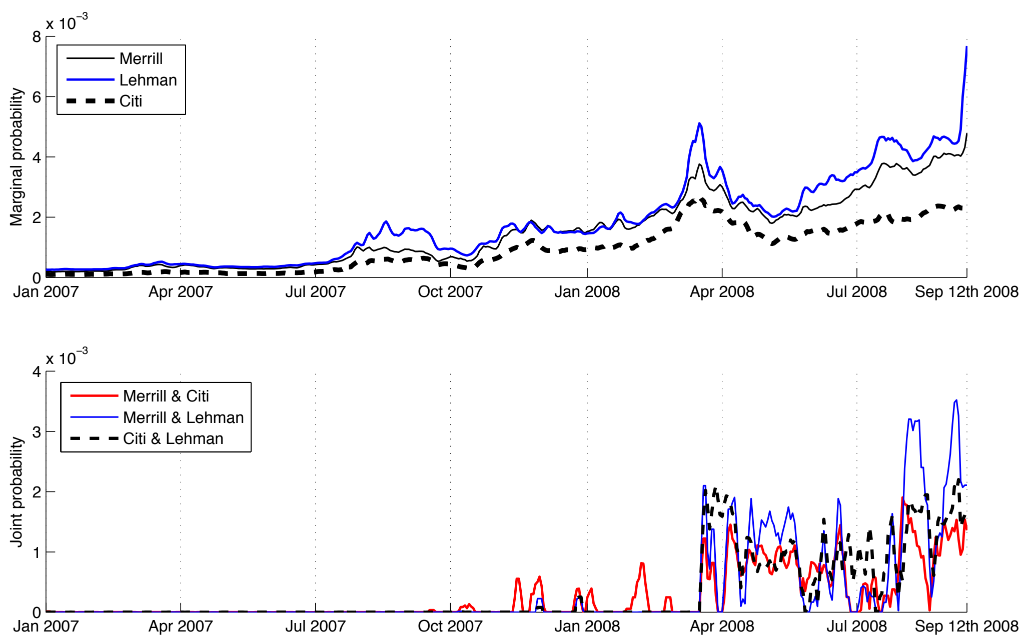

Figure 7: Marginal and pairwise default probabilities for selected banks, in the "representative" solution to the upper bound for the monthly probability that at least four banks default, under the liquidity process calibration to nonfinancial firms. To reduce noise, I plot a 3 day moving average of each series.

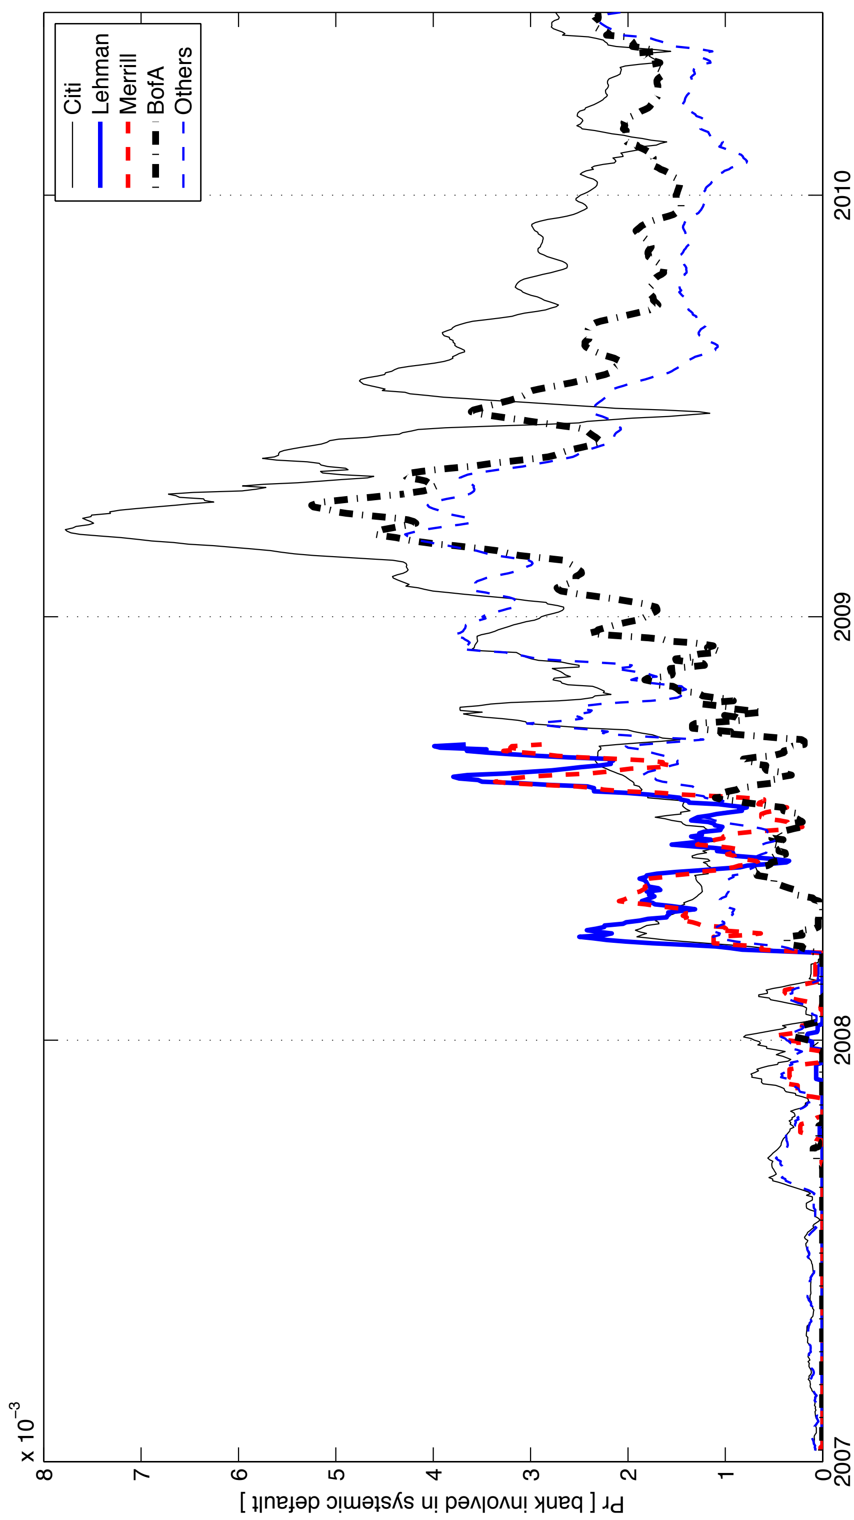

Figure 8: Individual contributions to systemic risk for selected banks, under the calibration of the liquidity process to that of nonfinancial firms, at the upper bound for the probability that at least 4 banks default. The Figure plots for each bank j the average monthly probability that a systemic default event occurs (at least 4 banks default) and institution j is one of them. To reduce noise, I plot a two week moving average of each series.