CDS as Insurance: Leaky Lifeboats in Stormy Seas

Eric Stephens
University of Alberta
eric.stephens@ualberta.ca

James R. Thompson*
University of Waterloo
james@uwaterloo.ca

February 2014

Abstract

What market features of financial risk transfer exacerbate counterparty risk? To analyze this, we formulate a model which elucidates important differences between financial risk transfer and traditional insurance, using the example of Credit Default Swaps (CDS). We allow for (heterogeneous) insurer insolvency, which captures the possibility that relatively risky counterparties may exist in the market. Further, we find that stable insurers become less stable as the price of the contract decreases. The analysis includes insured parties that have heterogeneous motivations for purchasing CDS. For example, some may own the underlying asset and purchase CDS for risk management, while others buy these contracts purely for trading purposes. We show that traders will choose to contract with less stable insurers, resulting in higher counterparty risk in this market relative to that of traditional insurance; however, a regulatory policy that removes traders can, perversely, cause stable counterparties to become less stable. We conclude with two extensions of the model that consider a Central Counterparty (CCP) arrangement and the consequences of asymmetric information over insurer type.

Keywords: financial risk transfer, counterparty risk, insurance, credit default swaps, regulation.

*We are grateful to Philip Bond, Neil Brisley, Alex Edmans, Ing-Haw Cheng, Stefano Giglio, Mike Hoy, Thor Koeppl, Af Mirza, Craig Pirrong, Adriano Rampini, Jano Zabojnik and seminar participants at the 2011 WFA (Sante Fe), 2011 FIRS (Sydney), 2011 LSE Paul Wooley Centre Conference, 2011 Risk Theory Society (Little Rock), Bank of Canada, Reserve Bank of New Zealand, University of Calgary (Haskayne), Duke University (Fuqua), Center for Economics Studies (Munich), University of Toronto, and University of Alberta for helpful comments.

1 Introduction

Counterparty risk has received considerable media attention since the beginning of the credit crisis in 2007. For example, there was public outrage over the use of U.S. tax payer money to pay (in full) the Credit Default Swap (CDS) claims that sellers, such as AIG had sold to many major banks.[1] In response to this and other episodes, policy-makers have been under pressure to implement regulatory reforms to mitigate counterparty risk. In the State of New York for example, legislation was tabled in 2008 to have CDS sellers classified and regulated as insurers.[2] While these issues are of widespread interest, the discourse is lacking in theoretical perspective. Our analysis provides a simple framework which shows how counterparty risk may increase in these markets and illustrates important differences between these and traditional insurance contracts. We focus on CDS since this is the most straightforward example, however the model is kept sufficiently general so that the results are applicable to a variety of financial contracts when they contain counterparty risk.

We capture the pervasiveness of insurer instability by allowing for the possibility that sellers of CDS contracts become insolvent. We consider two insurer types: a stable ('good') insurer and an unstable ('bad') insurer. When a contract is written between an insured party and an insurer, the insurer determines how to invest the insurance premium - between a liquid, low return asset which is always available to pay claims, and an illiquid, high return asset which is never available to pay claims. The bad insurer is characterized by assets in place (e.g., a portfolio) that can take either a high or low value. The probability of failure of this insurer is then independent of the investment choice so that the illiquid asset is favored whenever the return of that investment sufficiently high. The good insurer has more stable assets in place such that investment choice can affect the probability of failure. As such, this insurer invests at least some of the premium in the liquid asset to reduce the probability of failure when a claim is made. Given bankruptcy costs, and since the contract can affect the probability of failure of the good insurer but not the bad, the bad insurer is able to charge a cheaper price (premium) than the good.

We find that bad insurers can exist in equilibrium, even though all parties know that they are unstable. To show this, we allow the insured party to choose to contract with either the good or the bad insurer. This choice boils down to a tradeoff between the price and the degree of exposure to counterparty risk (probability of insurer insolvency). The resulting equilibrium can have good or bad insurers dominate the market. When the insured party is sufficiently averse to counterparty

[1] In a credit default swap, an insurer agrees to cover the losses of the insured if pre-defined credit events (e.g., default) happen to some debt instrument. In exchange, the insured agrees to pay an ongoing premium at fixed intervals for the life of the contract. A CDS written on the debt of a single company is typically bought and sold through a dealer. When the underlying debt is more complicated (and so requires a non-standard contract), the CDS is completed directly between the two parties. For example, the CDS contracts that destabilized AIG were mainly direct contracts with major banks, written on complex mortgage related securities. The estimated notional size of the CDS market in 1998 was 180 billion dollars, by 2004 this number had grown to 6 trillion, and by the end of 2008 it was 41 trillion dollars (Stulz 2009). Note that this is a notional amount and no doubt overestimates the absolute economic value of all contracts, but the relative growth has been rapid.

[2] http://ins.state.ny.us/circltr/2008/cl08_19.htm

risk, as we would expect in a standard insurance model, the bad insurer will not be able to cut its premium enough to attract the insured party. In this case, only the good insurer exists, and can extract positive profits. When the insured party has little aversion to counterparty risk, the bad insurer may control the market as insured parties become premium driven, rather than counterparty risk driven. When there are sufficiently large social costs to counterparty risk, the choice of the bad insurer in equilibrium is inefficient since insured parties do not internalize the social cost. We show in equilibrium that as prices fall, e.g., due to competition, counterparty risk increases since the total amount that the good insurer invests in the liquid asset goes down as the premium decreases. This result is similar in spirit to one found in the banking literature in which stability of the system can decrease when competition among banks increases (see Boyd and De Nicoló, 2005; or Vives, 2010, for a summary of this literature).

We then incorporate a key feature of the CDS market into the model - insured parties that can differ on their motivation to insure. To do this, we allow these parties to have different aversion to counterparty risk. Recall that the two key factors in the choice of insurer are counterparty risk and premium. Those insured parties who use CDS purely for trading purposes, and perhaps do not even own the underlying asset (i.e., have no insurable interest), are more likely premium driven. On the other hand, buyers who use CDS for risk management/hedging would internalize the counterparty risk more, and would be willing to pay relatively more for stable protection. Given that traditional insurance markets are usually viewed as having risk averse insured parties, the analogue to CDS would be a market composed entirely of buyers using the contracts for risk management purposes. As more participants use CDS purely for trading purposes, we find that the market will be serviced more by unstable insurers. This is because the traders prefer the lower premium that unstable insurers can offer. However, removing traders from the market may have unintended consequences. Although such a policy can reduce the number of unstable insurers in the market, it can also have the perverse effect of making the otherwise stable insurers riskier. This is because removing buyers from the market creates more competition among sellers, which drives down prices and in turn increases the counterparty risk of stable insurers. When there are social costs of counterparty risk, removing traders can be inefficient as a result of this effect.

In the first extension of the model, we consider the consequences of a central counterparty. A CCP acts as the buyer to every seller and the seller to every buyer. Participants in this market contribute to a fund designed to shelter each other from counterparty risk. Given that the counterparty risk to which an insured party is exposed is now that of the entire pool of insurers (through co-insurance), good insurers lose their comparative advantage. We consider the case in which there are a large number of insured parties and insurers. Given a CCP arrangement, counterparty risk is effectively pooled so that non-failing participants can absorb the losses of the failed ones. Therefore, insuring with a good insurer has little effect on the exposure of the insured to counterparty risk. Consequently, the insured party will contract with the bad insurer to obtain a better premium. In equilibrium, good insurers are pushed out of the market. This occurs because each individual insured party does not internalize the amount that their contract adds to the pooled counterparty risk, yielding an outcome similar in spirit to the classic problem of the commons. However, contrary to the problem of the commons, in this case central organization is the cause of and not the cure for this outcome. The solution to this problem is simple: the CCP should, to the extent possible, condition an insurer's contribution to the risk pool on their quality.

In the final extension of the model, we analyze asymmetric information on insurer quality. When stable insurers dominate the market under full information, asymmetric information is characterized by higher counterparty risk for two reasons: first, unstable counterparties pool with stable counterparties thereby lowering the average quality of counterparties and second, competition due to pooling drives down the premium thereby increasing the counterparty risk at the stable insurer. When unstable insurers dominate under full information, they may choose to pool with stable insurers and charge a higher premium. In this case counterparty risk decreases since stable insurers can now exist in the market driving up the average quality of counterparties. Alternatively, they may choose to reveal themselves when the pooling price is too low. In this case there is no change in counterparty risk from the full information case.

The features of CDS that we model are motivated by a number of recent studies and proposed policies. Arora et al. (2009) and Morkoette et al. (2012) provide evidence that counterparty risk exists in the CDS market, and the latter shows that it is priced.[3] Fitch (2009, 2010), Norden and Radoeva (2013) and Oehmke and Zawadowski (2013) provide strong evidence of the existence of pure trading motives for CDS, which is a key feature that we explore. Furthermore, we consider the consequence of removing players who use CDS for trading purposes since this closely mirrors policy that has been considered, and in some cases implemented, in recent years. Consider Germany's recent ban on the practice of buying CDS without owning the underlying risk, and China's intent on creating a CDS market with this same restriction.[4] We consider the imposition of a CCP in our model in light of the Dodd-Frank bill in the U.S., which mandates that a sizeable proportion of CDS trades go through clearinghouses. Europe appears also to be moving in a similar direction with the European Market Infrastructure Regulation (EMIR). Finally, asymmetric information in the CDS market of the type that we consider can arise because the financial institution selling the contract will likely know more about its own risk than the purchasing institution. This is particularly relevant in light of the opacity in the CDS market which became evident by the sudden and repeated downgrading of large CDS insurers such as Ambac, MBIA and AIG. Acharya and Bison (2013) provide a more in depth discussion and formal treatment of this opacity.

The paper proceeds as follows: Sections 2 and 3 describe the baseline model and equilibria. In Section 4, we allow insured parties to differ based on their motivation to purchase CDS. Section

[3] There is a debate currently taking place on the extent to which counterparty risk is priced in the CDS market. In the model to be outlined, we simply assume that the risk can be fully priced. As Arora et al. (2009) note, if the price mechanism is not be fully utilized to price counterparty risk, then collateral must be offsetting the risk. All of the results of our paper will follow through if we embed collateral into the model either in lieu of, or in addition to the price mechanism. The key for our paper is that unstable insurers can be penalized, if not through price, then through the posting of costly collateral.

[4] http://www.bloomberg.com/news/2010-09-13/china-plans-to-introduce-credit-default-swaps-by-year-end-official-says.html, http://dealbook.nytimes.com/2010/05/19/germany-bans-naked-shorts-and-c-d-s-s/

5 provides two extensions of the model: 5.1 explores the consequences of a central counterparty, while 5.2 considers asymmetric information over insurer quality. Section 6 provides a discussion on the robustness of two of our assumptions, and Section 7 concludes. Non-trivial proofs can be found in the Appendix.

Literature Review

This paper contributes to the literature on counterparty risk, credit default swaps and insurance. Thompson (2010) considers a case with endogenous counterparty risk in financial insurance. It is shown that an insurer has a moral hazard problem and may not invest in the best interest of the insured party. Furthermore, it is shown that truthful revelation of insured type can be attained because revelation affects the investment decision of the insurer, and consequently, the counterparty risk to which the insured is exposed. In contrast, we explicitly model multiple insurers and so can analyze the composition of insurers in the market. In another related paper, Acharya and Bisin (2013) show that due to the opacity of over-the-counter markets (where many CDS trade), counterparty risk can occur because insurers may take positions which increase their likelihood of default. In contrast, we model a situation in which insurers have varying degrees of stability and show that, regardless of whether CDS markets are opaque, unstable insurance can be a feature of the equilibrium. Neither Thompson (2010) nor Acharya and Bisin (2013) analyze the motivation to purchase CDS, or the commons problem that arises with a CCP, as is done in this paper.

In insurance economics, there is a small literature on insurer default which focuses mainly on contract size (see among others, Doherty and Schlesinger, 1990; Cummins and Mahul, 2003). For example, Cummins and Mahul (2003) determine the optimal indemnity in the case where the insurer and insured party have different beliefs about the probability that the insurer will fail. This analysis does not apply to our setting because of the inability of insured parties in the CDS market to separate on ex-ante contract size due to the non-exclusivity of contracts (i.e., a seller cannot preclude a buyer from purchasing insurance elsewhere). Therefore, separation à la Rothschild and Stiglitz (1976) cannot generally be achieved. This issue is detailed in Stephens and Thompson (2012).

2 Model: CDS as insurance

The purchaser, whom we refer to as a bank, owns a risky asset which it wishes to insure. We refer to this asset as a loan. We will not model anything unique to a bank, however as banks are the largest purchasers of these types of contracts, we use this terminology for ease of exposition. The providers of CDS are simply referred to as insurers. To give a preview of the model, we first outline the timing as summarized in Figure 1 below. There are three periods, in which we assume there is no discounting. At $t = 0$, an insurance contract is written by an insurer on the risky loan. The insurer is endowed with a portfolio (a random variable) and has an investment choice to make with the premium: how much to put in a liquid asset and how much to put in an illiquid asset. At $t = 1$, the bank submits a claim if needed and the insurer receives its portfolio draw. In addition, the part of the premium invested in the liquid asset is available. If there is a claim, the insurer fulfils it when solvent, otherwise it fails and returns nothing to the bank. At $t = 2$, if the insurer is solvent, the payoff to the premium invested in the illiquid asset is received.

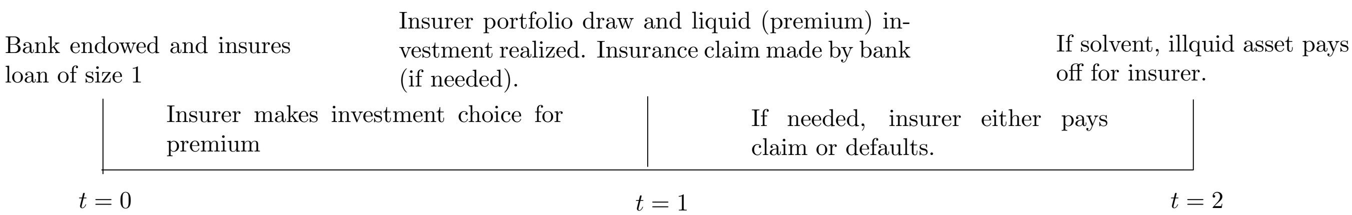

Figure 1: Timing of the Model

2.1 Banks

The fundamental characteristic of a bank is the desire to reduce risk. As in Thompson (2010), if the bank incurs a loss and has not insured this risk, it suffers the cost $Z \geq 0$.[5] If the bank has a loss for which it is insured, but the insurer cannot pay, we assume for simplicity that it also suffers the cost Z. This cost could represent a regulatory penalty for exceeding some risk level, or an endogenous reaction to a shock to the bank's portfolio; however, we will not model this here. It is this cost that makes the bank averse to holding risk.

The bank's loan yields return R_B with probability p, otherwise it defaults with probability $1-p$ and returns nothing. The size of the loan is normalized to 1 and we assume that the bank must insure this amount.[6] Therefore, if a claim is fulfilled, the bank will receive 1 and if it is not fulfilled, the bank is penalized Z. Denoting the premium (price) as P, and the probability that the insurer is solvent as q, the bank's expected payoff is

$$\pi_b = pR_B + (1-p)q - (1-p)(1-q)Z - P. \tag{1}$$

The first term represents the payoff to the loan when it yields a positive return. The second term represents the payoff when the loan returns nothing, but the insurer pays 1 through the claim. The

[5] See robustness Section 6.2 for a more formal discussion of Z.

[6] This is done for simplicity. One could imagine that the bank chooses different amounts of insurance depending on the insurer type (note that different insurer types will be discussed in the next subsection). An endogenous contract size will not affect the qualitative results to be presented. Note also that we are ignoring a potential moral hazard, wherein the bank may lose the incentive to monitor its loan when it completely insures. We assume full insurance for simplicity, but all the results of the paper would go through if we assumed that the bank insured only a fraction of its loan. That fraction could then be set such that moral hazard is eliminated. Even in the presence of a moral hazard problem, the only difference in our model would be that either p would decrease, R_B would decrease, or both. The insurer would simply alter its beliefs about the expected cost of a claim and the results of the model would follow through. For a more formal treatment of this moral hazard problem, see Bolton and Oehmke (2010), Parlour and Winton (2009) or Thompson (2007).

third term represents the penalty imposed on the bank when it submits a claim and the insurer is insolvent. The final term is the price paid for the contract. Note that in the event of a claim, the insurer fails with probability $1-q$ and for simplicity, pays nothing to the bank. As we will show below, $1-q$ is endogenous and is a function of P and the insurers investment choice (both of which are solved for in equilibrium).

2.2 Insurers

We allow for the possibility of insurer insolvency, and importantly, allow it to be heterogeneous across insurers. This represents our first departure from the literature. We model this heterogeneity by considering two insurance providers, one relatively stable and the other unstable, referred to simply as "(G)ood" and "(B)ad" insurers. It is worth pointing out that the title of *bad* insurer should be interpreted only to mean that this type has more counterparty risk. As we will soon show, the lower premium that bad insurers can charge may make them an attractive party with whom to contract.

To model the insurers, we extend the model of Thompson (2010) to include two insurer types and present a simplified environment without convex liquidation costs as is present in Thompson (2010). The good (bad) insurer is endowed with a portfolio of assets that produces a payoff represented by a draw from a distribution $F_G(\theta)$ ($F_B(\theta)$) at $t=1$.[7] Both distributions are assumed to have the same support, so that $\theta \in [\underline{\theta}, \overline{\theta}]$. Below we will place restrictions on these two distributions to ensure that a claim can be paid in at least some states of the world, and that an insurer cannot trivially eliminate all counterparty risk. If an insurer does not sell a contract, it is assumed to fail when its portfolio draw is between $[\underline{\theta}, 0]$. To facilitate the analysis, in such a case we assume that the insurer experiences a bankruptcy cost Γ.[8] The expected payoff of insurer $j \in \{G, B\}$ without the contract is given by:

$$\pi_j^{noinsure} = \int_{\underline{\theta}}^{0} (\theta - \Gamma) dF_j(\theta) + \int_{0}^{\overline{\theta}} \theta dF_j(\theta) \tag{2}$$

The first term represents the state of the world where the insurer fails. In that case the payoff is the (negative) draw plus the bankruptcy cost. The second term represents the state of the world where the insurer is solvent, in which case the payoff is simply the portfolio draw. When an insurer provides protection to the bank, it receives a premium up-front and makes an investment decision. The insurer may put money into a risk-free liquid asset at a rate of return 1 from $t=0$ to $t=1$. It may also put money into a more profitable, but illiquid risk-free asset which has a rate of return $r>1$, but is received only at $t=2$. For simplicity we assume that the illiquid asset has no pledgable

[7] A potentially interesting difference between CDS and pure insurance contracts is that the former insures more systematic risk while the latter, idiosyncratic risk. Systematic risk could be included in the model by allowing the bank loan to be correlated with $F_j(\cdot)$. The key requirement in the context of our model is that there is some counterparty risk and that the good insurer is more likely to pay when a claim is made.

[8] A bankruptcy cost helps illuminate the investment decision. Varying this parameter will be more intuitive than varying the return distributions when we consider how the insurers invest.

value at $t = 1$.[9] Insurer j chooses β_j, the percentage of the premium to put in the liquid asset. Each insurer knows the probability of a claim $1 - p$, and charges P_j for protection. We write the payoff function for each insurer when insuring a loan (recall that the loan size is normalized to 1).

$$\begin{aligned}\pi_j^{insure} = \quad & p\left[\int_{\underline{\theta}}^{-\beta_j P_j}(\theta - \Gamma + \beta_j P_j)dF_j(\theta) + \int_{-\beta_j P_j}^{\overline{\theta}}(\theta + (\beta_j + (1-\beta_j)r)P_G))\,dF_j(\theta)\right] \\ & + (1-p)\left[\int_{\underline{\theta}}^{(1-\beta_j P_j)}(\theta - \Gamma + \beta_j P_j)dF_j(\theta) + \int_{(1-\beta_j P_j)}^{\overline{\theta}}(\theta - 1 + (\beta_j + (1-\beta_j)r)P_G))\,dF_j(\theta)\right]\end{aligned} \tag{3}$$

The first line represents the state of the world where there is no claim (which occurs with probability p). The first integral represents the case in which the draw is sufficiently low so that the insurer fails. In this case, they receive the (negative) draw θ and are subject to the bankruptcy cost. For simplicity, we have assumed that they receive the proceeds from the liquid asset; however, they receive none of the illiquid asset since it cannot be liquidated early at $t = 1$ when the insurer fails. Notice the difference between default in this case, and default without the insurance contract in (2). Whereas without a contract, the insurer failed when the portfolio draw was below zero, now it fails when the draw is less than $-\beta_j P_j$. This is because the investment in the liquid asset will reduce the probability of insurer default by allowing some negative portfolio draws to be offset by the payoff from the liquid asset. The second integral represents the case in which the insurer is solvent. In this case, it receives both the portfolio draw θ as well as the return from both its liquid and illiquid investment. The second line represents the state of the world where there is a claim (which occurs with probability $1-p$). The first integral is the case in which the insurer is insolvent and so cannot pay the claim. Notice now that this occurs whenever the portfolio draw is less than $1-\beta_j P_j$. When a claim is submitted, the insurer can use its investment in the liquid asset; however, it needs to obtain a sufficiently high portfolio draw to pay the full amount (the contract/claim size of 1) or else it fails. In this case, they do not bear the cost of -1 since they do not pay anything to the bank and thus they receive the same payoff as had there been no claim. The second integral is similar to the case in which a claim was not submitted except now the insurer is less likely to be solvent, and the payoff they receive includes -1 because of the payment to the bank.

Note that because the premium is always paid upfront in our model, we only analyze one-sided counterparty risk. In reality, CDS contracts typically involve both an upfront payment and an ongoing premium, so that there a risk to the insurer that the bank is unable to pay the ongoing portion. The counterparty risk in which we study is presumably more interesting. The reason for this is that the bank is the party transferring risk to the insurer so that the amount that the bank can lose (the claim) will far exceed the ongoing premium. Modeling this two-sided counterparty risk would require a multi-stage game with a more complicated contract, however all that is required for our qualitative results to remain is that the bad insurer is still able to charge less than the good

[9] This implies that if insurer fails at $t = 1$, it receives nothing from the illiquid asset. This can be relaxed to allow it to be liquidated or borrowed against at $t = 1$, provided that it is done at a sufficient discount.

insurer.

To facilitate the analysis, we define the distributions F_G and F_B as simply as possible. We wish to analyze the case where the bad insurer's portfolio is riskier than the good insurer's, so we let the bad insurer have a payoff distribution with 'fat tails' whereas the good insurer has a more even distribution. For simplicity, let F_G be the uniform distribution, while F_B is an extreme case of 'fat tails': a two-point distribution that takes the value $\overline{\theta}$ with probability q_B and $\underline{\theta}$ with probability $1-q_B$. To highlight the most interesting case, let $\overline{\theta} > 1$ and $\underline{\theta} < -1$. Since $\overline{\theta} > 1$, even in the case where the bad insurer invests entirely in the illiquid asset, it will always be able to pay a claim given a draw of $\overline{\theta}$. Since $\underline{\theta} < -1$ the bad insurer cannot use the contract to avoid default in the event of the low draw, even if all of the premium was put in the liquid asset. In other words, the low draw is sufficiently low, so that bad insurer always defaults in that state. The extreme form of F_B simplifies the analysis but is not necessary for the qualitative results in the paper. A more general distribution is discussed in Robustness Section 6.1. To ensure that the bad insurer is in fact riskier, we assume throughout that $F_G(1) \leq 1 - q_B$. Henceforth, instead of using total profit with the insurance contract π_j^{insure}, we will use the incremental payoff to the insurers from the insurance contract, $\pi_j = \pi_j^{insure} - \pi_j^{noinsure}$. This transformation has no effect on the optimization problem we will solve below (since $\pi_j^{noinsure}$ is independent of β_j), and is required to find the equilibrium zero profit price.

The following lemma characterizes the optimal investment choices of the good and bad insurers, β_G^* and β_B^*.

Lemma 1
The good insurer optimal investment is: $\beta_G^* = \min\left\{1, \max\left\{0, \frac{\overline{\theta} - r(\overline{\theta}-(1-p)) + P_G r + \Gamma - (1-p) - \underline{\theta}}{2P_G r}\right\}\right\}$
The bad insurer optimal investment is: $\begin{cases} \beta_B^* = 0 & \text{if } r \geq \frac{1}{q_B} \\ \beta_B^* = 1 & \text{if } r < \frac{1}{q_B} \end{cases}$

Proof. *See Appendix.*

Notice that the bad insurers investment choice is independent of Γ. This is because the risk of failure with the contract is the same as the risk of failure without the contract. Therefore, they put the entire premium in the illiquid asset when the return of this asset, r, is sufficiently high. The good insurer's optimal investment is more complicated. It is straightforward to see that the higher Γ, the more of the premium that is put in the liquid asset. The reason for this is that investing in the illiquid asset causes the probability of failure (and so the probability of incurring the bankruptcy cost) to go up. Investing more in the liquid asset can reduce this cost. The min and max operators are required in the optimal solution to ensure that $\beta_G^* \in [0, 1]$. We wish to explore the most interesting case in which $\beta_G^* > 0$, so that the good insurer invests at least a part of its premium in the liquid asset. To this end, we now establish that there exists a bankruptcy cost $\hat{\Gamma}$ for which $\beta_G^* > 0$ whenever the bankruptcy cost exceeds $\hat{\Gamma}$.

Lemma 2 *There exists a* $\hat{\Gamma}$ *such that for any* $\Gamma > \hat{\Gamma}$, $\beta_G^* > 0$.

Proof. *See Appendix.*

The higher is the bankruptcy cost, the more the good insurer has to gain from reducing the probability of failure. Since only the liquid asset can be used to reduce the probability of failure, it follows that with a sufficient bankruptcy cost there will be at least some investment in the liquid asset. We make the following assumption to ensure that this is the case.

Assumption 1 $\Gamma > \hat{\Gamma}$.

3 Equilibrium

In equilibrium, the bank chooses an insurer with whom to contract. We assume for simplicity that the bank cannot split its contract over the two insurers.[10] In addition to the bank choosing the insurer with whom to contract, it can also choose whether to participate in the market at all. The participation constraints for the bank represents a choice between insurer j and no insurance. Note that for the remainder of the paper, we leave our mathematical expressions in general form and suppress the arguments of P and q as closed formed solutions are not possible. Participation requires:

$$P_j^0 \leq q_j(1-p)(1+Z), \tag{4}$$

where P_j^0 denotes the zero-profit price for insurer j. Expression (4) implies that when the insurer is charging the lowest price that it can, the benefit of insurance to the bank (the right hand side) exceeds the cost (the left hand side). We now define the equilibrium of the model with participation.

Definition 1 *An equilibrium with participation is a set of prices and an insurer choice such that:*

1. *Banks choose insurer type to maximize payoff.*

2. *Prices are determined through Bertrand competition among insurers.*

Modifying expression (1), we give the bank's payoff function when contracting with insurer type j.

$$\pi_b(j) = pR_B + (1-p)q_j - (1-p)(1-q_j)Z - P_j \tag{5}$$

We assume for simplicity that a bank which is indifferent between contracting with a good and bad insurer opts for the former. Therefore, the good (bad) insurer will dominate the market when $\pi_b(G) \geq \pi_b(B)$ ($\pi_b(B) > \pi_b(G)$). The following proposition summarizes the two equilibria with trade in the model. Existence of a parameter space which satisfies the required expressions is established in the proofs.

[10] One can imagine a transaction cost which induces this behavior, or a straightforward restriction on the parameter space will also accomplish this. Allowing the bank to split its contract would only complicate the analysis and would not change our qualitative results. We refer the interested reader to Stephens and Thompson (2012), where this assumption is explicitly relaxed.

Proposition 1 *When trade occurs there exist two types of equilibria, each of which is unique for a given parameter range.*

1. *The bank purchases insurance from the good insurer when*

$$(1-p)(1+Z)(q_G - q_B) \geq P_G^0 - P_B^0, \tag{6}$$

where the equilibrium premium for this case is

$$P_G^* = \min\{P_B^0 + (1-p)(1+Z)(q_G - q_B), (1-p)(1+Z)q_G\}. \tag{7}$$

2. *The bank purchases insurance from the bad insurer when*

$$(1-p)(1+Z)(q_G - q_B) < P_G^0 - P_B^0 \tag{8}$$

where the equilibrium premium for this case is

$$P_B^* = \min\{P_G^0 - (1-p)(1+Z)(q_G - q_B) - \epsilon, (1-p)(1+Z)q_B\}, \tag{9}$$

for ϵ small.

Proof. *See Appendix.*

The good insurer will dominate the market when the benefit of reduced counterparty risk, the left hand side of expression (6), more than compensates for the additional premia that the bank must pay, the right hand side of expression (6). To derive this case, we set both the good and bad insurer's premia to that which earns zero profit and determine when the good insurer is preferred by the bank. The equilibrium price is then determined in one of two ways. When the bad insurer offers sufficient competition (such that the bank's participation constraint is satisfied at P_B^0), the good insurer will raise its premium to $P_G^* = (1-p)(1+Z)(q_G - q_B) + P_B^0$, i.e., the point at which the bank is indifferent between the good and the bad insurer.[11] Alternatively, if the bank's participation constraint at the bad insurer is not satisfied at P_B^0, i.e., the bank would prefer not to contract rather than contract with the bad insurer, the equilibrium price is simply given by that which makes the bank just indifferent between purchasing the contract at the good insurer or not participating at all, $P_G^* = (1-p)(1+Z)q_G$. Again, this equilibrium can yield positive profit for the good insurer. As one might expect, condition (6) can hold when Z is sufficiently large, as more risk-averse buyers will favor the additional protection that the good insurer offers.

Conversely, the bad insurer dominates the market when the premium discount it can offer exceeds the cost of the additional counterparty risk that it poses to the bank. When the good

[11]It is important to note that the equilibrium price is such that the good insurer can earn positive profits. Bertrand competition forces the good insurer to cut its price just enough so that the bank prefers it over the bad insurer, however there is no reason for the good insurer to cut its price further to the point of zero profits.

insurer offers sufficient competition (such that the bank's participation constraint is satisfied at P_G^0), the bad insurer sets a premium that earns weakly positive profits and is just low enough to force the good insurer out of the market $P_B^* = P_G^0 - (1-p)(1+Z)(q_G - q_B) - \epsilon$. Alternatively, when the good insurer is sufficiently unattractive (so that the bank's participation constraint is not satisfied at P_G^0), the bad insurer simply sets the price that makes the bank indifferent to purchasing the contract $P_B^* = (1-p)(1+Z)q_B$. Condition (8) can be satisfied when r and Γ become large. To see this, note that as r becomes large, the zero profit premium of the bad insurer becomes small, and since the zero profit premium of the good insurer is increasing in Γ, we show in the proof that there exists a Z such that (8) is satisfied. This represents what we view as a key difference between traditional insurance and CDS: in traditional insurance, buyers of protection are risk-averse. Thus, in the context of our model, this implies that buyers are more likely to have a Z such that the good insurance equilibrium prevails. In CDS on the other hand, it is possible that the buyer may have a low Z, e.g., $Z = 0$ corresponding to risk neutrality. Therefore the bad insurance equilibrium may prevail in this case. We study this issue further in Section 4.

We now consider the effect that changes in price have on counterparty risk, which is a key feature of the model.

Lemma 3 *Counterparty risk at the good insurer is decreasing in price.*

Proof. *See Appendix.*

One could imagine a decrease in price arising from increased competition for contracts. In a richer model of both sides of the market, this is quite intuitive: when you increase the supply of insurance, price falls.[12] Thus competition, or more generally any market feature which affects prices will affect counterparty risk. Note that for simplicity, we have structured the model such that counterparty risk at the bad insurer is exogenous, thus only the good insurer is relevant for Lemma 3. The same result would hold at the bad insurer if we altered F_B to allow for endogenous counterparty risk.

3.1 Efficiency

In this paper, we analyze market forces which impact counterparty risk. Our environment is designed to study incentives and is not suited for a full general equilibrium welfare analysis, however in this section we consider the impact of a simple externality which we introduce on efficiency in risk transfer.[13] To do this, consider the case in which failure of the contract causes instability in the bank (the cost of which to the bank is, as before, Z). Such costs may not be borne solely by the

[12] This was done in an earlier version of the paper.

[13] There is actually an inefficiency inherent in the existing model without introducing any new features. There exists a moral hazard problem on the good insurer investment choice. This moral hazard problem arises because the insurer does not consider the cost of counterparty risk to the bank when it makes its investment decision. This problem has been studied in Thompson (2010) where it is shown that insurers invest too much in illiquid assets from the point of view of the bank. A similar problem has also been studied in Acharya and Bisin (2013). We opt to pursue an inefficiency better suited to the analysis in our paper.

bank, however. The literature has highlighted a variety of externalities associated with the failure of an individual financial institution. For example, bank instability has been linked to contagion through interconnectedness and fire sales, it has been shown to hamper economic development as well as reduce future GNP.[14] This suggests that there are potential costs to the counterparty risk that we study that the bank may not internalize. To capture these potential externalities, we assume that at the same time the cost Z is borne by the bank, there exists a social cost of sZ that the bank does not consider when making its choice, where $s \geq 0$.[15] Thus, we should think of s as incorporating the cost of bank instability to all others outside that bank, which, as discussed above, includes the cost to other financial institutions through linkages, and to society in general through reduced economic development and growth. To maintain tractability however, we assume the cost sZ is exogenous. To analyze efficiency, we consider a utilitarian social planner who, in addition to the payoffs of the bank and insurers, internalizes the sZ cost.[16] Thus, the planner maximizes welfare according to

$$W = \sum_{k=b,B,G} \pi_k - (1-p)(1-q)sZ, \tag{10}$$

where k indexes the bank, bad insurer and good insurer. The following proposition summarizes the efficiency of insurer choice inherent in the equilibrium from Proposition 1.

Proposition 2 *When* (6) *is satisfied, the insurer choice is efficient. When* (8) *is satisfied, the insurer choice is inefficient for s sufficiently large.*

Proof. *See Appendix.*

This result says that the choice of the bad insurer is inefficient whenever social costs of counterparty risk are sufficiently high (the condition for s can be found in the proof as equation (24)). To understand this result, consider the equilibrium in Proposition 1 in which the bank contracts with the bad insurer, i.e., (8) is satisfied. For the bank, the choice represents a tradeoff between the benefit of the lower premium at the bad insurer and the increased cost of counterparty risk, Z. The social planner on the other hand internalizes the full cost of counterparty risk, both Z and sZ. Therefore, it may force the bank to contract with the good insurer when the bank would have contracted with the bad in equilibrium. In this case the equilibrium insurer choice is inefficient. It is straightforward to see that an s can be chosen such that the planner will always force the bank to contract with the good insurer. In contrast, when (6) is satisfied so that the good insurer is chosen in equilibrium, a social planner would never force the choice of the bad insurer since that would incur additional cost due to s.

[14] See among others, Cifuentes et al. (2005), Colwell and Davis (1992), De Nicolò et al. (2012), and Kupiec and Ramirez (2013).

[15] It is relatively straightforward to allow for a social cost of default in the state in which a claim is not made; however we choose to do it this way to correspond with the cost of counterparty risk for the bank.

[16] The use of a utilitarian objective is made without loss of generality, the key insights will not change if we allow for a more general set of welfare weights.

4 Incentives to Insure

A fundamental difference between a standard insurance market and that for CDS are the incentives for purchasing these contracts. To our knowledge, there has not been a paper which analyzes the impact of buyer's incentives on the market for CDS. We can address this issue by a simple addition to the model, since the incentive to purchase insurance can be captured solely by Z. One would expect that those who purchase CDS for risk management purposes will view counterparty risk differently than those who purchase it for trading purposes (i.e., traders will internalize the cost of counterparty risk less). Further, although it does not directly follow that those who purchase CDS for risk management will own the underlying loan, it is reasonable to expect that those who do not own the underlying loan are more likely to purchase CDS for trading purposes than for risk management. This has interesting consequences for policy since ownership of the underlying loan can be easily observed.

Modifying the base model from Section 2, we create a market for CDS as simply as possible. In addition to the good and bad insurers, we let there be two types of banks who differ on Z, which we denote Z_L and Z_H (to be defined below). We will refer to the Z_L bank as a trader and the Z_H bank as a hedger. Note that we do not use the term "speculator" to describe the Z_L bank, as this can be a controversial label. Keynes (1930) and Hicks (1946) present the first treatment of such issues. The key variable that separates hedgers from "speculators" in these papers is risk aversion. Subsequent literature has also emphasized informational asymmetries and beliefs in addition to risk aversion when modeling speculation (see among others, Hirshleifer, 1975, 1977; Spiegel and Subrahmanyam, 1992; and more recently, Goldstein, Li and Yang, 2013). Given the interpretation of Z in our model as risk aversion (see Robustness Section 6.2), we focus on the case in which traders differ solely in this parameter. For simplicity, in our model we interpret those that own the underlying risk as hedgers, and those that do not as traders. Given that a full market microstructure model is beyond the scope of this paper, we simply note that the results below will obtain when there are divergent beliefs, provided that the hedger has a higher willingness to pay due to risk aversion.

We now investigate the effect that the market composition of hedgers and traders has on counterparty risk. The definition of equilibrium remains that of Definition 1. The following lemma determines the value of Z for which a bank is indifferent between insuring with a good or bad insurer.

Lemma 4 *Define $\hat{Z}$ as the level of Z for which the bank is indifferent between insuring with the good or bad insurer at the zero profit premium for each insurer. Thus, with P_G^0 and P_B^0, a bank with $Z < \hat{Z}$ will prefer to contract with the bad insurer and a bank for which $Z > \hat{Z}$ will prefer the good insurer. The expression for $\hat{Z}$ is given by $\hat{Z} = \frac{P_G^0 - P_B^0}{(1-p)(q_G - q_B)} - 1$.*

Proof. *See Appendix.*

We interpret $\hat{Z}$ by considering the two relevant components of a contract from the perspective of a bank: counterparty risk and premium. A bank trades off a higher premium against increased counterparty risk in its choice of insurer. A bank for which $Z < \hat{Z}$ is less averse to counterparty risk and so insures with the bad insurer, as it is able to offer a lower premium than the good insurer. A bank for which $Z > \hat{Z}$ is sufficiently averse to counterparty risk to compensate for the increase in premium at the good insurer. For what follows, we assume that $Z_L = 0 < \hat{Z}$ and $Z_H > \hat{Z}$ so that traders are risk-neutral and hedgers are risk-averse.

Using Lemma 4, we can explore the difference between a market for CDS and that for traditional insurance in a relatively simple way. When modeling a standard insurance market, it is customary to assume that the insured party has exposure to the underlying risk (i.e., has an insurable interest). An insured party is typically modeled as being risk averse and so willing to pay a risk premium when purchasing insurance. As discussed above, in the market for CDS, some buyers purchase protection purely for trading purposes. As the number of traders (i.e., $Z_L = 0$ types) increases, so does the relative amount of insurance sold by bad insurers. Since it is reasonable to assume that the CDS market has more traders than a traditional insurance market, it follows that the market for CDS will tend to have lower quality sellers.

The existence of traders and bad insurers implies that CDS markets are generally characterized by higher counterparty risk. Although this is a "mechanical" consequence of our framework, it adds a new element to the policy debate on the CDS market. Ideally, a policy maker whose mandate is to reduce counterparty risk could simply remove bad insurers; however, the quality of the counterparty is often not observable to the bank (we will explore this case in Section 5.2), so it may not be observable to a regulator.[17] An alternative policy is to remove the Z_L banks from the market, similar to the recent proposals described in the introduction which disallow CDS to be purchased by those who do not own the underlying loan. Although it is possible that those who own the loan could purchase CDS for trading purposes, it is more likely that this policy will reduce the number of buyers for which Z_L more than it would for buyers with $Z = Z_H$. Therefore, we analyze the case in which the $Z_L = 0$ bank can be eliminated.

Proposition 3 *Removing traders from the market eliminates the bad insurer, but increases counterparty risk at the good insurer*

Proof. *See Appendix.*

When the $Z_L = 0$ bank is removed, the bad insurer no longer has a bank with whom to contract. Instead, it competes for the Z_H bank, thereby driving down the premium that the good insurer can charge. This is very intuitive and amounts to a decrease in demand causing the price to decrease. In equilibrium, the good insurer still obtains the contract with the Z_H bank, but since the premium decreases, the counterparty risk increases as per Lemma 3. Therefore, this result should be viewed as a cautionary note. Removing traders may reduce the number of unstable counterparties in the market, but there is another, opposing effect in that stable counterparties may become less stable.[18]

[17] For a discussion on this issue see Pirrong (2009).

[18] In a previous version we defined the average market counterparty risk and showed that there exists conditions under which it increases with the policy.

We now consider the policy with respect to the social planner's problem.

4.1 Efficiency

To consider a policy which removes traders, we extend the definition of social welfare described in (10) to include two buyer types, so that welfare is

$$W = \sum_{k=L,H,B,G} \pi_k - (1-p)(1-q_G)sZ_H, \tag{11}$$

where L and H represent low and high type buyers. Note that since $Z_L = 0$, there is no private or social cost to default for this type (although one could be included by allowing Z_L to be strictly positive without changing the results). The goal of this brief section is simply to further stress that the increased counterparty risk shown in Proposition 3 can have negative effects on the market. In particular, the following corollary follows easily from that proposition.

Corollary 1 *Removing traders from the market increases the expected social cost of counterparty risk and decreases welfare for s sufficiently large.*

When the price at the good insurer decreases following the removal of the trader, the increase in counterparty risk causes the expected social cost to increase. In terms of welfare, the policy represents a transfer from the good insurer to the Z_H bank. Although it is individually optimal for the bank to pay a lower price for protection, the social planner internalizes both the cost to the insurer (i.e., the lower premium) and social cost, which the bank does not internalize.[19] These costs can dominate the benefit to the bank in the social planning problem causing welfare to decrease as a consequence of the policy. When s is sufficiently large, such a decrease is assured. In a general model, the welfare implications of such a policy will always be ambiguous. The above result is intended to highlight the fact that the counterparty risk channel from Proposition 3 can cause a decrease in welfare, relative to a case in which this channel is not present.

5 Extensions

5.1 Central Counterparties

CDS contracts have been slowly migrating from over-the-counter markets to more formal central counterparty (CCP) arrangements. In the wake of the credit crisis that began in 2007, law makers around the world have been tabling regulations to legally mandate this migration.[20] In the absence of a central counterparty, contracts take one of two forms. First, contracts can be negotiated through a dealer. In these types of transactions, a buyer purchases protection from a counterparty

[19]The other cost of the policy comes from the bad insurer whom no longer receives a positive payoff. Note that prior to the removal of the $Z_L = 0$ bank type, the equilibrium price is such that the trader makes zero profit from the contract. This is in contrast to the case in which the insurer makes zero profit, i.e., the case where there is only one bank and two insurers as is true after the policy is implemented.

[20]See Bliss and Steigerwald (2007) an for an in-depth discussion on CCPs.

located by a dealer or from the dealer itself. Second, trading may be done without a dealer, where a buyer approaches a seller directly. In a CCP arrangement the contract is initially between a buyer and seller as per usual, however after the terms have been agreed upon, the CCP simultaneously buys the contract from the seller and sells to the buyer. In other words, all transactions flow through a central counterparty which acts as the buyer to every seller and the seller to every buyer. In this arrangement, participants provide capital and post margins (collateral) that the CCP can use to cover default losses.[21] Therefore, a CCP is an attempt to mutualize default risk across participants.

In our model, bad insurers are forced to set a lower premium because they pose a greater risk of default. Importantly, the CCP forces a single premium on the market because traders view counterparty risk as being only that of a CCP. As discussed above, a CCP ordinarily requires capital (for a default fund) and collateral in case of contract non-performance. Typically, CCPs demand capital/collateral according to the quality of the asset being insured, but less so based on the quality of the counterparties (Pirrong, 2009). To best illuminate our result, we will consider the polar case in which a CCP cannot differentially penalize bad insurers.[22]

A comprehensive analysis of a CCP arrangement is beyond the scope of this paper, however our framework can be used to address an issue that has been largely ignored in the debate thus far. Extending the model from Section 4, we let there be a measure N banks, who each contract with an insurer. Within the banks, assume that there are a measure $N_G \leq N$ for which $Z = Z_H$ and a measure $N_B \leq N$ for which $Z = Z_L = 0$ where $N_G + N_B = N$. As in Section 4, we assume for simplicity that each active insurer contracts with its own bank and that there is an excess of insurers of both types so that competition drives premia down to that which earns zero profit. When there is no CCP, Lemma 4 implies that in equilibrium Z_H (Z_L) banks insure with good (bad) insurers, so that the measure of good (bad) insurers in the market is N_G (N_B).

We now analyze the imposition of a CCP on this market. To pool risk, we assume that each insurer contributes a fixed amount c regardless of their quality, to a default pool. Given a measure N active insurers, the total size of the pool is then $D = Nc$, where $c \in (0, 1]$. The CCP will pay out claims as long as it is solvent, but fails if the number of insurers which have defaulted on claims is too high. For simplicity, we normalize collateral to zero and assume that the CCP is unable to raise additional funds after insurers have defaulted. Therefore, the CCP itself will default when D insurers cannot pay their claim (recall that contracts are of size 1). The default risk of the CCP, denoted $1 - q_{ccp}$, can be characterized as

$$1 - q_{ccp} = \begin{cases} 0 & \text{if} \quad (1 - q_B)\, N_B + (1 - q_G)\, N_G \leq D \\ 1 & \text{if} \quad (1 - q_B)\, N_B + (1 - q_G)\, N_G > D. \end{cases}$$

[21] In theory, the CCP can require participants to make additional payments if needed to cover losses, however there is no consensus as to how well this mechanism would work in practice.

[22] In practice, CCPs can and sometimes do try to enforce higher capital charges (and higher collateral) to riskier counterparties. The relevance of this assumption is discussed below, but we note that the results of this section will survive provided that the CCP does not perfectly condition on counterparty quality (which in practice appears to be the case).

Since there is a measure of insurers, the counterparty risk of the CCP is deterministic.[23] We now turn to the equilibrium in this extended model. The definition of equilibrium remains the same as that in Definition 1. Each bank must choose its insurer type to maximize profit and prices are determined through Bertrand competition. As discussed, given that there is always excess insurers, Bertrand competition implies that prices are always those which earn zero profit.

When faced with claims that cannot be fulfilled by insurers, the CCP either has sufficient funds and never defaults, or always defaults. Given that each insurer must pay c to the CCP to participate, it follows that the cost must be borne by the banks in the form of a higher premium. When a bank makes a choice of an insurer with whom to contract, it is too small to change the counterparty risk of the CCP. Therefore, the advantage of the good insurer (i.e., lower counterparty risk) is absent so that a bank's choice will be driven solely by the premium. This leads to the main result of this section.

Proposition 4 *In the presence of a CCP, bad insurers will push good insurers out of the market when $P_B^0 < P_G^0$.*

Consider one Z_H bank switching from a good to a bad insurer. Given that each bank is of measure zero, default risk of the CCP remains the same. Therefore, the Z_H bank will switch when $P_G^0 > P_B^0$. It follows that every Z_H bank will unilaterally switch to the bad insurer, so that in equilibrium $N_G = 0$. This result is similar in spirit to the classic problem of the commons in that the imposition of a CCP results in banks which do not internalize the effect of their decisions on counterparty risk.[24] The resulting equilibrium level of counterparty risk, which we define as $1 - q_{ccp}^*$, then follows since $N_B = N$ and $N_G = 0$:

$$1 - q_{ccp}^* = \begin{cases} 0 & \text{if} \quad c \geq 1 - q_B \\ 1 & \text{if} \quad c < 1 - q_B. \end{cases}$$

It is worthwhile exploring the robustness of Proposition 4. The assumption of an infinite number of insurers implies that risk pooling by the CCP is perfect. Consider the case in which $c = 1 - q_B$. It is easy to show that no bank would wish to contract outside of the CCP with a bad insurer, and provided that the price at the bad insurer is sufficiently low, with a good insurer. This is because the amount by which the bank's premium increases under the CCP (c), is low relative to the decrease in counterparty risk from the CCP arrangement. This will not necessarily be the case if, for example, the bad insurers are more exposed to aggregate risk than the good insurers. With risk of this form, the CCP will need to charge more if it is to avoid default. However, Proposition 4 still holds because individual banks do not consider the impact of their decisions on the CCP so that the market will still be serviced solely by bad insurers. In this case, if aggregate risk was

[23] In a previous version of the paper, we had a finite number of banks which allowed for a probabilistic counterparty risk of the CCP. Our deterministic version provides a substantially less complicated analysis while yielding the same intuition into the problem.

[24] This result can also be interpreted as an example of the Lucas critique, in that policy-makers considering the imposition of a CCP must consider the reaction of market participants to the policy.

sufficiently high, it is straightforward to derive a case in which banks would wish to leave the CCP and contract bilaterally.

To eliminate the outcome described in Proposition 4, the CCPs could charge bad insurers proportionately more to participate. Since bad insurers would have to pass this charge to banks in the form of a higher premium, it would inhibit the ability to undercut the good insurer. In our model, this would be possible if the CCP were able to use premia as a signal of underlying counterparty risk. In reality, CCPs may not be able to perfectly deduce the quality of insurers in the market. Pirrong (2009) reports that this may not be possible since, even in a dealer market, other dealers can struggle to quantify counterparty risk and so it is unlikely that a CCP could perfectly quantify the risk. Nonetheless, our analysis suggests that CCPs should condition capital requirements and collateral on the quality of the counterparty to the extent possible.

A note regarding mark-to-market is in order. Currently, much of risk management within CCPs is done by setting collateral charges based on the quality of the underlying asset by marking to market daily. This mechanism is meant to alleviate counterparty risk, however it does so indirectly. With mark-to-market of this type, the CDS seller would have to post additional collateral if the quality of the underlying asset deteriorates. If, for example, the quality of the insurer falls at the same time as the underlying asset, the increase in collateral will help mitigate the counterparty risk. However, it could be that the underlying asset becomes safer at the same time as the insurer becomes riskier. In this case, the decrease in collateral exacerbates the counterparty risk. Therefore, it is clear that mark-to-market on the underlying asset does not invalidate Proposition 4.

Finally, we wish to make clear that this result obtains in a natural extension of our model and highlights a very specific point relevant to the debate over CCPs. There are many factors that should be considered in determining whether such an arrangement would be beneficial to the market. For example, a richer characterization of the benefits of diversification through co-insurance relative to the endogenously lower quality individual insurance that we consider. Further, there are other possible benefits such as netting that CCPs can provide.[25] A full welfare analysis of CCPs is an interesting direction for future research.

5.2 Unknown Insurer

We extend the baseline model from Section 3 to consider the consequences of asymmetric information regarding the quality of the insurance provider. With asymmetric information, the bad insurer can camouflage itself and offer a contract with the same premium as the good insurer. Although our results will hold under a wide range of off-equilibrium-path beliefs, to avoid undue complication we place simple and plausible restrictions on the bank's beliefs.

Assumption 2 *Upon observing a premium P less than P_G^0, the bank's beliefs are that the insurer is bad. For prices $P \geq P_G^0$, the banks beliefs are that either insurer is equally likely.*

[25]For an analysis on the efficiency of netting, see Duffie and Zhu (2011).

In the case in which the bank does not update its beliefs, an insurer with whom to contract is chosen randomly. We now extend Definition 1 to allow for asymmetric information.

Definition 2 *An equilibrium with participation is a set of prices and an insurer choice such that:*

1. *Banks choose insurer type to maximize payoff.*
2. *Prices are determined through Bertrand competition among insurers.*
3. *Bank beliefs are updated according to Assumption 2*

The following Proposition characterizes the impact of asymmetric information on counterparty risk.

Proposition 5 *Relative to the full information case:*

- *When $(1-p)(1+Z)(q_G - q_B) \geq P_G^0 - P_B^0$, the presence of asymmetric information unambiguously increases expected counterparty risk.*
- *When $(1-p)(1+Z)(q_G - q_B) < P_G^0 - P_B^0$, the presence of asymmetric information weakly decreases expected counterparty risk.*

Proof. *See Appendix.*

Note that the characterization and existence of the equilibria is established in the proof. The intuition behind this result is as follows. When the insurer type is known and the good insurer dominates ($(1-p)(1+Z)(q_G - q_B) \geq P_G^0 - P_B^0$), the good insurer charges the highest premium such that the bank still prefers to contract with it rather than the bad insurer (who offers P_B^0). When the insurer type is unknown, the good insurer is forced to cut its premium or else give up the entire market to the bad insurer who can undercut it without revealing itself. In equilibrium, bank beliefs and insurer competition imply that the premium charged by both insurers is P_G^0. Since the good insurer (weakly) reduces its premium, it follows from Lemma 3 that the good insurer becomes less stable. Furthermore, bad insurers now (potentially) participate in the market, so that expected counterparty risk unambiguously increases.

When the bad insurer dominates under full information ($(1-p)(1+Z)(q_G - q_B) < P_G^0 - P_B^0$), there are two equilibria that can arise when insurer type is unknown. First, the bad insurer may choose to reveal itself by setting $P_B^* < P_G^0$ and dominate the market. The bad insurer does this to obtain the insurance contract with certainty, rather than charging P_G^0 and allowing the good insurer to stay in the market, thereby reducing the chances it will obtain the contract. Since the counterparty risk of the bad insurer is independent of the premium, if it chooses to reveal itself and take the market, expected counterparty risk remains unchanged. Conversely, when the bad insurer prefers the higher premium P_G^0 over obtaining the contract with certainty, the presence of the good insurer causes market counterparty risk to fall since the expected counterparty risk to which the bank is exposed decreases.

6 Robustness

6.1 Insurer Portfolio Distribution and Investment Choice

We model heterogeneity between the two insurer types in a simple way - by assuming that F_B is a two-point distribution while F_G is uniform. Alternatively, the good (bad) insurer could receive a draw from a more general distribution. As in the paper, let the proportion that the good (bad) insurer invests in the liquid asset be given by β_G (β_B), with the remainder invested in the illiquid asset. Using the usual notation for premia, counterparty risk can now be defined in a similar way as in Section 2, $1-q_G = F_G(1-\beta_G^* P_G)$ and $1-q_B = F_B(1-\beta_B^* P_B)$, where the asterisk represents the optimal portfolio choice. We then need only impose the appropriate restrictions on the distribution functions to ensure $q_G(\beta_G^*) > q_B(\beta_B^*)$ so that our bad insurer has higher counterparty risk and that $q_G'(P) < 0$, namely, that price and counterparty risk move in opposite directions. It is not even required that $\beta_G^* > \beta_B^* > 0$ provided again that $q_G(\beta_G^*) > q_B(\beta_B^*)$. Note that when $\beta_B^* > 0$ and $q_B'(P) < 0$, the qualitative results remain unchanged but the analysis becomes more tedious since we must consider endogenous counterparty risk for the bad insurer as well.

In addition to relaxing the distributional assumptions we can also consider the investment choice itself. Although the liquid versus illiquid investment choice presented in this paper yields crisp results, we could instead have a risky versus riskless choice, or a more complicated portfolio problem involving the choice of assets of varying risk/liquidity, again provided that $q_G'(P) < 0$ and $q_G > q_B$.

6.2 Z and Risk Aversion

It is worthwhile to contrast the parameter Z in our model with standard utility assumptions made in most insurance papers. Typically, a non-linear utility function is used for the insured party that puts different weights/utility value on high and low outcomes. A standard risk averse utility function will put relatively more negative weight on the bad outcomes (e.g., an 'accident') versus the high outcome (e.g., no 'accident'). As such, insurance is purchased to protect the risk averse individual that can cost more than the expected loss from the accident. In our model, we use the simplest formulation possible that captures these motives. In particular, we put a weight Z on the bad outcome (i.e., the loan fails). As such, the utility in the good state (i.e., the loan does not fail) is simply equal to the monetary payoff. In our model we let $Z \in [0, \infty)$. To understand this range, we consider the condition under which a bank (with probability of default 1-p) is indifferent between purchasing and not purchasing insurance, i.e., its participation constraint.

$$\begin{aligned} pR_B + (1-p)q - (1-p)(1-q)Z - P &= pR_B - (1-p)Z \\ \Rightarrow P &= (1-p)q(1+Z) \end{aligned} \tag{12}$$

Therefore, when $Z = 0$, $P = (1-p)q$, which is the actuarially fair premium. In other words, the bank will pay at most the expected value of the coverage. This corresponds to the usual insurance result with a risk-neutral agent.[26] When $Z > 0$, the bank is willing to pay greater than the expected value in return for coverage. This represents the risk premium that an insurance provider can extract due to the risk aversion of the insured party.

7 Conclusion

In this paper we analyze features of credit default swaps which can exacerbate counterparty risk. We characterize when unstable insurers can exist in these markets and show that downward pressure on prices can increase counterparty risk. Furthermore, we show that when some buyers of CDS use the instrument purely for trading purposes (and potentially have no insurable interest), the market will be characterized by more unstable insurers; however, removing these traders can cause otherwise stable insurers to become less stable. We extend the model to consider the consequence of a central counterparty and show that in such an arrangement, stable insurers can be driven out of the market due to their inability to compete on premia. Finally, we consider the case when counterparty risk of the insurer is unknown to the insured party. We show that asymmetric information can cause otherwise stable insurers to become less stable due to pooling, and find that counterparty risk will increase when stable insurers dominate under full information, but can decrease when unstable insurers dominate under full information.

Appendix

Proof of Lemma 1

Using our distributional assumptions and simplifying, we write the investment problem of the good insurer.

$$\max_{\beta_G \in [0,1]} \quad \frac{\overline{\theta} P_G(\beta_G+(1-\beta_G)r)+P_G^2\beta_G(1-\beta_G)r-(1-p)P_G(1-\beta_G)r+\Gamma P_G\beta_G-(1-p)\Gamma-(1-p)(\overline{\theta}-(1-P_G\beta_G))-\underline{\theta}\beta_G P_G}{\overline{\theta}-\underline{\theta}} \tag{13}$$

Taking the FOC and solving for β_G^* yields:

$$\beta_G^* = \frac{\overline{\theta} - r(\overline{\theta} - (1-p)) + P_G r + \Gamma - (1-p) - \underline{\theta}}{2P_G r} \tag{14}$$

Implementing the restriction that $\beta_G^* \in [0, 1]$ yields the condition from the Lemma. Similarly, we write the investment problem of the bad insurer as

$$\max_{\beta_B \in [0,1]} \quad \beta_B P_B + q_B P_B(1-\beta_B)r - q_B(1-p). \tag{15}$$

Taking the FOC and observing the restrictions we obtain: $\begin{cases} \beta_B^* = 0 & \text{if } r \geq \frac{1}{q_B} \\ \beta_B^* = 1 & \text{if } r < \frac{1}{q_B} \end{cases}$ as desired.

[26] One can draw a parallel to the asset pricing literature in that when $Z = 0$, P represents the risk neutral probability of the state in which the loan defaults and the insurer is solvent.

■

Proof of Lemma 2

From Lemma 1, β_G^* is non-negative when $\overline{\theta} - r(\overline{\theta} - (1-p)) + P_G r + \Gamma - (1-p) - \underline{\theta} \geq 0$. Thus the following exogenous condition ensures $\beta_G^* > 0$.

$$\hat{\Gamma} = (r-1)[\overline{\theta} - (1-p)] + \underline{\theta} > (r-1)[\overline{\theta} - (1-p)] - P_G r + \underline{\theta} \tag{16}$$

■

Proof of Proposition 1

Case 1: Expression (6) is needed to ensure that the good insurer dominates the bad. Re-arranging yields:

$$1 + Z \geq \frac{P_G^0 - P_B^0}{(1-p)(q_G - q_B)}, \tag{17}$$

which holds for Z sufficiently large since the right hand side is finite. The participation constraint is $P_G^* \leq q_G(1-p)(1+Z)$, which again holds for Z sufficiently large. Bertrand competition between the good and bad insurer then determines the equilibrium price. If the participation constraint is slack, then the good insurer's optimal premium P_G^*, is that which satisfies (6) with equality (where P_G^0 is replaced by P_G^*), as given in the statement of the proposition. If the participation constraint binds, then $P_G^* = q_G(1-p)(1+Z)$. Note that the case in which the bank prefers to pay a higher premium than P_G^* is uninteresting and can be ruled out simply by allowing the bank a technology to retain the funds. Alternatively, we can show that a parameter range exists under which this equilibrium exists and under which the payoff for the bank is decreasing in price.

Case 2: To prove existence, consider the case in which r becomes arbitrarily large (thus $\Gamma > \hat{\Gamma}$ becomes arbitrarily large as per Lemma 2). Thus, $P_B^0 \to 0$ and $P_G^0 \to 1-p$ and so expression (8) becomes:

$$1 + Z < \frac{1}{q_G - q_B}. \tag{18}$$

The participation constraint of the bank at the bad insurer is $P_B^* \leq q_B(1-p)(1+Z)$, which is implied by $P_G^0 \leq q_G(1-p)(1+Z)$; since the bank participates at the good insurer when the good insurer charges the zero profit price, then it must participate with the bad insurer since the bad insurer dominates in this equilibrium. Since $P_G^0 \to 1-p$, we re-write the participation condition as:

$$1 + Z \geq \frac{1}{q_G}. \tag{19}$$

To see that (18) and (19) can be simultaneously satisfied, simply consider the case in which $1+Z = \frac{1}{q_G}$. The equilibrium premium P_B^* is then the maximum for which (8) is still satisfied (where P_B^0

is replaced by P_B^*) or is given by the participation constraint, whichever is lower. Note that since counterparty risk at the bad insurer is independent of price, a bank would never wish to pay any more than P_B^*.

■

Proof of Lemma 3

Using the uniform assumption, write the probability of insurer default conditional on a claim as:

$$1 - q_G = \frac{1 - \beta_G P_G - \underline{\theta}}{\overline{\theta} - \underline{\theta}} \tag{20}$$

Differentiating with respect to P_G yields:

$$\frac{d(1 - q_G)}{dP_G} = -\frac{\frac{d\beta}{dP_G} P_G + \beta}{\overline{\theta} - \underline{\theta}} \tag{21}$$

We need to show that equation (21) is negative for any $\beta_G \in (0, 1]$. When $\beta_G = 1$, $\frac{d\beta_G}{dP_G} = 0$ so that (21) is clearly negative. When $\beta_G \in (0, 1)$ we use (14) to obtain:

$$\frac{d\beta_G}{dP_G} = \frac{-\beta_G + \frac{1}{2}}{P_G} \tag{22}$$

Plugging (22) into (21) yields

$$\frac{d(1 - q_G)}{dP_G} = -\frac{1}{2(\overline{\theta} - \underline{\theta})} < 0. \tag{23}$$

■

Proof of Proposition 2

Case 1 - (6) is satisfied: In equilibrium, the bank chooses the good insurer (with price $P_G^* \geq P_G^0$) over the bad insurer (with price $P_B^* = P_B^0$). If the social planner forces the bank to choose the bad insurer at price $P_B^* = P_B^0$, then the insurers are worse off since positive profit is not possible, the bank is worse off since the bad insurer is individually sub-optimal, and the probability of incurring the social cost increases since counterparty risk increases. Thus welfare (10) unambiguously decreases. Case 2 - (8) is satisfied: In equilibrium, the bank chooses the bad insurer (with price $P_B^* \geq P_B^0$) over the good insurer (with price $P_G^* = P_G^0$). Define $\pi_b^B(P_B^*)$ as the expected payoff of the contract to the bank in equilibrium and $\pi_b^G(P_G^0)$ as the expected profit of the contract to the bank for the bank if it is forced to insure with the good insurer. Further, define $\pi_B(P_B^*)$ as the expected payoff to the bad insurer from the contract in equilibrium. Note that the expected payoff of the good insurer if the social planner forces the contract to be written with it is zero since it charges its zero profit price. It follows that forcing the bank insure with the good insurer is welfare improving when:

$$s > \frac{\pi_B(P_B^*) + \pi_b^B(P_B^*) - \pi_b^G(P_G^0)}{Z(1-p)(q_G - q_B)}. \tag{24}$$

Since payoffs are finite and $q_G - q_B \neq 0$, there exists a finite s under which (24) is satisfied.

■

Proof of Lemma 4

The bank's payoff from insuring with the good and bad insurers are given as follows.

$$\pi_b(G) = pR_B + (1-p)q_G - (1-p)(1-q_G)Z - P_G \tag{25}$$

$$\pi_b(B) = pR_B + (1-p)q_B - (1-p)(1-q_B)Z - P_B \tag{26}$$

We now define $\hat{Z}$ as that which equates these expressions.

$$\hat{Z} = \frac{P_G - P_B}{(1-p)(q_G - q_B)} - 1 \tag{27}$$

Inserting the zero profit premia yields the expression characterized in Lemma 4.

■

Proof of Proposition 3

We first consider the equilibrium prior to the removal of the Z_L bank. The Z_H bank is preferred by both insurers, as this type is willing to pay a higher premium for insurance, yet poses no additional risk. Thus, by Lemma 4, we restrict our attention to the case when the good insurer contracts with Z_H and the bad insurer with Z_L. Given this, a unique set of equilibrium premia are determined by the following set of participation and incentive constraints.

$$q_B(1-p)(1+Z_L) \geq P_B \tag{PCL}$$

$$q_G(1-p)(1+Z_H) \geq P_G \tag{PCH}$$

$$P_G - P_B \geq (1-p)(1+Z_L)(q_G - q_B) \tag{ICL}$$

$$P_G - P_B \leq (1-p)(1+Z_H)(q_G - q_B) \tag{ICH}$$

The inequality PCL (PCH) ensures that the Z_L (Z_H) bank will purchase insurance from the bad (good) insurer, rather than go without. Inequality ICL (ICH) ensures that the Z_L (Z_H) bank contracts with the bad (good) insurer rather than the competitor. Expanding ICH, we have

$$P_G \leq q_G(1-p)(1+Z_H) - [q_B(1-p)(1+Z_H) - P_B]. \tag{28}$$

The second term on the right hand side is negative since

$$q_B(1-p)(1+Z_H) - P_B > q_B(1-p)(1+Z_L) - P_B \geq 0, \tag{29}$$

where the second inequality follows from PCL. Thus, (28) shows that PCH is redundant and can be ignored. Furthermore, in equilibrium, ICH must be satisfied with equality, otherwise the good insurer could increase the premium and still attract the Z_H bank. This implies

$$P_G - P_B = (1-p)(1+Z_H)(q_G - q_B) > (1-p)(1+Z_L)(q_G - q_B), \tag{30}$$

so that ICL can also be ignored. Finally, in equilibrium the bad insurer will increase its premium until the Z_L bank is just indifferent to purchasing the contract or not so that PCL is satisfied with equality. To summarize, the equilibrium premia in this situation are $P_B^* = q_B(1-p)(1+Z_L)$ and $P_G^* = (1-p)(q_G - q_B)(1+Z_H) + P_B^*$.

We now consider the equilibrium when the Z_L bank is removed. From Lemma 4, this will drive the bad insurer out of the market but changes the equilibrium premium of the good insurer. With only one bank to compete over, the bad insurer cuts its premium to P_B^0. Thus, the good insurer sets $P_G^{**} = (1-p)(q_G - q_B)(1+Z_H) + P_B^0$. Since $P_G^{**} \leq P_G^*$, Lemma 3 implies that the good insurer will become (weakly) less stable.

■

Proof of Proposition 5

Case 1: $(1-p)(1+Z)(q_G - q_B) \geq P_G^0 - P_B^0$. Bertrand competition and the bank beliefs imply that the only permissible equilibrium is pooling wherein the price is P_G^0 and each insurer obtains a contract with probability $\frac{1}{2}$. Assume to the contrary that $P > P_G^0$. Then, either the good or bad insurer could cut its price by some small ϵ and obtain the contract with certainty, thus $P > P_G^0$ cannot be an equilibrium. It is straightforward to see that for any $P < P_G^0$, the bad insurer is revealed. Since $(1-p)(1+Z)(q_G - q_B) \geq P_G^0 - P_B^0$ implies that the bad insurer does not exist in the market under full information, it cannot be that $P < P_G^0$. Thus, $P = P_G^0$. With full information over insurer type and all insurance provided by the good insurer, the premium is $P_G^* \geq P_G^0$ as characterized in Proposition 1. Lemma 3 implies that counterparty risk increases when the premium decreases. Thus, comparing the full information equilibrium to the asymmetric information equilibrium, expected counterparty risk increases since both the bad insurer and the good insurer at the lower equilibrium premium have higher counterparty risk than the good insurer under full information.

Case 2: $(1-p)(1+Z)(q_G - q_B) < P_G^0 - P_B^0$. There are two permissible equilibria: a pooling and a revealing equilibrium. As above, the pooling equilibrium is given by the price P_G^0 and each insurer obtains the contract with equal probability. The payoff of the bad insurer in this case is given by: $\frac{1}{2}\pi_B(P_G^0)$. If the bad insurer sets any price below P_G^0 then it reveals itself and the equilibrium price is given by Proposition 1. The payoff in that case is: $\pi_B(P_B^*)$. Therefore, when $\pi_B(P_B^*) < \frac{1}{2}\pi_B(P_G^0)$ only pooling can exist, when $\pi_B(P_B^*) > \frac{1}{2}\pi_B(P_G^0)$, only revealing can exist, where the knife-edge case can be defined either way. To see that there exists a parameter range under which the pooling equilibrium can exist, let $(1-p)(1+Z)(q_G - q_B) = P_G^0 - P_B^0 - \epsilon$ for ϵ small. In this case, the bad insurer that reveals earns zero profit since the bank is indifferent to it and the good insurer when they both charge zero profit prices. Since $(1-p)(1+Z)(q_G - q_B) > 0$, it follows that $P_G^0 > P_B^0$ and thus the bad insurer will earn positive expected profit by pooling. In terms of counterparty risk, with full information only the bad insurer exists in the market. With asymmetric information, the possibility of contracting with the good insurer causes expected counterparty risk to fall.

To see that parameters exist such that the revealing equilibrium can exist, let r become arbitrarily large (thus $\Gamma > \hat{\Gamma}$ becomes arbitrarily large as per Lemma 2). In this case $P_B^0 \to 0$ and $P_G^0 \to 1-p$. The bank participation constraint with pooling implies that $P_G^0 \leq (1-p)(1+Z)q_j$, where in this case, q_j represents the average of q_B and q_G. Thus, if the bad insurer pools and changes $P_G^0 = 1-p$, the bank does not participate for any $Z \leq 1/q_j - 1$. When this holds, the bad insurer earns nothing. Therefore, in this case, revelation is optimal since $P_B^* \geq P_B^0$ and thus the bad insurer can earn positive profit. This equilibrium is identical to that under full information and so counterparty risk remains unchanged.

■

8 References

Acharya, Viral and Alberto Bisin, "Counterparty Risk Externality: Centralized Versus Over-the-Counter Markets," *Forthcoming, Journal of Economic Theory*, (2013).

Allen, Franklin and Douglas Gale, "Financial Contagion," *Journal of Political Economy*, 104 (2000), 1–33.

Arora, Navneet, Priank Gandhi, and Francis A. Longstaff, "Counterparty Credit Risk and The Credit Default Swap Market," UCLA Working Paper, (2009).

Bliss, Robert R. and Robert Steigerwald, "Derivatives Clearing and Settlement: a Comparison of Central Counterparties and Alternative Structures," *Economic Perspectives*, 4Q (2006), 22–29.

Bolton, Patrick and Martin Oehmke, "Credit Default Swaps and the Empty Creditor Problem," *Review of Financial Studies*, 24 (2011), 2617–2655.

Boyd, John H. and Gianni De Nicoló, "The Theory of Bank Risk Taking and Competition Revisited," *The Journal of Finance*, 60 (2005), 1329–1343.

Cifuentes, Rodrigo, Gianluigi Ferrucci, and Hyun Song Shin, "Liquidity Risk and Contagion," *Journal of the European Economic Association*, 3 (2005), 556–566.

Colwell, R. J. and E. P. Davis, "Output, Productivity and Externalities - the Case of Banking," Bank of England Working Paper, (1992).

Cummins, J. David and Olivier Mahul, "Optimal Insurance With Divergent Beliefs About Insurer Total Default Risk," *The Journal of Risk and Uncertainty*, 27 (2003), 121–138.

De Nicolò, Gianni, Giovanni Favara, and Lev Ratnovski, "Externalities and Macroprudential Policy," IMF Staff Discussion Note, (2012).

Doherty, Neil and Harris Schlesinger, "Rational Insurance Purchasing: Considerations of Contract Nonperformance," *Quarterly Journal of Economics*, 105 (1990), 243–253.

Duffie, Darrell and Haoxiang Zhu, "Does a Central Clearing Counterparty Reduce Counterparty Risk?," *Review of Asset Pricing Studies*, 1 (2011), 74–95.

Fitch Ratings, "Global Credit Derivatives Survey: Surprises, Challenges and the Future," www.fitchratings.com (2009).

Fitch Ratings, "Credit Derivatives Survey: Respondents Opine on Public Perceptions, Regulations, Sovereigns and More," www.fitchratings.com (2010).

Goldstein, Itay, Yan Li, and Liyan Yang, "Speculation and Hedging in Segmented Markets," *Forthcoming, Review of Financial Studies*, (2013).

Hicks, John R., “Value and Capital,” (Oxford: Clarendon Press, 1946).

Hirshleifer, J., “Speculation and Equilibrium: Information, Risk, and Markets,” *Quarterly Journal of Economics*, 89 (1975), 519–542.

Hirshleifer, J., “The Theory of Speculation Under Alternative Regimes of Markets,” *Journal of Finance*, 32 (1977), 975–999.

Keynes, John M., “A treatise on Money,” in II: The Applied Thoery of Money (London: Macmillan and Co., 1930).

Kupiec, Paul H. and Carlos D. Ramirez, “Bank failures and the cost of systemic risk: Evidence from 1900 to 1930,” *Journal of Financial Intermediation*, 22 (2013), 285–307.

Morkoetter, Stefal, Johanna Pleus, and Simone Westerfeld, “The impact of counterparty risk on credit default swap pricing dynamics,” *Journal of Credit Risk*, 8 (2012), 63–88.

Norden, Lars and Kristina Radoeva, “What do we know about speculation in the CDS market?,” *Erasmus University Working Paper*, (2013).

Oehmke, Martin and Adam Zawadowski, “The Anatomy of the CDS Market,” *Columbia University and Boston University Working Paper*, (2013).

Parlour, Christine and Andrew Winton, “Laying off Credit Risk: Loan Sales versus Credit Default Swaps,” *Journal of Financial Economics*, 107 (2013), 25–45.

Pirrong, Craig, “The Economics of Clearing in Derivatives Markets: Netting, Asymmetric Information, and the Sharing of Default Risks Through a Central Counterparty,” Available at SSRN: http://ssrn.com/abstract=1340660, (2009).

Rothschild, Michael and Joseph Stiglitz, “Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information,” *Quarterly Journal of Economics*, 125 (1976), 629–649.

Spiegel, Matthew and Avanidhar Subrahmanyam, “Informed Speculation and Hedging in a Noncompetitive Securities Market,” *Review of Financial Studies*, 5 (1992), 307–329.

Stephens, Eric and James R. Thompson, “Separation Without Mutual Exclusion in Financial Insurance,” University of Alberta and University of Waterloo Working Paper, (2012).

Stulz, René M., “Credit Default Swaps and the Credit Crisis,” Finance Working Paper No. 264, ECGI (2009).

Thompson, James R., “Credit Risk Transfer: To Sell or to Insure,” Queen’s University Working Paper No. 1131, (2007).

Thompson, James R., “Counterparty Risk in Financial Contracts: Should the Insured Worry About the Insurer?,” *Quarterly Journal of Economics*, 90 (2010), 1195–1252.

Vives, Xavier, “Competition and Stability in Banking,” IESE Business School Working Paper, (2010).