Concentrated Capital Losses and the Pricing of Corporate Credit Risk*

Emil N. Siriwardane†

Job Market Paper‡

February 12, 2015

Abstract

Using proprietary credit default swap (CDS) data from 2010 to 2014, I show that capital fluctuations for sellers of CDS protection are an important determinant of CDS spread movements. I first establish that markets are dominated by a handful of net protection sellers, with five sellers accounting for nearly half of all net selling. In turn, a reduction in their total capital increases CDS spreads. Capital fluctuations of the largest five sellers account for over 10 percent of the time-series variation in spread changes, a significant amount given that observable firm and macroeconomic factors account for less than 17 percent of variation during this time period. I then demonstrate that the concentration of sellers creates fragility — higher concentration results in more volatile risk premiums. I also employ a number of complementary approaches to address identification, such as using the 2011 Japanese tsunami as an exogenous shock to the risk bearing capacity of CDS traders. My findings are consistent with asset pricing models with limited investment capital, but also suggest that both the level and distribution of capital are crucial for accurately describing price dynamics.

*The views expressed in this paper are those of the author's and do not necessarily reflect the position of the Depository Trust & Clearing Corporation (DTCC), the Office of Financial Research (OFR), or the U.S. Treasury. DTCC data is confidential and this paper does not reveal any confidential information.

I am very grateful to my dissertation committee: Robert Engle (chair), Viral Acharya, Xavier Gabaix, and Stijn Van Nieuwerburgh. In addition, I have benefited from conversations with Tobias Adrian, David Backus, Daniel Barth, Nishani Bourmault, Matteo Crosignani, Itamar Dreschler, Greg Duffee, Darrell Duffie, Andrea Eisfeldt, Nikunj Kapadia, Matteo Maggiori, Patricia Mosser, Holger Mueller, Martin Oehmke, Thomas Philippon, Sriram Rajan, Alexi Savov, Philipp Schnabl, Kevin Sheppard, and Sumudu Watugala. I am also extremely indebted to Joe Bishop, Rob Capellini, Jin Ohm, and Valerie Wells for all of their help on this project. I thankfully acknowledge grant support from the Macro Financial Modeling Group (MFM) and the NYU Stern Center for Global Economy and Business.

†Affiliation: NYU Stern School of Business and the Office of Financial Research, U.S. Department of Treasury. Address: 44 West 4th St., Floor 9-Room 197K. New York, NY 10012. E-mail: esiriwar@stern.nyu.edu.

‡For the latest version, see http://people.stern.nyu.edu/esiriwar/pdf/Siriwardane_JMP.pdf

1 Introduction

The 2007-09 financial crisis highlighted the role that limited investment capital plays in determining equilibrium asset prices, a feature that has since been embedded in many popular asset pricing models.[1] A common element of these models is that capital cannot flow frictionlessly to investment opportunities, and in turn, risk premiums adjust with the available amount of risk bearing capital in the market. In this paper, I explore two ways in which capital fluctuations for traders in the credit default swap (CDS) market impact the pricing of credit risk.[2] First, I show that increases in the level of credit spreads are driven by decreases in the total amount of risk bearing capital for sellers of CDS protection. Second, I show that the volatility of default risk premiums increases as the concentration of CDS protection sellers rises. My findings suggest that both the *level* and *distribution* of investment capital in a market are critical for accurately describing asset price dynamics.

I measure capital fluctuations for each trader using changes in the mark-to-market value of all of its CDS positions.[3] In order to conduct my analysis, I use a proprietary dataset of over 600 million CDS positions written on 5700 underlying reference entities and 900 CDS indices. The proprietary data is provided by the Depository Trust & Clearing Corporation (DTCC), and contains detailed information on the portfolios of nearly 1700 counterparties. The DTCC supplies trade processing and registration services for all major dealers of CDS, meaning I effectively observe the entire U.S. market from 2010 to 2014. To the best of my knowledge, I am the first to map true economic exposures via CDS for every counterparty and reference entity in the U.S. based on individual transaction data.

Through this data, I discover a salient feature of CDS markets: they are dominated by a handful of net protection buyers and sellers, with sellers twice as concentrated as buyers. The top five sellers account for nearly half of all net selling, or in other words, 50 percent of net selling is in the hands of less than 0.1 percent of the total number of CDS traders. I refer to these large CDS players as mega-sellers, mega-buyers, or more generally, mega-players.

Because mega-sellers represent such a large share of the market, I begin by studying

[1] e.g. Froot and O'Connell (2008), Mitchell, Pedersen, and Pulvino (2007), He and Krishnamurthy (2013), Brunnermeier and Sannikov (2014), or Duffie and Strulovici (2012).

[2] In a CDS contract, the buyer of insurance pays a premium to a seller for protection against corporate default. The buyer and seller in the swap are called "counterparties." The insurance contract covers the default of an underlying firm, or "reference entity."

[3] Hence, I treat capital fluctuations for CDS trading desks as the relevant state variable for pricing, and I argue why this is a sensible assumption in Section 4. In Section 7.1, I provide more institutional details and some empirical evidence consistent with this interpretation.

how their capital impacts CDS pricing. I formally explore the relationship between capital fluctuations and credit spreads using a panel regression with log-CDS spread changes as the dependent variable. I address identification in this setting by testing whether losses from one portion of a mega-seller's CDS portfolio affect pricing for unrelated reference entities. As an example of my strategy, I test whether changes in the CDS spread of Ford Motor Company can be explained by mega-seller capital losses coming from positions taken outside of the auto-industry. I also control for a large number of reference entity characteristics and macroeconomic variables that may drive movements in CDS spreads.

I find that mega-seller capital changes have a substantial effect on CDS spreads. Capital fluctuations of the five largest sellers in the market account for nearly one-ninth of the variation in weekly CDS spread movements. To put this in perspective, observable firm-level and macroeconomic factors explain only one-sixth of spread variation over the same time period. Following with intuition, capital losses raise the effective risk aversion of sellers, thereby increasing the premium they require for bearing default risk. When the five largest sellers incur a one standard deviation capital loss, the level of CDS spreads rises by 2.8 percent per week. This elasticity is economically large, as the standard deviation of weekly spread movements is 6 percent for the average firm in my sample. Consistent with many theoretical models of limited investment capital, I also demonstrate a non-linear relationship between risk bearing capital and CDS pricing — capital losses impact spreads more than capital gains.

In isolation, these findings do not necessarily imply that concentration itself impacts premiums, as opposed to just the total level of capital in the market. The second objective of my paper is therefore to show why concentration, or the distribution of risk bearing capital, is indeed important for pricing. High concentration creates fragility because an idiosyncratic capital shock to an important seller can have a sizable impact on aggregate risk bearing capital. Consequently, the volatility of risk premiums increases as sellers become more concentrated.

As one way to establish a link between volatility and concentration, I estimate what I call the aggregate price of credit risk, denoted by Π_t, from a panel of CDS spreads. In simple terms, Π_t captures the average component of credit spreads that cannot be explained by fundamental default risk.[4] Next, I estimate a GARCH volatility model for log-changes in Π_t, and I refer to this estimated volatility series as σ_t^{π}. To quantify the concentration of sellers

[4] Gilchrist and Zakrajsek (2012) refer to their version of essentially the same quantity as the "excess bond premium."

in CDS, I compute a standard Herfindahl measure for natural sellers of CDS protection.

I then test how σ_t^π responds to changes in the concentration of sellers. Volatility rises substantially as sellers become more concentrated. A one standard deviation increase in seller concentration (via the seller Herfindahl) is associated with a 2 percent increase in the volatility of the price of credit risk. An increase of this magnitude is large given that risk prices are relatively smooth. Put differently, a one standard deviation increase in seller concentration results in a 1.75 standard deviation event for σ_t^π. These findings highlight why the distribution of risk bearing capital — not just the level — is important for asset price dynamics.[5]

For completeness, all of my empirical tests consider the effects of both buyer and seller capital fluctuations, though my exposition focuses primarily on sellers of CDS because I do not observe the joint bond-CDS portfolio for buyers. Thus, it is difficult to interpret how a buyer's mark-to-market CDS loss impacts his entire portfolio. On the other hand, it is more likely that large seller losses affect their overall wealth because directly hedging a sold CDS position requires costly shorting of the underlying bond. I discuss these issues further in Section 2. Furthermore, net sellers are much more concentrated than buyers, so their risk bearing capacity should have a greater influence on CDS pricing. As expected, I find little evidence of a relationship between credit spreads and buyer capital movements. In addition, buyer concentration does not appear to impact the volatility of the price of credit risk.

To bolster the interpretation of my main results, the latter portion of the paper further addresses identification issues in a few ways. First, I use the 2011 Japanese tsunami to study how an exogenous shock to mega-seller risk bearing capacity affected CDS spreads on U.S. firms. My proprietary data reveals that U.S. counterparties had large net exposure to Japanese firms prior to the tsunami, a necessary condition for mega-players to propagate the shock to U.S. firms. I find no evidence of mega-buyers transmitting the shock of the tsunami to U.S. reference entities; however, firms whose primary protection sellers were highly exposed to Japanese firms saw their CDS spreads rise 2.4 percent in the week after the tsunami, relative to reference entities whose main sellers had very low exposure to Japan. To emphasize the importance of concentrated positions, I make use of the fact that one seller — seller J — had a particularly outsized exposure to Japanese firms. I then compare U.S.

[5]In Appendix A, I show that capital losses for all natural sellers, as opposed to just mega-sellers, are followed by increases in the price of credit risk. I also offer more evidence that mega-seller capital fluctuations effectively represent the entire market. On the other hand, there is no response in risk prices to dollar losses at the "average" CDS seller. If the distribution of risk bearing capital was unimportant, this would not be the case.

firms based only on J's share of their selling (or buying). I find that U.S. firms where seller J had a larger share of selling also experienced larger spread increases after the tsunami struck.

For all of the preceding analyses, I used aggregated CDS spread quotes as the relevant dependent variable, but my last identification technique also finds evidence that seller capital losses impact pricing using executed transactions.[6] More specifically, I exploit variation in CDS premiums for a single buyer purchasing protection from multiple sellers on a fixed reference entity and date. By controlling for all unobservable characteristics of the buyer and the reference entity, I am able provide causal evidence that seller losses impact subsequent CDS spreads. My results indicate that sellers who have incurred larger capital losses charge relatively higher premiums for CDS protection.[7]

Related Literature

My paper adds to a rapidly growing literature that studies the behavior of asset prices when capital cannot flow frictionlessly to investment opportunities. The importance of capital market frictions is central to theoretical models of limits to arbitrage (Shleifer and Vishny (1997); Kyle and Xiong (2001)), slow moving capital (Duffie (2010); Acharya, Shin, and Yorulmazer (2013); Duffie and Strulovici (2012)), and financial intermediary-based asset pricing (Brunnermeier and Sannikov (2014); He and Krishnamurthy (2013); Adrian and Boyarchenko (2013)). On the empirical side, examples of previous studies on asset pricing with limited investment capital include Froot and O'Connell (2008), Gabaix, Krishnamurthy, and Vigneron (2007), Mitchell, Pedersen, and Pulvino (2007), Coval and Stafford (2007), Chen, Joslin, and Ni (2014b), Acharya, Schaefer, and Zhang (2014), and Adrian, Etula, and Muir (2014). My first major finding augments the aforementioned theoretical and empirical work by showing that fluctuations in seller capital play a sizable role in generating time variation in CDS spreads.[8] Relative to the existing empirical research in these fields, my

[6] When I refer to "CDS spreads" or "aggregated CDS spread quotes", I mean those coming from standard data vendors like Markit and Bloomberg, who report a composite CDS spread for each reference entity. The composite spread is a function of many dealer quotes.

[7] For a related application in the banking literature, see the "within-firm" estimator applied by Khwaja and Mian (2008) or Chodorow-Reich (2014).

[8] My findings are closely related to Acharya, Schaefer, and Zhang (2014), who find that CDS spreads rose for firms outside of the auto-industry following the downgrade of GM and Ford in 2005. Their explanation is that financial intermediaries were reluctant to bear additional credit risk, presumably because their existing inventories were exposed to GM and Ford. Chen, Joslin, and Ni (2014b) also find a very similar result in options markets, where the premium for out-of-the-money options (and other asset classes) increases as financial intermediaries reduce their desire to sell put options to institutional investors.

paper also provides more direct evidence that limited investment capital impact asset prices, a contribution that is made possible by the richness of my data on CDS positions.

Many asset pricing models with capital market frictions utilize a representative intermediary or investor who faces some impediment to investment. However, as my results show, considering the impact of multiple investors with heterogeneous market shares (or constraints) is helpful for understanding higher moments of risk premiums. Thus, incorporating concentration effects into theoretical models of limited investment capital may improve the empirical performance of these models.

My paper also adds to a large literature on how credit risk is priced in the economy. Examples of work in this area include Duffee (1998), Elton, Gruber, Agrawal, and Mann (2001), Collin-Dufresne, Goldstein, and Martin (2001). A common conclusion of these studies is variables that measure fundamental default risk, such as those from structural models of credit, explain a surprisingly low amount of credit spread variation. More recently, structural models of credit have introduced time varying risk premiums for corporate bonds. Some examples of this approach are Chen, Collin-Dufresne, and Goldstein (2009) or Chen, Cui, He, and Milbradt (2014a). These papers focus on firm exposure to macroeconomic or liquidity risks, whereas my results propose that changes in seller risk bearing capital partially account for movements in credit spreads. Because some large CDS sellers are themselves intermediaries, the general theme of my paper is also in line with previous studies that link financial-intermediary activity to corporate bond pricing (e.g. Green, Hollifield, and Schurhoff (2007), Newman and Rierson (2004)).

In terms of real effects, Gilchrist and Zakrajsek (2012) have a related paper demonstrating that shocks to the price of credit risk lead to significant declines in consumption, investment, and output. These authors conjecture that shocks are driven by the risk bearing capacity of financial intermediaries. I add to this conjecture by showing that capital shocks to mega-sellers, some of whom are financial intermediaries, can explain shocks to credit risk prices. More specifically, in Appendix A a simple variance decomposition indicates that one-fifth of unexpected changes in the price of credit risk can be attributed to shocks to the capital of natural CDS sellers. Consequently, my findings have important implications not just for asset pricing, but also for real activity.[9]

My work is also related to a recent theoretical paper by Atkeson, Eisfeldt, and Weill

[9]My findings also connect to recent research asking whether CDS markets are redundant to corporate bond markets. Oehmke and Zawadowski (2014) give a theoretical foundation for non-redundant CDS markets, and the evidence in my paper broadly supports their claim.

(2014a), whose model predicts that banks with large preexisting risk exposure, or those with low effective risk bearing capacity, charge a high spread for selling CDS in equilibrium. One way to map my results to their model is to define their traders with low risk bearing capacity as mega-sellers who have previously lost money in CDS. Moreover, these authors introduce exit decisions following an unexpected negative shock to CDS traders. When new counterparties are hard to find — as I show is true by virtue of the concentration in CDS — their model suggests there are cases where negative idiosyncratic shocks can lead to increases in prices. My finding about how concentration can induce fragility is broadly consistent with this intuition.

Lastly, a significant innovation in this paper is a comprehensive set of facts describing CDS markets in the United States. What separates the data in this study is the granularity of my information. Unlike previous work, I observe the full identities of counterparties, terms of trade, and (nearly) all outstanding CDS exposures going back to 2010. Less detailed forms of the DTCC data I employ have been used by Schachar (2012), Chen, Fleming, Jackson, Li, and Sarka (2011), Oehmke and Zawadowski (2013), Atkeson, Eisfeldt, and Weill (2013), and Duffie, Scheicher, and Vuillemey (2014). Because I am able to determine the ultimate net CDS buyers and sellers for each reference entity, my empirical work should also aid ongoing efforts to understand how derivatives markets may affect financial stability. Duffie, Li, and Lubke (2011) and references therein provide an overview of these issues. A lot of regulatory focus has been on improving transparency and reducing counterparty risk in over-the-counter (OTC) markets, particularly through central clearing of OTC derivatives. My paper highlights a different aspect of the market; namely, limited participation. The existence of mega-sellers and mega-buyers poses a different set of issues that may or may not be solved by recent regulatory efforts.

The remainder of the paper proceeds as follows. Section 2 gives a brief description of the data and methods used in this paper, with details found in a separate Data Appendix. Section 3 presents the main stylized facts that form the basis of the rest of the paper. Section 4 establishes my two main results about how the level and concentration of capital affect CDS premiums. In Appendix A, I conduct additional robustness tests that support the primary conclusions drawn in Section 4. Section 5 provides more empirical support of my main results by using the 2011 Japanese tsunami as an exogenous shock to the risk bearing capacity of CDS market participants. Section 6 presents an alternative way I mitigate identification concerns through actual CDS transactions. Finally, in Section 7, I discuss why capital losses may impact effective risk aversion and also provide some concluding remarks.

2 Data Description and Empirical Work

2.1 Data on CDS Transactions and Positions

As mentioned, the data I work with is provided by the DTCC, who provides trade processing services for every major dealer in credit default swaps. I have access to two complementary subsets of the DTCC's database: transactions and positions. Transactions represent flows in CDS, and positions represent stocks. For example, a new transaction between two counterparties also results in a new outstanding position between them. Future amendments or additional transactions on the same underlying reference entity simply add to their existing position. In practice, computing positions from transactions is quite complicated and is done using the DTCC's own internal algorithms.

For both transactions and positions, my data contains full information on the counterparties in the trade, the pricing terms, the swap type, the notional amount, the initiation date, and so forth. Within the DTCC's trade repository data, I am privy to any transaction or position that meets one of two conditions: (i) the underlying reference entity is a U.S. firm or (ii) at least one of the counterparties in the swap is registered in the U.S.. In addition, my data includes all North American index CDS transactions and positions (i.e. any where the reference entity is in the "CDX.NA." family). Taken together, my data effectively covers the entire CDS market for U.S. firms. The data begins in 2010 and is updated continuously on a weekly basis. I truncate my analysis in June 2014.

This dataset is enormous and contains over 40 million index positions and 600 million single name positions. To be as precise as possible, I have carefully documented each step of my data processing in a separate Data Appendix. When necessary, I also provide additional detail about the underlying data in the empirical analysis contained in the main text.

Index Swaps Index swaps constitute nearly half the gross notional of the entire CDS market, so accounting for exposures via index swaps is crucial for understanding true credit risk exposures. CDS index products contain a basket of single name swaps. For example, suppose I sell $100 of notional on an index swap that contains 100 different single names. If one of the names defaults, I have to then pay out $1 in notional to the buyer of the index swap. After I make my payment, there are 99 names in the index remaining. A future default of one of these names also results in a payout of $1. Writing $100 in protection via an index is therefore equivalent to writing 100 different single name swaps, each worth $1 in notional. To think about the amount of credit risk exposure to a single reference entity, I

am careful to consider exposures via single name swaps and index swaps. Full details of this procedure can be found in the Data Appendix, Section 1.3.

An important caveat is that this process does not consider the liquidity advantages of index swaps over single name swaps. However, because my principal concern is with the allocation and pricing of credit risk, I ignore any liquidity components that might otherwise complicate netting index positions against single name positions.

A Comment on Only U.S. Reference Entities For the remainder of the paper, there are applications where it is important to be aware of the limitations of my data. In particular, it is often necessary that I restrict myself to U.S. reference entities because I am not certain that I see the entire market for non-U.S. reference entities. Data Appendix Section 1.4 contains a detailed methodology of how I filter only U.S. reference entities. I take a conservative approach to creating this subset because many reference entities in the raw data do not have a listed country. In addition, I exclude reference entities that are written on mortgage backed securities, which I detail in Data Appendix Section 1.4.2.

A Comment on Bearing Credit Risk Through CDS I will often make statements about ultimate sellers of CDS protections, and in particular the largest net sellers, bearing corporate credit risk. However, because I do not observe the total portfolio of any counterparty, I am implicitly assuming that these sellers are not hedging their sold protection using other instruments. As noted earlier, the most direct hedge would be shorting a corporate bond, but evidence suggests it is very costly, especially at a large scale (Nashikkar and Pedersen (2007)). At a minimum, it seems safe to assume that net sellers are not fully hedging by shorting the underlying cash instrument. This assumption is most plausible for the largest net sellers because the quantity of bonds they would need to short to fully hedge are large.[10] An alternative way to hedge a sold position might be, for instance, to buy a deep out-of-the-money put option. While I also cannot rule this out, my empirical analysis of the price of credit risk provides additional evidence why the largest ultimate sellers of protection are actually bearing credit risk overall. If losses to the CDS portfolio do not change total wealth, theoretical models would not predict a relationship between CDS portfolio losses

[10]In addition, I have done my best to spot check that some of the largest sellers are not fully hedging their sold protection. When the seller is a dealer, I have compared their positions to their reported short selling of all corporate bonds in regulatory reports. When the seller is a hedge fund, I have verified that their advertised strategies are explicitly to take credit risk through selling CDS. I cannot report the specific results of these inquiries for legal reasons.

and credit risk pricing.

2.2 Market Wide Data on CDS Pricing

I obtain CDS spreads from the data vendor Markit, which reports a composite CDS spread term structure for a large number of reference entities. They compute their composite CDS spread using quotes from over 30 major market participants. Quotes are translated to a composite spread through Markit's own internal algorithms. For this reason, the Markit CDS spread reflects both quotes and realized transactions. This distinction will be important to keep in mind when I use transactional data to address identification.

I will additionally rely on information on the physical measure of default. As is relatively standard, I proxy for the physical likelihood of default using Moody's Expected Default Frequency (EDF) database. I describe each of these data in Section 1.2 of the Data Appendix.

2.3 Notation and Terminology

Throughout the paper, I will make use of the following terminology and notation. A reference entity is the underlying firm on which a credit default swap is written. For instance, if Hedge Fund ABC sells protection to Hedge Fund DEF for the default of firm X, then "X" is the reference entity in the transaction. The counterparties in a transaction are the buyer and the seller of insurance. Continuing with the previous example, the two counterparties would be Hedge Fund ABC and Hedge Fund DEF. $NS(c, r, t)$ denotes the *net* amount of protection sold by counterparty c on reference entity r on date t. For example, suppose that, as of date t, Counterparty c has sold 100 of notional insurance on Reference Entity r, but also has bought 25 of notional insurance on Reference Entity r. The net amount sold by c on r as of date t is then $100 - 25 = 75$. Thus, negative values of $NS(c, r, t)$ indicate the counterparty is a net buyer of protection. $\mathcal{C}_t$ is the set of all counterparties that have open positions in the CDS market as of date t. Lastly, $\mathcal{R}_t$ is the set of reference entities traded in the CDS market as of date t.

3 Facts About Credit Risk Sharing in CDS Markets

Before exploring how CDS trader capital affects the pricing of credit risk in Section 4, I first document the existence and market share of mega-players in CDS markets. I start by quantifying the size, or the amount of net risk transferred, by CDS. I then build a simple

measure that captures the aggregate concentration of net buyers and net sellers of protection. I also document additional facts regarding the network structure and risk flows in the Online Appendix, Section 1.

3.1 How Much Credit Risk Is Actually Transferred?

How big is the CDS market? Knowing the size of the overall CDS market, and of a particular reference entity, is important for at least two reasons. The first and most basic reason is that the true size of the entire CDS market has been difficult to pin down quantitatively due to data constraints. Indeed, the total gross notional of outstanding positions has been computed by a variety of sources. However, the gross notional does not provide much information about the amount of risk transferred.[11]

The second reason is to quantify the market share of buyers and sellers of each reference entity. In turn, this requires me knowing the net amount of risk outstanding for each reference entity r on date t, which I compute as follows:

$$NO(r,t) := \sum_{c \in \mathcal{C}_t} \max\left(NS(c,r,t), 0\right) \tag{1}$$

$NO(r,t)$ is analogous to the face value of debt outstanding in bond markets. By symmetry, it is also equivalent to summing the net amount bought across counterparties who are net buyers overall. The total amount of net outstanding in the market is then computed by summing $NO(r,t)$ over all reference entities:

$$NO(t) := \sum_{r \in \mathcal{R}_t} NO(r,t) \tag{2}$$

The left axis of Figure 1 plots my estimate of the total net outstanding $NO(t)$ through time.[12] In early 2010, the size of the U.S. CDS market was just under $2 trillion. It has

[11] To illustrate why, consider two transactions. In the first, Counterparty ABC sells $100 of protection on Reference Entity X to Counterparty DEF. In the second, Counterparty ABC buys $100 of protection on Reference Entity X from Counterparty DEF. The gross notional outstanding is $100 + 100 = 200$. But, the net exposure of Counterparty ABC to Counterparty DEF is zero. There is no actual credit risk transferred between the two counterparties. I analyze the gross notional and related metrics in the Online Appendix, Section 1.

[12] Section 1 of the Online Appendix provides a more refined look at the size of U.S. CDS markets. I compute my estimate by taking the average size between the: (i) the entire dataset and (ii) reference entities that I can definitely classify as a U.S. firm. The former is an upper bound. Conversely, the latter is a lower bound because I am conservative in classifying firms as U.S. based.

Figure 1: Net Notional Outstanding in CDS Markets

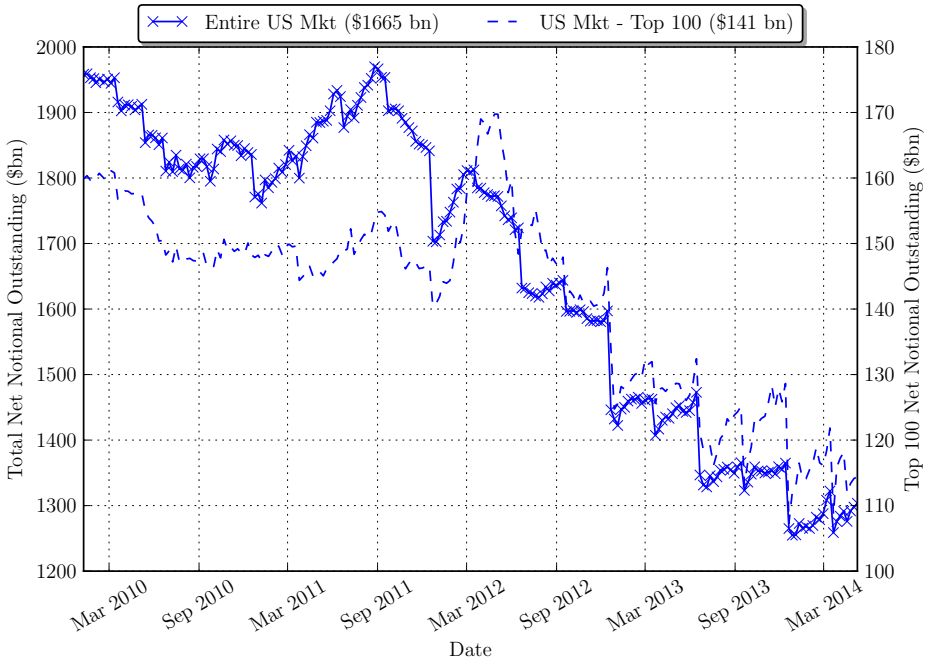

Notes: The left axis of this figure plots the net notional outstanding in U.S. CDS markets, computed as $NO(t) = \sum_{r \in \mathcal{R}_t} \sum_{c \in \mathcal{C}_t} \max\left(NS(c,r,t), 0\right)$. The size of the U.S. market is the average of the net notional using all positions in the dataset and positions for reference entities that I can definitely classify as based in the U.S.. I take the average of the two because the former is an upper bound on the size of the U.S. market, and the latter is a lower bound. A U.S. reference entity is defined according to the methodology in the Data Appendix. The right axis of this figure shows the net notional outstanding for the largest 100 reference entities. In the legend, time-series averages are in parentheses.

declined 33 percent to $1.3 trillion as of May 2014. The average total net notional outstanding over the entire sample is $1.7 trillion.

Despite the downward trend in the size of the CDS market, the amount of credit risk transferred is still large. As a rough comparison of magnitude, the size of the U.S. corporate bond market is approximately $9 trillion, so CDS markets are anywhere from 15-20 percent of the size of corporate bond markets. These results are echoed by Oehmke and Zawadowski (2013), who find the ratio of CDS net notional to debt outstanding is, on average, 19.7 percent. However, they consider net outstanding notional through single name CDS only. My estimates take into account positions via single name CDS *and* index CDS, while also encompassing a wider range of reference entities.

The dotted line on the right axis of Figure 1 also shows that the net notional outstanding is concentrated in the top 100 reference entities. For example, in May 2014, the top

100 reference entities represented less than 2 percent of all traded reference entities, but accounted for 11 percent of total net notional outstanding.[13] I will keep this fact in mind when computing, say, the average market share of a counterparty across reference entities. Clearly, more weight must be given to sellers and buyers who have a significant presence in the market for the largest reference entities.

3.2 Concentration of Buyers and Sellers of CDS Protection

In reality, there are many ways to measure the concentration of buyers and sellers of CDS. In this section, I use a simple measure that can be interpreted as the market share of aggregate net selling. In the Online Appendix, I take a complementary approach that looks at concentration within a reference entity, then commonality of buyers and sellers across reference entities. No matter the route taken, the end conclusion is the same: a small set of buyers, and an even smaller set of sellers, are responsible for most of the credit risk transferred by CDS.

To start define counterparty c's market share in a single reference entity r as:

$$MS_S(c,r,t) = \frac{NS(c,r,t)}{NO(r,t)}$$

where, again, $NS(c,r,t)$ is the net amount sold by c on r, and $NO(r,t)$ is the net notional outstanding on r. If $MS_S(c,r,t) = 20$ percent, then c accounts for 20 percent of all selling in r. Conversely, if $MS_S(c,r,t) = -20$ percent then c accounts for 20 percent of all buying in r. The subscript S serves as a reminder that positive values reflect sellers of protection.

Next, to compute the *aggregate* share of selling for each counterparty, I take a size-weighted average across all reference entities:

$$\begin{aligned} MS_S(c,t) &:= \sum_{r \in \mathcal{R}_t} \omega_{rt} MS_S(c,r,t) \\ \omega_{rt} &= NO(r,t) / \sum_r NO(r,t) \end{aligned} \tag{3}$$

where I use a size-weighted average instead of an equal-weighted average to offset the influ-

[13] To put more structure on the cross-sectional size distribution, I used data from February 28, 2014 to estimate a power law coefficient via the rank 1/2 estimator of Gabaix and Ibragimov (2011). The estimated power law coefficient for CDS size is 0.48 (with t-statistic of 171.708), thereby confirming the largest reference entities play an outsized role in CDS markets.

Figure 2: Share of Top Five Aggregate Sellers and Buyers of CDS Protection

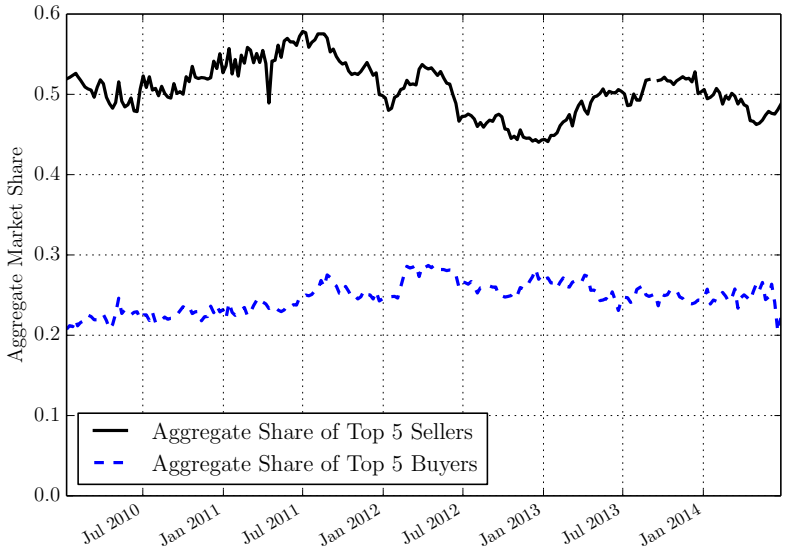

Notes: This figure plots the aggregate share of the top five sellers and buyers of CDS protection through time. The share of a single counterparty c in a given reference entity is the proportion of net selling by c in that reference entity. The aggregate share of net selling by c is the size-weighted average share across all reference entities. The top five sellers are those with the largest aggregate share, and the top five buyers are those with the most negative aggregate share. I convert the market share of buyers to a positive number because my definition assigns negative shares to net buyers.

ence of extremely small reference entities.

$MS_S(c,t)$ is a parsimonious measure of the importance of c as a seller in the aggregate economy. If c is a seller in the largest reference entities, then $MS_S(c,t)$ will be large and positive. Similarly, if c is a buyer in the largest reference entities, then $MS_S(c,t)$ will be very negative. Notice, though, if a counterparty offsets net positions across reference entities (i.e. sells in one name, and buys in another), then its aggregate share will tend towards zero.

In turn, I define the top five aggregate sellers at each point in time as the traders with the largest $MS_S(c,t)$. The top five buyers are the five counterparties with the most negative $MS_S(c,t)$. Figure 2 then plots the total share of the top five sellers and buyers, respectively, through time. For illustration, I have converted the market share of buyers to a positive number because, again, my definition assigns negative shares to net buyers.

Net sellers of CDS are highly concentrated. According to my definition of market share, the top five sellers account for 50 percent of all protection sold. Put differently, because there are about 1700 counterparties in the market, 50 percent of all selling is in the hands of less than 0.1 percent of potential counterparties. Buyers are also concentrated, albeit only half

Figure 3: Persistence of Top Five Aggregate Sellers and Buyers

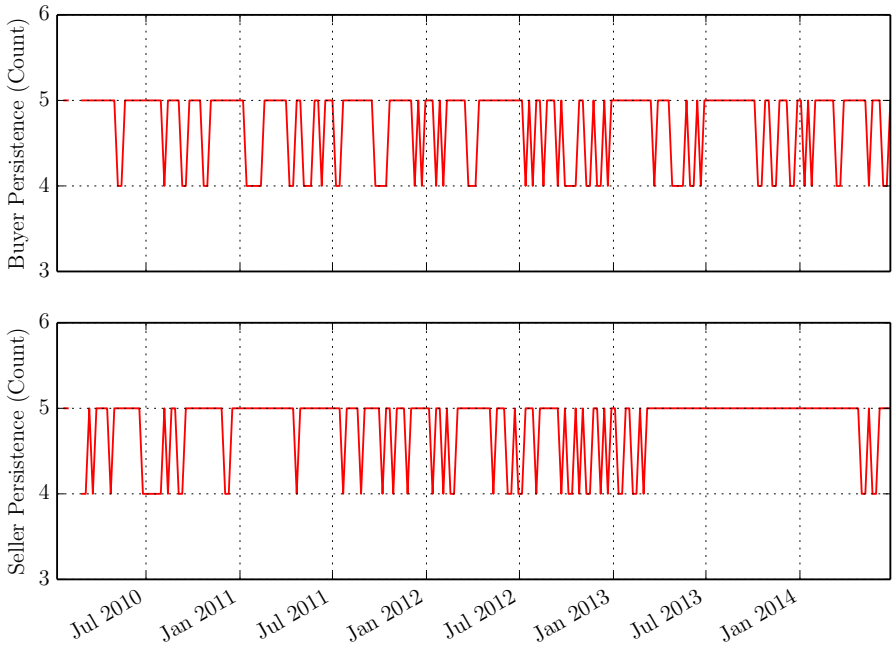

Notes: This figure plots the persistence of the aggregate share of the top five sellers and buyers of CDS protection through time. For week t, I count the number of the top five buyers who are also in the top five in week $t-1$. I do the same for the persistence of top sellers.

as concentrated as sellers. The top five buyers are responsible for roughly 20-25 percent of net buying in the aggregate. In addition, the share of the top five sellers and top five buyers is relatively constant throughout my sample period.

The identities of the top five buyers and sellers are also persistent through time. Figure 3 plots, for both buyers and sellers, the count of top five counterparties that remains the same from time $t-1$ to t. For example, in week t, I count the number of top five sellers who were also in the top five in the previous week. On average, 94 percent of the top five buyers and 96 percent of the top five sellers remain constant from week to week. Thus, not only are CDS markets highly concentrated with a few mega-buyers and mega-sellers, but this organizational feature of the market is also fairly static through time.

3.3 Who Bears Credit Risk and Who Buys Protection?

Given the size and concentration of the CDS market, it is natural to ask: who are the mega-sellers and mega-buyers of credit protection? I answer this question by assigning every

Figure 4: Aggregate Proportion of Buying and Selling, by Counterparty Type

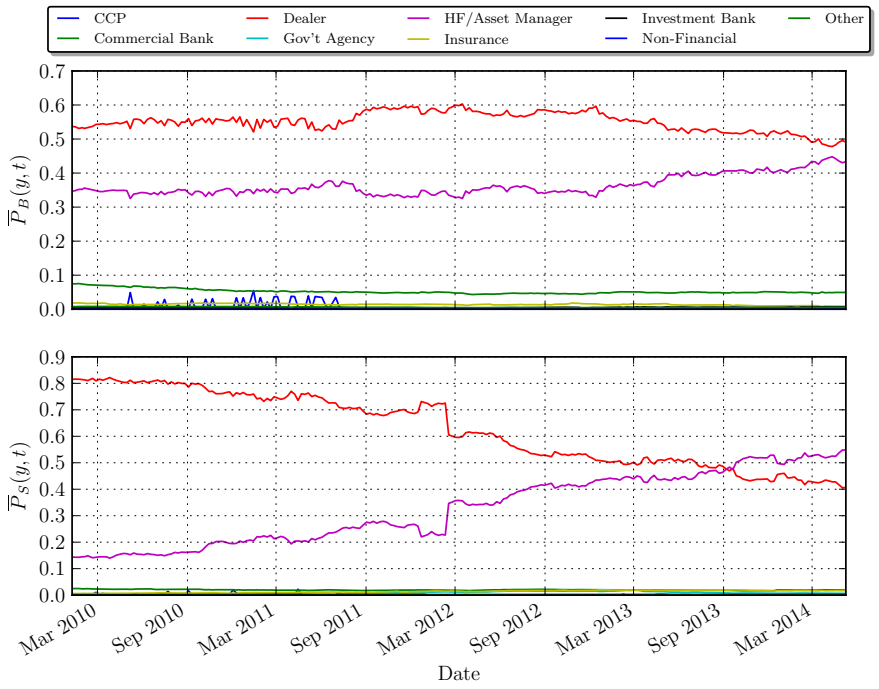

Notes: This figure plots the aggregate proportion of net selling and buying done by each counterparty type. For each reference entity, I compute the proportion of net buying and selling by each counterparty type. To aggregate, I compute the size-weighted average, across reference entities, of the proportion bought and sold by each type.

counterparty in my dataset to one of the types listed in Table 3 of the Online Appendix. Examples of types are commercial banks, insurance companies, and dealers.

Next, for each reference entity and date, I compute the proportion of net buying and selling done by each type. For instance, I compute what proportion of GE's net outstanding is sold by insurance companies. The computation is analogous to calculating the market share of an individual counterparty in a reference entity, except I do so for a counterparty type. Finally, I create an *aggregate* index of the proportion bought and sold by each type y, which I denote by $\overline{P}_B(y,t)$ and $\overline{P}_S(y,t)$, respectively. Each aggregate index is simply type y's size-weighted average market share across all reference entities.[14]

Figure 4 plots both $\overline{P}_S(y,t)$ and $\overline{P}_B(y,t)$ for all counterparty types through time. The top panel begins with the aggregate proportion of buying by counterparty type. Dealers and hedge funds/asset managers (HFAMs) are the two largest buyers. In the aggregate, dealers have consistently purchased approximately 55 percent of protection, with the remaining

[14]More details of these computations are also found in the Online Appendix, Section 1.3.3.

buying going to HFAMs.

The aggregate proportion of selling by counterparty types appears in the bottom panel of Figure 4. In contrast to buyers, the composition of sellers has dramatically changed since 2010. At the beginning of the sample, dealers accounted for 80 percent of all protection sold in U.S. CDS markets, with this share heavily skewed towards less than five dealers (approximately 50 percent of aggregate selling). Nonetheless, the total proportion sold by dealers has declined by almost half, with dealers accounting for almost 40 percent of total selling by the end of the sample.[15] Instead, HFAMs have grown into a much larger role in bearing credit risk via selling in U.S. CDS markets. More specifically, less than five HFAMs account for nearly 30 percent of all net selling of protection as of the first quarter of 2014.

Why Are Markets So Concentrated? Intuitively, concentration develops naturally in any market with high fixed entry costs.[16] CDS markets are costly to enter for a few reasons. First, trading CDS requires back-office processing of trades and risk management to manage existing positions. To this point, many smaller hedge funds will pay their dealer an additional fee in return for the dealer handling the oversight of trades. Moreover, establishing a CDS desk requires substantial information acquisition (Merton (1987)), not only in terms of hiring traders and managers with expertise in credit risk, but also specifically in credit derivatives. Second, CDS trading is similar to banking in the sense that relationships are "sticky".[17] For example, in the Online Appendix I show that the average non-dealer trades with only three counterparties. In lieu of the costs of building new trading relationships, it is no surprise that trading activity in all over-the-counter markets is dominated by a handful of dealers who can use their existing relationships from many lines of business. Third, operating a CDS desk is costly from a funding standpoint. Since the 2007-09 crisis, it has been common practice for CDS positions to be marked-to-market every day. Consequently, CDS desks need a stable source of funding in order to survive daily fluctuations in mark-to-market values. There are large economies to scale in terms of funding, and as a result, large dealers and hedge funds naturally emerge as key players in the market.

Of course, there are a multitude of additional reasons why CDS markets are concentrated. The purpose of this paper is not to answer this question, but rather to understand how limited

[15] The 40 percent can be further decomposed as follows: less than five dealers account for 26 percent of all total selling, with other dealers accounting for the remaining 14 percent.

[16] By concentration, I mean the large net buyers and net sellers that I have documented in this section. This is a slightly different than the concept of concentration put forth in Atkeson et al. (2014a), who are more focused on large intermediaries who are on both sides of many trades.

[17] See, for example, Chodorow-Reich (2014).

capital in the market ultimately affects pricing in CDS. Nonetheless, my findings do shed some light on the question of concentration. For instance, because I find that HFAMs have become a dominant seller of CDS protection, it is unlikely that relationships play a first order role in concentration; if relationships were primarily driving concentration, dealers would always be mega-players. On the other hand, the fact that limited capital does seem to impact prices suggests that funding frictions may be important for explaining the existence of mega-players in CDS. While concentration is certainly an interesting and important topic, further inquiry is outside of the scope of my paper.

4 Capital Fluctuations and the Pricing of Credit Risk: Main Results

The two core findings of this paper are: (i) a decrease in the level of seller capital leads to an increase in CDS spread levels and (ii) an increase in the concentration of sellers generates more volatile risk premiums. In this section, I develop both of these points empirically. Before proceeding, I discuss how I measure risk bearing capital.

Using the Market Value of CDS Positions to Measure Changes in Capital

I define the risk bearing capital of an individual trading desk (counterparty in my case) as the capital available to the desk for the purposes of initiating and maintaining new investments.[18] Naturally, the risk bearing capital of sellers in the market is just the sum of all sellers' capital.

In my empirical work, I treat changes in the mark-to-market value of each counterparty's CDS positions (CDS profits and losses, or P&L) as a direct proxy for changes in their risk bearing capital. I focus on the risk bearing capital at the CDS desk, as opposed to the entire trading entity (e.g. hedge fund or dealer), for a few important reasons. The first reason is best seen by a simple example. Consider a multi-strategy hedge fund that trades in many different asset classes, one of which is CDS. It is not clear that the capital of the entire hedge fund is representative of the capital that the CDS desk adds to the market. An arguably better view is that the CDS desk is allocated some capital upfront and this capital grows or declines based on the performance of the desk. Thus, at least at shorter frequencies,

[18] Even in swaps, capital is required to initiate new trades because of initial margin payments and upfront payments that make the swap NPV zero. Maintaining an existing trade requires capital to make payments on net bought positions, variation margin payments, and in the case of net sellers, potential default payments.

fluctuations in the risk bearing capital of a desk can be proxied by fluctuations in the mark-to-market value of the positions it takes. Mitchell, Pedersen, and Pulvino (2007) provide evidence consistent with this story by showing that information barriers within a firm can lead to capital constraints for specific trading desks who have experienced mark-to-market losses.

Of course, purchases of CDS protection are often used to hedge underlying corporate bond positions. In the case of a counterparty who buys CDS to hedge their corporate bond portfolio, the capital of the CDS desk alone likely does not capture the true dynamics of their risk bearing capital. However, the wealth of the CDS desk should capture the risk bearing capital of large protection sellers because it is unlikely their positions are hedged with other securities. At the end of the day, this debate can be resolved empirically. If the P&L of the CDS trading desk tracks changes in risk bearing capital well, then P&L should also help explain price movements. Consistent with this argument, I find that seller capital, not buyer capital, impacts prices.

From an institutional perspective, it is natural to think CDS desk-specific capital is the correct state variable for pricing (e.g. effective risk aversion). Trading desks at large dealers and hedge funds are subject to risk limits (e.g. value-at-risk), which may tighten with prolonged losses. More importantly, poor portfolio performance means CDS traders have less capital to make variation margin payments on mark-to-market losses. If raising new capital on short notice is costly, losses will naturally constrain the ability to take on new risk, thereby raising the effective risk aversion of protection sellers.[19] Section 7.1 provides more institutional details and some empirical evidence consistent with this interpretation.

The final reason I use P&L as a measure of risk bearing capital is practical. Recent empirical research on slow moving capital and limited intermediary risk bearing capacity has used leverage as a measure of risk bearing capacity.[20] The theoretical underpinnings of this work suggest leverage is a sensible metric because it is a proxy for the wealth available for bearing risk. In some sense, I have a more direct measure of this wealth because I have proprietary data on actual positional holdings, which means I can compute the dollar value

[19] See Atkeson, Eisfeldt, and Weill (2014a) for a theoretical example of risk limits. Another potential channel for losses to affect pricing follows from Froot, Arabadjis, Cates, and Lawrence (2011), who demonstrate that loss aversion for institutional investors affects future trading. Froot and O'Connell (2008) also develop a model where costly external financing of intermediaries leads to above-fair pricing of catastrophe insurance. They find their effect to be particularly strong after large losses.

[20] e.g. He and Krishnamurthy (2013), Adrian and Boyarchenko (2013), and Brunnermeier and Sannikov (2014). Even within this literature, it is unclear whether leverage should be measured using book values or market values.

of each counterparty's CDS portfolio. Moreover, I find that hedge funds are a large player in CDS, but leverage measures for these entities are either non-existent or poorly measured. Using the P&L of counterparties to measure of changes in their risk bearing capital crucially allows me to account for traders who do not have better available alternative proxies for wealth.

Computing the mark-to-market value of each counterparty's CDS portfolio is itself a computationally challenging task. It requires me to mark over 600 million CDS positions to market for each day in my sample period. To keep the problem manageable, I choose the simplest possible methodology, with the details found in Section 2 of the Data Appendix.

4.1 Risk Bearing Capital and the Level of CDS Premiums

The first message of my paper is that risk premiums in CDS depend on the total amount of capital behind natural CDS protection sellers. By natural sellers, I mean those counterparties who are most often net sellers of protection.[21] I make this point in the most straightforward fashion by asking whether CDS spread movements are explained by simultaneous movements in the risk bearing capital of natural CDS protection sellers.

In addition, I begin my analysis by focusing on the risk bearing capital of mega-sellers because their capital fluctuations effectively represent capital fluctuations of all sellers. This statement follows directly from the fact that sellers of CDS are extremely concentrated, as shown in Section 3. In Appendix A, I study capital fluctuations of all natural sellers, though starting with mega-seller capital also sets the stage for examining why the distribution of risk bearing capital is relevant for pricing.

The difficulty in explaining movements of CDS spreads with changes in capital is reverse causality: are capital fluctuations (i.e. mark-to-market changes) causing CDS spread movements or vice versa? One way to circumvent this issue is by testing whether losses in one part of a mega-seller's portfolio influence pricing of other, unrelated portions of their portfolio. My approach is similar in spirit to Froot and O'Connell (2008), who show that losses to a large seller of catastrophe reinsurance affect the pricing of this insurance. Their identification technique examines, for example, whether a hurricane in Florida causes prices to rise for freeze damage insurance in New England.

[21] I provide a formal definition of natural sellers in Appendix A. Intuitively, it is the set of counterparties who are most often net sellers of protection. Much of my empirical work also considers the effect of buyer capital fluctuations (emanating from their CDS portfolios) and offers strong support of the idea that seller capital is the relevant variable for pricing.

The following regression implements a similar concept in the context of CDS markets:[22]

$$\Delta \log (CDS_{rt}) = a_r + \boldsymbol{\beta}_1^{'} \Delta \mathbf{Z}_{rt} + \boldsymbol{\beta}_2^{'} \Delta \mathbf{X}_t + \zeta_s OCF_{rt}^s + \zeta_b OCF_{rt}^b + \varepsilon_{rt} \tag{4}$$

where CDS_{rt} is the 5-year CDS spread of reference entity r at time t. I obtain these spreads from the data vendor Markit, and relegate further details of the underlying data to the Data Appendix. a_r is a reference entity fixed effect that absorbs any time invariant firm characteristics. $\mathbf{Z}_{rt}$ is a vector containing Moody's 5-year expected default frequency (EDF) and Markit's expected loss-given-default (LGD); I choose these firm-level controls based on reduced form models of credit risk.[23] In some versions of regression (4), $\mathbf{Z}_{rt}$ also includes the CDS spread implied by options markets. To compute an option-implied CDS spread for reference entity r, I translate the price of out-of-the-money put options to CDS spreads using the methodology of Carr and Wu (2011). The details of this procedure are contained in Appendix D.4. The important advantage of using option-implied CDS spreads is they control for a large number of unobservable firm-level and macroeconomic factors that may drive credit spreads.

$\mathbf{X}_t$ is a set of observable macroeconomic variables that may also cause CDS spread movements. I choose these controls based on theoretical models of credit risk and previous research on credit spread variation.[24] These variables are the log equity-to-price ratio for the S&P 500, VIX, TED, CFNAI, 10 year Treasury yield, 10-year-minus-2-year Treasury yield, and the CBOE Option Skew index. After first differencing these aforementioned controls, I also include the excess market return of the CRSP value-weighted index.[25] In some specifications, I replace the vector $\mathbf{X}_t$ with a time fixed effect to ensure that the point estimates in the regression are not biased by any unobservable macroeconomic factors.

The important variables in regression (4) are the OCF measures, which stand for "outside

[22]The two approaches are not, however, directly comparable. Continuing with the hurricane example, Froot and O'Connell (2008) consider demand effects by controlling losses to buyers of hurricane insurance. The logic is after a hurricane, buyers will update their probability models and demand more hurricane insurance. To capture this idea, I have estimated the regression with the losses of the top five buyers of r's CDS specifically. I find the effect of mega-seller risk bearing capacity to be basically the same.

[23]I use the term reduced-form in the spirit of the work by Jarrow and Turnbull (1995) and Duffie and Singleton (1999). The popular alternative to this approach are so-called structural models of credit, a la Merton (1974).

[24]e.g. Duffee (1998), Bai and Wu (2012), Collin-Dufresne et al. (2001), Ericsson et al. (2009), or Tang and Yan (2013)

[25]It may seem redundant to include the log-change in the S&P 500 index, but I do so in order to account for higher frequency (weekly) equity movements. The earnings-to-price ratio is monthly and taken from Robert Shiller's website.

capital fluctuations". For example, OCF^s_{rt} captures the capital fluctuations of mega-sellers, with the caveat that these fluctuations are due to changes in the market value of *positions on reference entities outside of r's industry*. Formally, OCF^s_{rt} is computed as:

$$OCF^s_{rt} = \sum_{c \in TS_{t-1}} \Delta V_{c,-r,t}$$

where $V_{c,-r,t}$ is the mark-to-market value of counterparty c's portfolio for all reference entities outside of the same industry as r. TS_{t-1} is the top five aggregate sellers of protection, or what I refer to as mega-sellers, at time $t-1$. OCF^b_{rt} is the same variable, but for the top five aggregate buyers. I include outside capital fluctuations of mega-buyers as a first check of whether their capital levels have any effect on pricing.

Because this regression accounts for firm-level fundamentals (via $\mathbf{Z}_{rt}$) and macroeconomic factors (via $\mathbf{X}_t$ or a time fixed effect), I argue that any impact of OCF on spread changes is attributed to limited capital of sellers and buyers. Table 1 contains the results of regression (4).

It is best to view Column (1) of Table 1 as a benchmark. It is a regression of CDS spread changes on all observable reference entity and macroeconomic controls. The bottom line from Column (1) is that my control variables can only capture 16.4 percent of spread variation on their own.

Column (3) adds the outside capital variables to the baseline regression with firm and macroeconomic controls. As is clear from the point estimates and their standard errors, outside capital fluctuations for large sellers — not large buyers — are an important determinant of spread changes. A one standard deviation capital loss to large sellers on positions from outside of r's industry results in an increase of 2.7 percent in the level of r's CDS spread.[26] To put this in perspective, the standard deviation of spread movements across all firms in my sample is about 6 percent. Thus, a one standard deviation outside loss for mega-sellers creates an effect on spreads that is on the order of 50 percent of a standard deviation. I view this as a lower bound on the effect of seller capital losses on prices, given that I exclude losses coming from positions on firms in r's industry for the purpose of identification.

Column (3) also indicates that capital losses help explain an additional 11 percent of spread variations, which is large given that observable macroeconomic and firm fundamentals

[26] In this setting, the level of the CDS spread is analogous to a price level for a stock. Similarly, the log-change in the CDS spread is analogous to a return for a stock. So in other words, a one standard deviation outside loss for important sellers results in a 2.7 percent "return" for the CDS spread, or an increase in the level of CDS spreads by 2.7 percent.

Table 1: Losses Transmit Across Important Sellers' Portfolios

Dep. Variable	$\Delta \log(CDS_{rt})$					
	(1)	(2)	(3)	(4)	(5)	(6)
OCF^s_{rt}			-0.027		-0.028	-0.028
			(-12.5)**		(-11.2)**	(-8.39)**
OCF^b_{rt}			0.003	0.004	-0.001	-0.000
			(1.93)*	(2.20)**	(-0.55)	(-0.14)
$OCF^s_{rt} \times 1_{OCF^s_{rt} \geq 0}$				-0.022		
				(-8.6)**		
$OCF^s_{rt} \times 1_{OCF^s_{rt} < 0}$				-0.031		
				(-13.1)**		
EDF and LGD	Yes	Yes	Yes	Yes	Yes	Yes
Macro Variables	Yes	No	Yes	Yes	No	No
Time FE	No	Yes	No	No	Yes	Yes
Option-Implied CDS	No	No	No	No	No	Yes
Adj. R^2	16.4	28.9	27.5	27.4	33.2	39.0
N	65,272	65,884	61,869	61,869	62,459	29,412

Notes: This table presents the results of the regression: $\Delta \log(CDS_{rt}) = a_r + \iota_t + \beta' \Delta \mathbf{Z}_{rt} + \zeta_s OCF^s_{rt} + \zeta_b OCF^b_{rt} + \varepsilon_{rt}$. OCF^s_{rt} is the change the mark-to-market value of mega-sellers' CDS portfolio, excluding reference entities in r's industry. OCF^b_{rt} is the same measure, for mega-buyers. Mega-sellers (buyers) are those with the five most positive (negative) market shares, defined in Equation (3). Variables have been standardized to have a mean of zero and variance of one. All standard errors are double-clustered by reference entity and time. **,* indicates coefficient is statistically different than zero at the 5 percent and 10 percent confidence level, respectively. Data is weekly and spans March 2010 to May 2014.

explain only 16.4 percent on their own. Another way to view the incremental R^2 in column (3) versus column (1) is that capital fluctuations for the five largest sellers of CDS protection can explain about one-ninth of the variation in CDS spreads.[27]

Column (4) splits outside capital fluctuations for sellers into two components: one where OCF^s_{rt} is positive and one where it is negative. This specification allows risk premiums to interact non-linearly with risk bearing capital, a common feature of many intermediary-based asset pricing models.[28] The point estimates in column (5) show that reductions in risk bearing capital are more influential than gains, so the direction of the non-linearity fits with intuition.[29]

[27] Column (2) includes a time fixed effect but omits the OCF variables. The R^2 in this regression is 28.9. The fixed effect makes this is a more conservative benchmark. However, I view the proper benchmark as column (1) because these are *observable* variables.

[28] e.g. He and Krishnamurthy (2013) or Brunnermeier and Sannikov (2014).

[29] In this specification, the standard error on the point estimate for mega-buyer outside capital fluctuations is small, most likely because the panel gives me good statistical precision. Still, the economic magnitude is negligible.

Columns (5) and (6) add controls that reinforce the stability of the point estimates in column (3). Column (5) removes macroeconomic controls from the regression and replaces them with a time fixed effect to absorb any unobservable characteristics common to the cross section of firms. Importantly, the point estimate on OCF_{rt}^{s} remains essentially unchanged. Because one might still be concerned that I am omitting important firm-level characteristics, column (6) adds option-implied CDS spreads to the regression. The sample size in this specification is cut in half because of an imperfect match between CDS data and options prices. Still, the key message is that, even after incorporating information on the firm implied by options markets, the effect of large sellers' outside losses on spread movements is statistically significant and about 2.8 percent in magnitude.

In isolation, I cannot use the results in Table 1 to claim that capital fluctuations at mega-sellers are special per se, relative to capital fluctuations in the aggregate. The reason is that mega-sellers represent a large portion of the market (nearly half), so their capital is effectively the total capital in the market. My findings so far only indicate that the total risk bearing capital of sellers moves CDS premiums.

To further illustrate the distinction, consider the following thought experiment: hold the total level of risk bearing capital fixed, but vary the distribution of capital within natural sellers of protection. In this case, most theoretical models would suggest that the *level* of risk premiums should not change, even if all of the capital was allocated to a small set of traders. My next task is to argue why concentration, or the distribution of risk bearing capital, is also important for pricing.

4.2 Concentration of Risk Bearing Capital and the Volatility of Risk Premiums

Perhaps the most obvious reason to care about concentration is fragility. If CDS markets are dominated by a handful of important sellers, then a capital shock to one of these key players will have a sizable effect on the total amount of risk bearing capital, and presumably, prices.[30] As a result, even though the level of risk premiums may be unaffected by the distribution of risk bearing capital, the volatility of risk premiums will depend on concentration. A similar concept for macroeconomic growth has been studied recently by Gabaix (2011) and Kelly, Lustig, and Van Nieuwerburgh (2014).

[30]Later, I use the 2011 Japanese tsunami as a case study to illustrate this idea.

4.2.1 The Aggregate Price of Credit Risk

As a high-level way to quantify the volatility of risk premiums, I begin by estimating an aggregate measure of CDS premiums via the following panel regression:

$$\log(CDS_{rt}) = a_r + \boldsymbol{\beta}' \mathbf{Z}_{rt} + \pi_t + \varepsilon_{rt} \tag{5}$$

where a_r is a reference entity fixed effect and $\mathbf{Z}_{rt}$ captures reference entity fundamentals. The firm-level variables I use for this exercise are Moody's 5-year EDF and Markit's LGD.

The key variable in regression (5) is π_t. Intuitively, at each point in time, it measures the cross-sectional average portion of log-spreads not captured by firm fundamentals. Gilchrist and Zakrajsek (2012) refer to a similar quantity as the "excess bond premium".[31] As a result, I interpret π_t as the log of the price of credit risk. In Appendix D.1, I also use a highly stylized reduced form model of credit risk that is consistent with my interpretation. The empirical details behind my estimate of π_t are also found in Appendix D.1. Keep in mind that, because it is estimated via a time fixed effect, the fitted value of π_t is relative to its own level at the beginning of the sample.

My ultimate objective is to measure the volatility of $\Pi_t := \exp(\pi_t)$, but the level of Π_t is itself interesting. The solid blue line corresponding to the left axis of Figure 5 plots my estimate of the price of credit risk through time. It is clear that there is significant time-series variation in the level of the aggregate price of credit risk. As my analysis from Section 4.1 shows, some of this variation can be attributed to fluctuations in seller risk bearing capital. The two major peaks in this series occur during the summer of 2012, and the late fall/winter of 2011. Macroeconomic news in the fall and winter of 2011 was headlined by concerns over the spread of the European sovereign debt crisis, as well as a downgrade in the credit rating of United States debt. The summer of 2012 also had many major macroeconomic events, most notably the worsening of the European sovereign debt crisis, increased political turmoil regarding the U.S. debt ceiling, and the expiration of the Bush tax-cuts (the "fiscal cliff"). As is clear from the trend at the end of the time series, the price of credit risk basically returned to early 2010 levels by mid-2014.

[31] In fact, these authors estimate the excess bond premium in a very similar fashion. They run a version of the panel regression (5), and then compute fitted credit spreads. They then define the excess bond premium as cross-sectional average difference between actual and fitted credit spreads. π_t is effectively the (log of the) same quantity, but I estimate it in a single step using a time fixed effect. Gilchrist and Zakrajsek (2012) also include some additional firm controls such as a credit rating indicator. The firm fixed effect in my sample accomplishes basically the same thing, given I have a much shorter sample than theirs and ratings do not change vary much in my sample.

Figure 5: Aggregate Price of Credit Risk

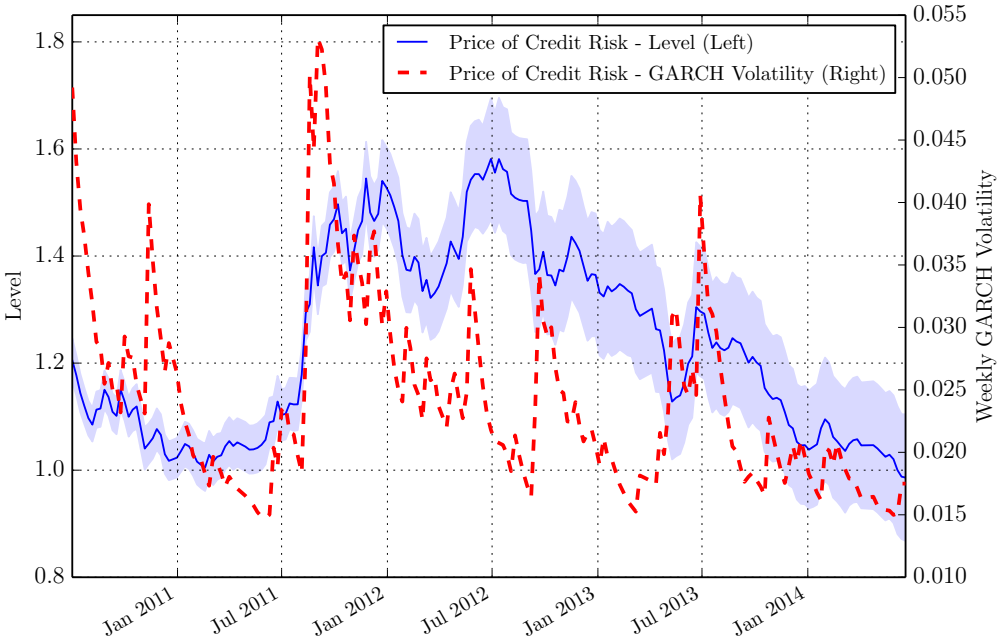

Notes: The left axis of this figure plots the price of risk, $\widehat{\Pi}_t = \exp(\widehat{\pi}_t)$, estimated from the following regression: $\log(CDS_{rt}) = a_r + \beta_1 \log(EDF_{rt}) + \beta_2 \log(LGD_{rt}) + \pi_t + \varepsilon_{rt}$. EDF_{rt} and LGD_{rt} are the Moody's 5 year expected default frequency and industry loss-given-default, respectively, for reference entity r. a_r is a firm fixed-effect and π_t is a time fixed-effect. The shaded region represents 95 percent confidence bands, where the standard errors were computed using the Delta method. All standard errors were clustered by reference entity and time. The right axis of this figure plots the GARCH(1,1) estimate for the series $\Delta\pi_t$.

4.2.2 Volatility in the Price of Credit Risk and Seller Concentration

The red dashed line in Figure 5 (right axis) plots an estimate of the *volatility* of the price of credit risk. I provide the full details of this estimate shortly, but the key observation is there is also significant time-series variation in the volatility of Π_t. My second major point in this paper is that this volatility is an increasing function of the concentration of CDS protection sellers.

To illustrate the mechanism, I assume a highly stylized representation of equilibrium risk premiums in CDS. Suppose that the price of credit risk, Π_t, at any given point in time is a decreasing function of the total amount of capital held by natural sellers, denoted by C_t:

$$\Pi_t = f(C_t)$$

where $f^{'} < 0$.[32] Additionally, suppose that the capital of each natural seller s in the market is $C_{s,t}$ and that individual capital evolves according to:

$$C_{s,t+1} = C_{s,t}\left(1 + \varepsilon_{s,t+1}\right)$$

where $\varepsilon_{s,t+1}$ is an independent shock to seller s's risk bearing capital. To focus on concentration, I assume the volatility of $\varepsilon_{s,t+1}$ is the same across sellers and given by σ_{ε}. It is certainly possible to allow richer dynamics in capital evolution, but this simple structure is enough to demonstrate the main idea. After using the identity that $C_t = \sum_s C_{s,t}$, a simple first order Taylor approximation yields the following formula for the growth rate of the price of credit risk:

$$\begin{aligned} \frac{\Delta\Pi_{t+1}}{\Pi_t} &= \sum_s \omega_{st}\varepsilon_{s,t+1} \\ \omega_{st} &= \frac{C_{st}}{\kappa + C_t} \end{aligned} \tag{6}$$

where κ is a constant from the Taylor approximation.[33] ω_{st} measures the contribution by seller s to the overall stock of risk bearing capital. In a concentrated market, ω_{st} will be large for the major sellers.

Next, I compute the volatility of the growth in the price of credit risk as:

$$\sigma\left(\frac{\Delta\Pi_{t+1}}{\Pi_t}\right) = H_t^S \sigma_{\varepsilon}$$

where $H_t^S = \left(\sum_s \omega_{st}^2\right)^{1/2}$ is a measure of the concentration of sellers in the economy and is effectively a standard Herfindahl index.

Intuitively, when natural sellers of protection become more concentrated, their idiosyncratic capital shocks do not "wash out" in the aggregate, thereby generating excess volatility in the price of credit risk.[34] This is one sense in which the concentration of risk bearing capital — and not just the level — matters for pricing.

[32] See Duffie and Strulovici (2012) for a complete model with this equilibrium feature.

[33] Specifically, $\kappa = \left[f(\overline{C})/f^{'}(\overline{C})\right] - \overline{C}$, where $\overline{C}$ is the expansion point. In a simple case where $\Pi_t = -bC_t$ for some $b > 0$, then $\kappa = 0$.

[34] Indeed in my setup, if H_t^S tends to zero as the number of sellers goes to infinity, then the price of credit risk will have no expected volatility. It is easy to add a common set of factors to each seller's capital stock evolution in order to generate a baseline level of volatility. In either case, I define "excess" volatility relative to the benchmark case when idiosyncratic shocks die out in the aggregate.

With this motivating example in mind, I estimate the volatility of log-changes in the price of credit risk, $\Delta\pi_t$, using a standard GARCH(1,1) model. I denote the estimated volatility series for $\Delta\pi_t$ by σ_t^π. In addition, I compute the Herfindahl for natural sellers and buyers at each point in time as follows:

$$\begin{aligned} H_t^S &= \left[\sum_{c\in A_t^s} \left(MS_S(c,t)\right)^2\right]^{1/2} \\ H_t^B &= \left[\sum_{c\in A_t^b} \left(MS_S(c,t)\right)^2\right]^{1/2} \end{aligned}$$

where $MS_S(c,t)$ is counterparty c's aggregate market shares from Section 3.2.[35] A_t^s and A_t^b are the set of natural sellers and buyers at time t, respectively. I define natural sellers as those with a positive market share and natural buyers as those with a negative share.

I then estimate a simple regression to determine the effect of concentration on volatility:

$$\log(\sigma_{t+1}^\pi) = a + \rho\log(\sigma_t^\pi) + \psi_s H_t^S + \psi_b H_t^B \tag{7}$$

where I use log values to avoid potential econometric issues stemming from the fact that σ^π is non-negative. I include the buyer Herfindahl in case the concentration of buyers also affects the volatility of risk prices. Table 2 presents the results of this regression analysis.[36]

Consistent with other results in this paper, the concentration of buyers does not appear to impact the volatility of Π_t. On the other hand, as sellers become more concentrated, the price of credit risk becomes more volatile. The point estimate on H_t^S is statistically significant at conventional levels and is economically large. A one standard deviation increase in the seller Herfindahl results in a nearly 2 percent increase in the level of σ_{t+1}^π. To put this in perspective, the standard deviation of the level of σ_{t+1}^π is 1.1 percent. Thus, a one standard deviation increase in the seller Herfindahl results in a nearly 1.75 standard deviation event for σ_{t+1}^π.

[35] Note that my definition of aggregate market share means the total weights will not sum up to 1. This is not an issue for my analysis, since I just need a coarse measure of seller concentration.

[36] In unreported results, I also use quasi-maximum likelihood to estimate a GARCH model for the price of risk that directly incorporates the seller and buyer Hefindahls into the volatility recursion. The results are essentially the same, with only the seller Herfindahl showing economic and statistical importance for the volatility of $\Delta\pi_t$. The t-statistic on the seller Herfindahl in the GARCH model is 2.7. I opted to present the results via regression (7) because of its simplicity.

Table 2: Concentration and Volatility of the Price of Credit Risk

Dep. Variable	$\log\left(\sigma^{\pi}_{t+1}\right)$
	(1)
$\log\left(\sigma^{\pi}_{t}\right)$	0.87
	(28.1)
H^S_t	0.019
	(2.10)**
H^B_t	0.014
	(1.45)
Adj. R^2	83.0
N	198

Notes: This table displays the results of estimating the following regression: $\log(\sigma^{\pi}_{t+1}) = a + \rho \log(\sigma^{\pi}_{t}) + \psi_s H^S_t + \psi_b H^B_t$. σ^{π}_t is obtained from a GARCH(1,1) model applied to log-changes in the price of credit risk, $\Delta\pi_t$. H^S_t is the Herfindhal of natural sellers at time t, where weights in the Herfindahl are computed using the aggregate market share measure from Section 3.2. H^B_t is the same Herfindahl measure for buyers. Sellers are defined as those with positive aggregate market share and buyers are those with negative share. Data is weekly and spans July 2010 to May 2014. All variables have been transformed to have a mean of zero and variance of one. **,* indicates coefficient is statistically different than zero at the 5 percent and 10 percent confidence level, respectively.

4.3 Robustness

In Appendix A, I provide additional empirical evidence intended to bolster the interpretation of my main results.

Total Seller Capital and the Price of Credit Risk First, I show that capital losses for all natural sellers, as opposed to just mega-sellers, are followed by increases in the price of credit risk, Π_t. Next, I use a simple variance decomposition to show that one-fifth of unexpected changes in Π_t can be explained by capital shocks to natural sellers of CDS. Within this framework, I then show that mega-seller losses have nearly the same impact on Π_t as do losses to all natural sellers, thereby confirming that mega-seller capital fluctuations effectively represent the entire market. This also relates to my point about fragility: the risk bearing capital of only a few protection sellers plays a large role in CDS pricing.

The Distribution of Capital and the Price of Credit Risk On the other hand, there is no response in Π_t to losses at the average CDS seller, which suggests that the distribution of capital is important. If the only the total amount of risk bearing capital mattered, then a dollar loss should have the same impact on prices, regardless of where it comes from. At the end of this analysis, I also provide an in-depth discussion of why my results indicate that it

is seller capital, not buyer capital, that is most important for CDS pricing.

In the last exercise in Appendix A, I explicitly demonstrate that concentration induces fragility by relating prices to mega-seller idiosyncratic shocks. To do so, I employ a three-step methodology. In step one, I estimate time-series regressions of each counterparty's capital fluctuations on the a set of factors that capture a large portion of CDS spread dynamics. The factors are the principal components of a panel of CDS spread changes. I interpret the regression residuals from this step are idiosyncratic dollar shocks for each counterparty. In step two, for each reference entity and date, I compute a weighted average of seller (buyer) idiosyncratic dollar shocks, where I give more weight to large sellers (buyers). I call these variables "buyer and seller granular residuals" in the spirit of Gabaix (2011). Finally, in step three, I find that negative shocks to large sellers cause CDS spreads to increase, which reaffirms the importance of concentration for pricing.

To summarize, my empirical results lead to two broad conclusions. First, the total amount of risk bearing capital for natural sellers of CDS plays an important role in generating time variation in the price of credit risk.[37] Second the distribution of risk bearing capital is also important for CDS pricing. Because CDS markets are so concentrated, the total amount of risk bearing capital in the market is largely a function of only a handful of sellers, and in turn, a shock to one of these key players has large aggregate pricing effects. Consequently, as sellers become more concentrated, default risk premiums become more volatile.

5 A Natural Experiment: The 2011 Japanese Tsunami

I now turn to a natural experiment that will further support my argument that both the magnitude and distribution of capital losses are important for pricing. The event I focus on is the Japanese tsunami of March 2011, which was the result of a magnitude 9.0 earthquake off the coast of Tohoku. The tsunami occurred on a Friday, and had a significant impact on the risk of the country as a whole. For example, Japan's sovereign CDS spread increased by 50 percent from 80 to 115 basis points on the following Monday. Appendix C.1 contains additional background information on the tsunami, and its after-effects.

[37] A natural objection could be that seller capital only impacts CDS pricing, but not actual bond pricing. In other words, my results may pertain only to the CDS-bond basis, as opposed to the actual price of credit risk. Still, I have verified that the CDS-bond basis is quite small during my sample period, which implies that CDS markets and bond markets did not diverge from each other. This analysis is found in Appendix B. In this appendix I also conduct a series of tests that show seller capital losses impact actual bond prices, not the CDS-bond basis.

5.1 U.S. Reference Entity Exposure to Japan via Sellers and Buyers

To clarify the logic of my approach, suppose Hedge Fund A had sold a great deal of CDS protection on Japanese firms, but Hedge Fund B had not. After the tsunami, capital losses accrue to Hedge Fund A because the Japanese firms they have written protection on are now fundamentally more risky; however, the same does not hold true for Hedge Fund B.[38] My hypothesized mechanism then suggests the U.S. firms for whom Hedge Fund A is large seller will experience increases in their CDS risk premiums. On the other hand, U.S. firms where Hedge Fund B is a large seller will not see their spreads rise.

In Appendix C.1, I verify that U.S. counterparties had large CDS exposures to Japanese firms. I also show that the tsunami caused non-negligible mark-to-market losses for many U.S. counterparties. This is crucial, since my econometric approach requires the shock of the tsunami to materially affect the risk bearing capacity of large players in the U.S. market.

5.1.1 Measurement

To formalize the preceding thought experiment, I construct measures of how exposed a U.S. reference entity r was to the tsunami *through its sellers and buyers*:

$$\begin{aligned} \Gamma_{S,r} &:= \sum_{c \in \mathcal{S}(r)} \left[\frac{NS(c,r)}{NO(r)} \right] \times NS(c, Japan) \\ \Gamma_{B,r} &:= \sum_{c \in \mathcal{B}(r)} \left[-\frac{NS(c,r)}{NO(r)} \right] \times NS(c, Japan) \end{aligned} \tag{8}$$

All of my measures are computed as of March 11, 2011, so I omit time dependencies for brevity. Here, $NS(c, Japan)$ is the net amount sold by counterparty c on Japanese firms. $\mathcal{S}(r)$ and $\mathcal{B}(r)$ are the set of sellers and buyers, respectively, of reference entity r. $\Gamma_{S,r}$ is the weighted average exposure of r's sellers to Japan. The term in brackets is the weight, and is the proportion of total net outstanding for r that is sold by c. $\Gamma_{B,r}$ carries the same intuition for buyers, and is the weighted average exposure of r's buyers to Japan.[39] When referring to both $\Gamma_{B,r}$ and $\Gamma_{S,r}$ in tandem, I will often just abbreviate using Γ.

In the absence of identification issues, we would then expect firms with high levels of

[38] Alternatively, as I show below, the tsunami caused credit spreads for Japan to rise by 50 percent. The mark to market losses to sellers would therefore be unexpected, and in turn reduce their ability to bear additional credit risk.

[39] The negative sign in the definition of $\Gamma_{B,r}$ is just to make sure the weights are positive and sum to 1, as opposed to negative and sum to -1.

$\Gamma_{S,r}$ to experience a rise in their risk premiums. The sellers of "high $\Gamma_{S,r}$" reference entities experience adverse shocks to risk bearing capacity from the tsunami, and in turn increase the premium they require for selling CDS in U.S. reference entities.[40]

5.2 Transmission of the Japanese Tsunami to U.S. CDS Spreads

To tease out my main hypothesis, I estimate variants of the following OLS regression:

$$\Delta \log(CDS_{r,1}) = a + \phi_1 \Gamma_{S,r} + \phi_2 \Gamma_{B,r} + \boldsymbol{\beta}' X_r + \varepsilon_r \tag{9}$$

where X_r is a vector of observable reference entity characteristics that I will discuss shortly. $\Delta \log(CDS_{r,1})$ is the log-change in r's CDS spread in the week following the tsunami. To reiterate, I consider only U.S. reference entities. There are certainly identification issues with attributing changes in CDS spreads after the tsunami with high levels of Γ, as the regression (9) would suggest. One obvious example is that sellers with large Japanese exposures also specialize in U.S. reference entities that are fundamentally linked to the Japanese economy. In Appendix C, I fully frame the identification issues and rule out this "specialization" hypothesis for both buyers and sellers of U.S. reference entities.

X_r controls for changes in observable reference entity fundamentals following the tsunami, I use the change in Moody's 5-year EDF, the change in Markit's LGD, and the equity return of the firm. Including the equity return of the firm is compelling from the perspective of structural models of credit, where any shock to credit spreads is the same as a shock to equity. In many ways, including the equity return of each reference entity allows me to dramatically reduce the number of necessary control variables, since any residual changes in CDS spreads must be driven by something independent of equity market movements.[41]

Because certain industries may have been more exposed to Japanese firms, X_r also contains a fixed effect corresponding to each reference entity's NAICS code. I also include level of CDS spreads for each reference entity on 3/11/2011 to allow for the possibility that Γ captures sellers/buyers who specialize in riskier credits. Finally, I include the 90-day running

[40] I am able to categorize the shock of the tsunami as a negative shock to sellers since, as I argued earlier, sellers of CDS protection are unlikely to be hedged in their position. The effect of the tsunami to large buyers CDS on Japanese firms is less clear. Indeed, the rise in Japanese-related risks that accompanied the tsunami would positively impact buyers' CDS portfolios, but if they owned Japanese bonds then this effect would be offset.

[41] In Appendix C.4, I run a simple regression of $\Gamma_{i,r}$ on these reference entity characteristics. My results indicate that reference entity characteristics explain almost no variation in either $\Gamma_{i,r}$ measure, which is consistent with my identification scheme.

Table 3: Transmission of Japanese Tsunami to U.S. Reference Entities via Large Sellers

Dependent Variable	$\Delta \log(CDS_{r,1})$					
	(1)	(2)	(3)	(4)	(5)	(6)
$\Gamma_{S,r}$	0.014		0.015	0.013		0.013
	(2.57)**		(2.63)**	(2.27)**		(2.41)**
$\Gamma_{B,r}$		0.002	0.003	0.001		0.003
		(0.35)	(0.70)	(0.28)		(0.67)
$\Gamma^{avg}_{S,r}$					0.002	-0.002
					(0.29)	(-0.27)
$\Gamma^{avg}_{B,r}$					-0.003	-0.004
					(-0.84)	(-0.73)
Control Variables	No	No	No	Yes	Yes	Yes
Total N	288	288	288	288	288	288
Adj. R^2	1.9%	0%	2.4%	20.4%	17.5%	21.0%

Notes: The table presents results from the regression: $\Delta \log(CDS_{r,1}) = a + \phi_1 \Gamma_{S,r} + \phi_2 \Gamma_{B,r} + \boldsymbol{\beta}' X_r + \varepsilon_r$. The dependent variable is the change in CDS spread for U.S. reference entities. $\Gamma_{S,r}$ and $\Gamma_{B,r}$ are the share-weighted average CDS exposure of r's net sellers and buyers, respectively, to Japanese firms. Exposure is defined as the net amount of protection sold on Japanese firms. $\Gamma_{S,r}$ and $\Gamma_{B,r}$ have been standardized to have unit variance and zero mean. The control variables are (for each reference entity r): the change in the 5-year Moody's expected default frequency, the change in Markit's loss-given-default, the weekly equity return, the 90-day trailing correlation of (changes in) r's CDS spread with the country of Japan's CDS spread, the 90-day trailing volatility of r's CDS spread, a fixed effect based on the NAICS code of each reference entity, and the level of the CDS spread for r on the day of the tsunami. The time span is the week from March 11, 2011 to March 18, 2011. *, ** represent statistical significance at a 10 and 5 percent level, respectively.

volatility of each reference entity's CDS spread (in log-changes); this allows for the possibility that reference entities who experienced large spread movements post-tsunami are those that have larger volatility. Table 3 summarizes the results of running variations of regression (9).

Consistent with the results in Section 4.1, there is no evidence of a transmission channel via buyers of CDS. Indeed, the coefficient on $\Gamma_{B,r}$ is small and insignificant in all specifications.

The coefficient on $\Gamma_{S,r}$ indicates a strong, positive effect of seller exposure to Japan and subsequent U.S. CDS spread movements. Column (1) estimates a bivariate specification, and columns (3)-(4) sequentially add other control variables. As expected, the coefficient on $\Gamma_{S,r}$ remains stable throughout. Including the full set of controls increases the explanatory power of the regression substantially, and the reduction of the residual variance only serves to reduce the standard error of the point estimate on $\Gamma_{S,r}$.

The magnitude of the transmission effect is large. Consider a U.S. reference entity whose sellers were in the 90th percentile in terms of their exposure to Japanese firms. Similarly,

consider a U.S. reference entity whose sellers were in the 10th percentile. Firms in the 90th percentile saw their spread levels increase 2.4 percent, relative to the 10th percentile, in the week following the tsunami.[42]

Isolating Concentration $\Gamma_{S,r}$ and $\Gamma_{B,r}$ are useful because they simultaneously capture if a reference entity's major sellers were also faced with a capital shock from the tsunami. For this reason, though, they do not allow us to separate the importance of concentration versus total capital losses for a reference entity's spread movements. A simple example illustrates the distinction. Consider two reference entities, r_A and r_B, who have the same two sellers S_1 and S_2. S_1's share of selling in firm r_A is 90 percent, which means that S_2's share is 10 percent. Conversely, S_1 and S_2 have an equal share of selling in r_B. Finally, suppose S_1 had net exposure of 100 to Japanese firms and S_2 had exposure of 10. In this example, the total exposure of r_A and r_B's sellers is the same since they have the same two sellers. Still, we might expect that r_A will be more sensitive to the shock of the tsunami because its primary seller had large exposure to Japanese firms.

I flesh this thought experiment out in the data in two ways. To start, I construct alternative versions of $\Gamma_{S,r}$ ($\Gamma_{B,r}$) by taking simple averages of seller (buyer) exposures to Japan:

$$\begin{aligned}
\Gamma_{S,r}^{avg} &:= \sum_{c \in \mathcal{S}(r)} \left[\frac{1}{\|\mathcal{S}(r)\|} \right] \times NS(c, Japan) \\
\Gamma_{B,r}^{avg} &:= \sum_{c \in \mathcal{B}(r)} \left[\frac{1}{\|\mathcal{B}(r)\|} \right] \times NS(c, Japan)
\end{aligned}$$

where the $\|\cdot\|$ operator denotes the size of a set. $\Gamma_{S,r}^{avg}$ ignores any possible concentration and allows me to compare two reference entities that were, on average, similarly exposed to Japan through their sellers.

Column (5) in Table 3 suggests that reference entities whose sellers had higher average exposure to Japan did indeed see their spreads rise very slightly, but the standard error of the point estimate on $\Gamma_{S,r}^{avg}$ is relatively large. Column (6) includes all Γ variables, both equal and share-weighted versions, in the regression. Even after controlling for the average exposure of each reference entity's sellers to Japan, the point estimate on $\Gamma_{S,r}$ is still economically large and statistically significant. These results highlight that it is critical to consider the

[42] i.e. $\Delta \log(CDS_{r \in 90,1}) - \Delta \log(CDS_{r \in 10,1}) = 0.024$ where, for instance, $\Delta \log(CDS_{r \in 90,1})$ is the log-CDS spread change for firms in the 90th percentile.

Table 4: Concentrated Exposures and Japanese Tsunami Transmission

Dependent Variable	$\Delta \log(CDS_{r,1})$
	(1)
$\omega_{J,r}$	0.085
	(2.8)**
Control Variables	Yes
Total N	175
Adj. R^2	6.56%

Notes: The table presents results from the regression: $\Delta \log(CDS_{r,1}) = a + \eta_J \omega_{J,r} + \boldsymbol{\beta}' X_r + \varepsilon_r$. The dependent variable is the change in CDS spread for U.S. reference entities. $\omega_{J,r}$ is the share of counterparty J in the net selling of reference entity r. J is the counterparty who had the largest exposure to Japanese firms prior to the tsunami. The control variables are (for each reference entity r): the change in the 5-year Moody's expected default frequency, the change in Markit's loss-given-default, the weekly equity return, the 90-day trailing correlation of (changes in) r's CDS spread with the country of Japan's CDS spread, the 90-day trailing volatility of r's CDS spread, a fixed effect based on NAICS industry code, and the level of the CDS spread for r on the day of the tsunami. The time span is the week from March 11, 2011 to March 18, 2011. *, ** represent statistical significance at a 10 and 5 percent level, respectively.

combined effect of concentration and capital losses when explaining spread dynamics.

As a second way to reinforce the importance of concentration, I take advantage of the fact that there was one seller in particular who had an extremely large exposure to Japanese firms just prior to the tsunami (see Appendix C.1). I denote this seller by the index J. The regression I estimate is then:

$$\Delta \log(CDS_{r,1}) = a + \eta_J \omega_{J,r} + \boldsymbol{\beta}' X_r + \varepsilon_r$$

where $\omega_{J,r}$ is the share of J in the net selling of r and X_r is the same set of reference entity controls used throughout this section. Fixing the seller and only varying J's share across reference entities allows me to focus on how concentration interacts with pricing. In addition, in Appendix C.4 I verify that $\omega_{J,r}$ is once again not just a proxy for reference entities with high fundamental exposure to the Japanese economy.

Table 4 collects the results of this regression. As the these results show, reference entities where J had a larger share of selling also experienced larger spread increases after the tsunami hit. To give an economic sense of magnitude, I compare reference entities where J had a high share (90th percentile of $\omega_{J,r}$) to reference entities where J had a low share (10th percentile of $\omega_{j,r}$). High $\omega_{J,r}$ firms saw their CDS spread levels increase by 2 percent following the tsunami, relative to low $\omega_{J,r}$ firms. These results further highlight why the distribution of exposures – in addition to the level — is important for price dynamics

5.3 How Long Does the Effect Last?

The last question I try to answer in my analysis of the Japanese tsunami is how it takes for the shock of the tsunami to dissipate within U.S. reference entities. There are a few ways to answer this question, and I take a simple graphical approach. Of the 288 reference entities I analyze, I divide reference entities into two sets. Reference entities in the "low Γ_S" set are those reference entities whose sellers are below the median Γ_S. In other words, this is the group of reference entities whose sellers were least exposed to Japan. Reference entities in the "high Γ_S" set are those reference entities whose sellers were most exposed to Japan (above the median Γ_S). Within each set, I then compute a weighted-average CDS spread. The weights assigned to each reference entity within each bucket are proportional to $\Gamma_{S,r}$. In the high Γ_S set, reference entities whose sellers were most exposed to Japan get the most weight; in the low Γ_S set, reference entities whose sellers were least exposed to Japan get the most weight. Weighting in this fashion puts additional emphasis on transmission of the tsunami via mega-sellers, and is akin to a "difference in differences" approach. Figure 6 plots the difference in the time series of each set's weighted-average CDS spread.

The main message of the exercise is quite clear from this depiction. Roughly speaking, the spread difference between the two groups of reference entities returns to normal on March 25, 2011, so two weeks after the tsunami (the beginning of the grey shaded region is the day of the tsunami). The two week window is not that short when considering the steps taken by the Japanese government in the aftermath of the tsunami. In particular, the Fukushima Daiichi nuclear disaster was a major catastrophe caused by the tsunami. In the days following the initial shock, it was unclear whether a full-fledged meltdown of the nuclear power plant would ensue. Before Japanese officials were able to stabilize the plant, many experts opined that a full meltdown would be akin to dropping a nuclear bomb on the site. It is easy to imagine the potential for this outcome played a major role in the heightened effective risk aversion for U.S. sellers who were exposed to Japanese firms. By March 25, the Fukushima plant had been relatively stabilized, and the likelihood of a full meltdown dropped sharply. This resulted in a rebound of risk bearing capacity, which was accompanied by the spread difference between high Γ_S and low Γ_S reference entities returning to normal.

The length of the effect should also be interpreted in conjunction with the size of Japanese firm exposure to the overall exposures of each mega-seller or mega-buyer. As mentioned, I present rough magnitudes of this relative size in Appendix C.1, but loosely speaking, for mega-sellers of U.S. CDS, Japanese exposures were roughly 4 to 6 percent of their overall portfolio. While this certainly not trivial, one would also not expect the shock of the tsunami

Figure 6: How Long Did the Shock of the Tsunami Last?

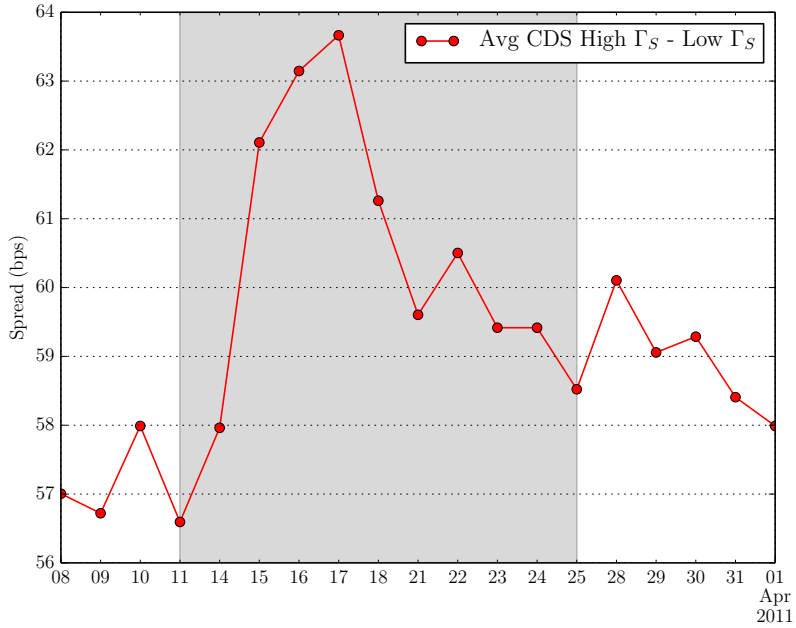

Notes: This figure shows the difference between the weighted-average CDS spread of high $\Gamma_{S,r}$ reference entities and low $\Gamma_{S,r}$ reference entities. High (low) $\Gamma_{S,r}$ reference entities are those whose sellers were most (least) exposed to Japanese firms. The weights assigned to each reference entity within each bucket are proportional to $\Gamma_{S,r}$. The shaded region begins on the day of the tsunami, March 11, 2011. The shaded region ends on March 25, 2011.

to have an enormous effect on their overall risk bearing capacity. In reality, the shock of the tsunami is just one example of a multitude of shocks that are constantly changing the risk bearing capacity of mega-sellers in the market.

6 Identification Using Transaction Prices

In this section, I use empirical methods common to the banking literature to provide an alternative way of identifying the causal effect of mega-seller losses on CDS spreads. My empirical approach is closely related to the within-firm estimators used by, for example, Khwaja and Mian (2008), Chodorow-Reich (2014), or Schnabl (2012). Broadly speaking, these papers try to measure how negative shocks to banks are transmitted to the real economy via lending relationships. In a similar vein, I am trying to measure how depressed risk bearing capacity of mega-sellers is transmitted to CDS prices.

My analysis here is inherently different from the previous portions of the paper and it is critical to keep this in mind when interpreting my results. One difference is that my current

approach uses data on *transactions*, whereas my previous exercises used data on CDS spreads reported by Markit. The latter are representative of *quotes*, not actual transaction prices. A second distinction is that my previous analysis focused on the capital and concentration of ultimate natural sellers of protection. To compute ultimate natural sellers of protection in the market, I accounted for exposures through single name and index positions. When looking at transactions, I will not be concerned with whether the seller in a transaction is also a natural seller of protection. My objective in examining transactions is only to check whether there is micro-level evidence that capital losses for a seller are indeed followed by higher transaction prices.

6.1 Empirical Strategy

To frame the identification problem, suppose CDS_{it} is the spread in transaction i on date t. The determinants of this spread can be characterized by an arbitrary function, $g(\cdot)$:

$$\log\left(CDS_{it}\right) = g\left(H_t^r, H_t^b, RB_{it}^s; X_{it}\right)$$

where H_t^r are determinants of CDS spreads that are firm-specific or common to the macroeconomy; H_t^b are buyer characteristics that determine the CDS spread; RB_{it}^s is the risk bearing capacity of the seller in transaction i, which I will define shortly; X_{it} is a vector of transaction characteristics. These characteristics are the maturity m, log-notional n, and fixed coupon of the transaction f. I work in logs to avoid econometric issues that arise from the fact that CDS spreads are bounded below by zero.

I am interested in estimating $\eta = \partial \log(CDS_{it})/\partial RB_{it}^s$, with the expectation that $\eta < 0$. To estimate this effect consistently, I must — at a minimum — control for all unobserved heterogeneity in the cross-section of buyers, reference entity, and dates. I do so by introducing a buyer-reference entity-time fixed effect. Formally, this means I estimate the following empirical model:

$$\log\left(CDS_{it}\right) = \alpha_{rtb} + \eta RB_{it}^s + \varepsilon_{it} \tag{10}$$

where α_{rtb} is the aforementioned fixed effect for a reference entity-buyer-time triplet. To further illustrate the logic of the regression, suppose Hedge Fund A purchases CDS on GE on date t from multiple sellers. Because I am looking within a buyer, date, and reference entity, any variation in the transacted CDS spread must be driven by seller characteristics or other observable characteristics in the transaction. The argument then goes that sellers

with lower risk bearing capacity charge a higher premium.

The risk bearing capacity measure, RB^s_{it}, is a function of the CDS portfolio value of seller s in transaction i, though I choose not to use the actual portfolio value, V_{st}, for two reasons. On the econometric side, it is a near unit root and, unlike in my previous analysis, I cannot take first differences across transactions. Second, because I am fixing a buyer, reference entity, and date, there will be substantial scale effects if a buyer faces two very differently sized sellers.

To address these issues, I construct RB^s_{it} by giving extra weight to dollar losses on recently open positions and normalizing dollar values by the total gross size of each counterparty's CDS portfolio. The full details are located in Appendix D.3. The resulting RB^s_{it} measure is still to be interpreted as a measure of portfolio performance, but it has better econometric properties than the level of dollar values and makes a comparison of different sized sellers more reasonable because it is normalized by gross values.

6.2 Data Description

As mentioned, the data I use comes from actual transactions. The transaction information provided by the DTCC contains information on the upfront money paid by the buyer to the seller at the inception date, or vice versa. A swap is, in theory, initially a zero NPV trade; the CDS spread we typically think of is the premium the buyer pays to the seller in order to make the value of default protection exactly equal to the value of the premiums.

In practice, things are slightly more complicated. It is standard in CDS markets for the buyer to instead pay the seller a fixed coupon, f. If the fair value CDS spread is, for instance greater than f, then it means the buyer of protection is paying less than she should for the default insurance she receives. In this case, the buyer pays the seller an amount upfront that makes the total NPV of the swap zero. In my data, I observe this upfront amount. In Section 3 of the Data Appendix, I document a simple way to transform an observed upfront amount to an implied hazard rate for the swap. This is important, as it allows me to compare transactions within the same (b, r, t) group that have different maturities and fixed coupon amounts. In the same appendix, I also provide additional information for how I filter transactions and build the final panel data to test the specification in Equation (10).

I consider transactions on all types of CDS (single name, index, and tranche) because my within buyer-reference entity-time estimator already distinguishes product types. This is another reason the results I will present are not comparable to those in Section 4.

Table 5: CDS Transaction Prices After Seller and Buyer Losses

Dependent Variable	$\log(CDS_{it})$	
	(1)	(2)
RB^s_{it}	-0.019	-0.011
	(-4.3)**	(-3.1)**
$\log(Notional_{it})$		0.034
		(4.1)**
$Maturity_{it}$		-0.07
		(-4.4)**
$FixedRate_{it}$		0.51
		(18.0)**
FE	(b, r, t)	(b, r, t)
Total N	31,872	31,872
Within R^2	.1%	20.8%

Notes: This table contains the result of the regression: $\log(CDS_{it}) = \alpha_{rtb} + \eta RB^s_{it} + \varepsilon_{it}$. CDS_{it} is the implied CDS-spread (or rather hazard rate) in transaction i. RB^s_{it} is the risk bearing capacity of the seller s in the transaction, defined as s's total annualized CDS portfolio return. RB^b_{it} is analogously defined. Transactions span February 2010 to June 2014, and are filtered according to the steps outlined in Section 3 of the Data Appendix. All standard errors clustered by seller. *, ** represent statistical significance at a 10 and 5 percent level, respectively.

6.3 Results

Table 5 presents the results of the regression of Equation (10).

Column (1) of the table indicates a one standard deviation increase in seller losses results in an implied CDS spread increase of roughly 2 percent. Column (2) adds other observable features of each transaction such as the notional, maturity, and fixed coupon. Interestingly, the implied CDS spread is increasing in the notional of the transaction, which indicates that sellers who take on a bigger exposure earn an added premium for doing so. One possible explanation is that traders often face risk limits, so they require an added premium to fill larger orders. This idea also rationalizes why we may observe a buyer facing multiple sellers on the same day and reference entity in the first place. If a buyer wishes to fill a large order size, it may have to contact multiple sellers to fully fill the quantity.[43] The more important takeaway is that the coefficient on RB^s_{it} is still large, negative, and significant after the inclusion of these observable transaction characteristics. A one standard deviation loss to sellers still results in an 1.1 percent increase in the implied CDS spread of a transaction, indicating that capital losses affect pricing even at a micro-level.

[43] In addition, buyer-seller relationships are "sticky" as in banking. In the Online Appendix, I show that the average non-dealer faces only three different counterparties.

An additional wrinkle is that when the risk bearing capacity of a seller is low, he may choose to not participate in the market. I confirm this using a simple logit regression in the Online Appendix. In this case, we won't observe any transactions with this counterparty, but that does not mean Markit-reported CDS spreads cannot be high. This is because the Markit-reported CDS spread is a function of transactions *and* quotes. Indeed, one way that sellers can withdraw from the market is to post very high quotes to discourage buyers. As a result, the effect I wish to draw should be even harder to detect using transactional data, which makes my results more likely to be conservative. In this sense, the results from the 2011 Japanese tsunami in Section 5 are more useful for analyzing how a shock to the risk bearing capacity of mega-players affects quotes (i.e. those reported by Markit).

Is It Just Counterparty Risk?

A natural set of objections to the results in Table 5 relate to the issue of counterparty risk. Counterparty risk is the risk that the seller (or buyer) will default on their payments in the CDS. Indeed, Arora, Gandhi, and Longstaff (2012) find some evidence of counterparty risk affecting prices, but the size of their effect is quite small. Nonetheless, consider two cases where this might be driving CDS spreads. The first is RB^s_{it} is itself a proxy for counterparty risk — sellers who have lost a lot of money are bad counterparties. In this case, though, we would expect the coefficient η in the regression to be positive, not negative. Sellers with high RB^s_{it} are good counterparties, and thus can earn a higher premium from buyers. Clearly, the results of Table 5 contradict this story.

The second case is that I have an omitted variable bias. In this case, the "true" regression specification should be:

$$cds_{it}(b, r, t, s, m, n, f) = \alpha_{rtb} + \eta RB^s_{it} + \gamma CP^s_{it} + \varepsilon_{it}$$

where CP^s_{it} measures counterparty risk of the seller in transaction i, and is high when s is a *bad* counterparty. To think about the direction of the omitted variable bias, first notice that we would expect $\gamma < 0$ because bad counterparties must accept lower CDS premiums. Second, it should be the case that $corr(RB^s_{it}, CP^s_{it}) < 0$ because sellers with high risk bearing capacity are not likely to be bad counterparties. Denote the estimated coefficient on RB^s_{it} when omitting CP^s_{it} by $\hat{\eta}$, and the true parameter value by η. Classical OLS theory says if $\gamma < 0$ and $corr(RB^s_{it}, CP^s_{it}) < 0$, then $\eta < \hat{\eta}$. In other words, the true coefficient on RB^s_{it} is even more negative, which suggests sellers with low risk bearing capacity charge even higher

premiums than the estimates I have reported.

7 Conclusion

7.1 Why Might Capital Be Slow Moving?

A basic message of my paper is that limited risk bearing capital for sellers of CDS protection strongly affects prices — losses at the CDS desks of sellers are followed by subsequent increases in default risk premiums. In this subsection, I discuss why capital losses may impact risk bearing capacity, or in other words why capital in CDS markets appears "slow-moving" in the sense of Mitchell, Pedersen, and Pulvino (2007) or Duffie and Strulovici (2012). A theoretical treatment is given by Acharya, Shin, and Yorulmazer (2013), who show that slow moving capital arises naturally when agents trade off holding capital for future arbitrage profits against their current investment opportunities.

One channel for losses to impact risk aversion is psychological. When traders have lost significant amounts of money, they may become persistently less aggressive in their future trading strategies.[44] Moreover, there are institutional details that suggest risk bearing capacity and capital losses are tightly linked.

Risk Limits

Risk management teams of financial intermediaries and hedge funds impose risk limits on the positions taken by traders. There are various inputs that go into deciding the size of these risk limits such as sensitivity to yield curve changes or credit spread changes (e.g. DV01).[45] These risk limits vary with time. At the onset of a trade, each desk is often given a notional size of the position they can take. If the position performs well, this limit might be expanded. On the other hand, if the position does poorly, the limit is tightened. A contraction of risk limits signifies a decline in risk bearing capacity. Still, conversations with regulators and practitioners suggest risk limits adjust slowly (e.g. monthly or quarterly) and only with large losses, so their impact on risk bearing capacity is likely to be strongest when combined with other institutional frictions.

[44] A quote from one practitioner: If a market downturn is approaching, "senior management may remind their troops to be lean, not be a hero (by giving tight bid-asks), and try to stay neutral."

[45] Recently, Atkeson, Eisfeldt, and Weill (2014a) use risk limits as a friction to complete risk sharing in a theoretical study of OTC markets. See references therein for even more background.

Value-at-Risk (VaR)

A related way in which portfolio performance relates to risk bearing capacity is through value at risk constraints, as in Adrian and Shin (2013). Theoretical treatments of VaR constraints use the wealth of the trader as a state variable. Naturally, this wealth depends on the performance of their portfolio, so as this declines, the constraint becomes tighter. From a regulatory perspective, dealers are required to hold capital against the VaR of their trading portfolios, which is itself computed based on forecasting models and historical portfolio performance. It does not seem far fetched to assume that VaR estimates are revised upward after a poor trading week or month. Furthermore, an important input to VaR-based risk management is volatility. Thus, if portfolio losses coincide with increases in expected portfolio volatility, then VaR constraints will also concurrently tighten.[46]

Funding Constaints

Perhaps the most likely way for capital losses to affect risk bearing capacity is through funding. As described in Duffie, Scheicher, and Vuillemey (2014), swap desks arrange in advance with their treasury desk for a pool of cash to be used for margin purposes. This precautionary buffer gets used up as the portfolio loses money because the desk must make variation margin payments. Furthermore, if VaR estimates also increase after large portfolio losses, then regulated entities will have to raise additional capital for regulatory purposes.

In turn, if raising additional capital on short notice is costly, then theory suggests that desk losses will translate to reduced risk taking capacity. For example, Froot and O'Connell (2008) use the notion of costly external financing to demonstrate how the internal funds of intermediaries are important for equilibrium asset prices.[47] Froot and Stein (1998) also build a model of financial intermediaries that links portfolio decisions with costly external finance. In both settings, costly external finance and depleted capital make intermediaries more risk averse with respect to their internal wealth levels.[48]

Supporting Evidence

[46] This claim is supported by a great deal of empirical evidence that documents a negative correlation between volatility and returns in equity markets.

[47] Interestingly, they also find that the market share of intermediaries will be increasing in the cost of external funding, due to increasing returns to scale. This theoretical result seems appropriate for CDS markets as well.

[48] Furthermore, the issue of how funding costs affect derivatives valuations has been controversial post-crisis, with most of the debate centered around "funding valuation adjustments (FVA)" for valuing credit derivatives. Hull and White (2014) provide a detailed overview of the potential issues.

I now present some supporting evidence of a link between capital losses and risk bearing capacity, though a complete investigation is outside the scope of this paper. For illustration, suppose that the funding friction channel is what connects losses and risk bearing capacity. Standard economic logic says that obtaining new funding on short notice must be costly, otherwise the desk would simply replenish the capital pool with no other consequences. So for losses to affect risk bearing capacity, the CDS desk must be losing money precisely when obtaining new funding is difficult.

To flesh this idea out empirically, I compute the correlation between changes in the mark-to-market value of a counterparty c's CDS portfolio and the growth rate of c's market leverage.[49] I restrict myself to the subset of dealers in my sample because I can only compute leverage measures for these counterparties.

Presumably, dealers with high measured leverage face tighter borrowing constraints. If the funding friction hypothesis is correct, then negative changes in the CDS portfolio depress risk bearing capacity because they are also accompanied by increases in external funding constraints for the dealer (i.e. leverage).[50] Furthermore, this link should be the tightest when the desk has already experienced significant losses. As such, I also compute conditional correlations for each dealer, where I condition on whether mark-to-market values are in their 6-month 10 percent tail.[51] Figure 7 illustrates the results, with each pair of bars in the plot corresponding to a single dealer.

The leftmost counterparties in Figure 7 are some of the largest net buyers on average for the entire market. For most of these dealers, changes in CDS portfolio value are positively correlated with changes in leverage; when the CDS desk is losing money, the firm is also in deleveraging. Thus, we would not expect a decrease in risk bearing capacity to result from CDS desk losses, given that the dealer is less constrained overall. This intuition is consistent with my finding that losses for mega-buyers do not move CDS spreads.

The important portion of Figure 7 is the three rightmost dealers. These dealers are consistently three of the largest net overall sellers of protection in the market, with their net position dwarfing the other dealers who are also sellers. Unconditionally, changes in portfolio value for these three mega-sellers are about 50 percent negatively correlated with changes in

[49] I measure leverage as the ratio of book value of debt to market value of equity.

[50] One might worry that leverage is endogenous and does not accurately reflect a tightening of borrowing conditions for a dealer. I have rerun my analysis using other measures of financial distress like inverse-equity volatility (Atkeson, Eisfeldt, and Weill (2014b)) and SRISK (Brownlees and Engle (2014)). The primary conclusions are robust to these alternatives.

[51] I compute rolling 6-month changes in the CDS portfolio value, then use the 10 percent quantile as the conditioning event.

Figure 7: Correlation Between CDS Portfolio Changes and Market Leverage

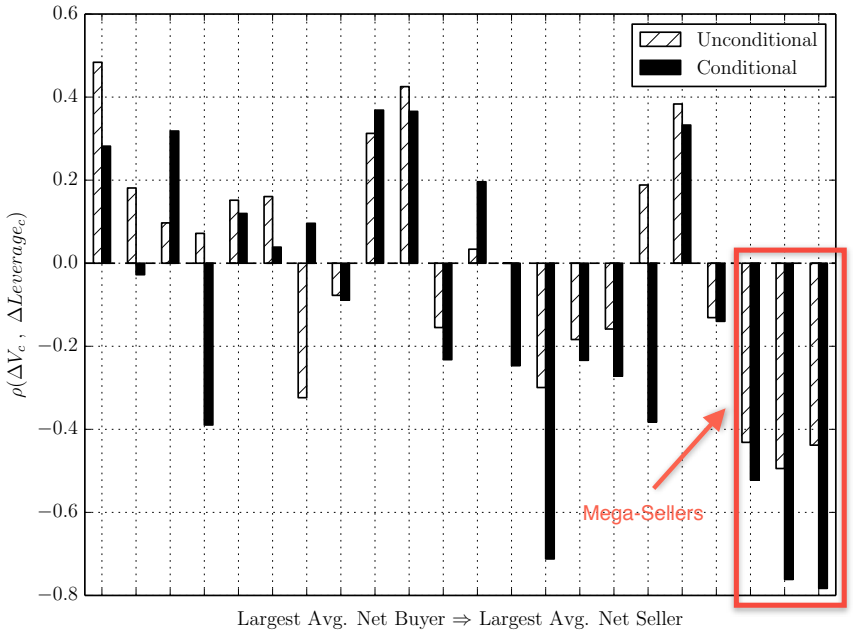

Notes: The figure computes, for each counterparty c who is a dealer, the correlation between log-changes in market leverage, $\Delta Leverage_c$, and the change in the mark-to-market dollar value of the portfolio, ΔV_c. In addition, I compute the conditional correlation between these two series. The conditioning event is that the trailing 6-month change in $V_{c,t}$ is in its 10 percent tail. On the x-axis, each pair of bars corresponds to a single dealer. Dealers are ranked based on their average net exposure, which for simplicity is computed by summing across all reference entities and counterparties.

leverage. In other words, when the desk loses additional money, the dealer's overall leverage is concurrently increasing. This broadly supports the notion that replenishing the desk's variation margin pool is costly because losses occur when new funding is relatively difficult to obtain, or, new capital is slow moving.

The correlations jump substantially when I condition on poor portfolio performance. When the CDS desks at these three dealers are doing particularly badly, changes in dollar portfolio value are nearly 80 percent negatively correlated with changes in leverage. This is exactly what we should see in the data if the funding frictions channel is correct. Of course, this evidence is observationally equivalent to both a risk limit and VaR story. It is entirely possible (and highly probable) that as a dealer's leverage is rising, it simultaneously contracts its risk limits and updates its VaR, which may itself require obtaining additional funds. In the present context, the exact avenue through which CDS losses depress risk bearing capacity is a matter of interpretation. In reality, all channels are probably at play.

The size of the CDS portfolios for these mega-sellers is also quite large. For confidentiality

reasons, I cannot report detailed descriptive statistics of portfolio size. Still, to give a sense of size, for each dealer I compute the ratio of the absolute value of their CDS portfolio value to their equity. Across the three dealers and through time, the average CDS-to-equity ratio is 5 percent, but at some points in the sample it reaches nearly 12 percent. These magnitudes are sizable given the wide range of activities and asset holdings of these institutions. It is not hard to argue that fluctuations in the value of the CDS desk for these mega-sellers also have a substantial impact on their overall wealth.

7.2 Final Remarks

The major contribution of my paper is twofold: first, I show that capital fluctuations for sellers of CDS protection explain a significant amount of variation in default risk premiums. Intuitively, when seller capital declines, risk premiums rise. This finding is consistent with many theoretical models of limited investment capital, though is particularly striking due to the fact that a handful of CDS sellers account for most of the market. Second, I show why the distribution of capital, as opposed to only the total amount of capital, matters for pricing. When sellers become more concentrated, the market becomes more fragile through an increase in the volatility of risk premiums.

These two findings have important policy implications. For one, the fact that the level of risk bearing capital in a market impacts risk premiums provides a potential rationale for outside capital injections (e.g. bailouts) when sellers are under-capitalized. Still, this argument is obviously subject to the usual set of moral hazard objections (as in the case of too-big-to-fail subsidies). Even in the presence of moral hazard issues, a deeper point is that the distribution of capital injections can also have a substantial impact on pricing. This follows directly from the fact that concentration leads to higher volatility. In this sense, it is unlikely that every bailout is equal — where capital injections flow is also important.

The stylized facts I document about CDS markets also raise potential policy red flags in terms of financial stability and market design. For instance, consider the highly controversial government bailout of AIG, who at the time was a large seller of CDS protection on the U.S. housing market. One of the cited reasons for the bailout in the popular press was that a failure of AIG could have substantial trickle down effects via counterparty risk. As alluded to in the introduction, much of the post-crisis regulation has aimed to reduce counterparty risk and increase transparency in derivatives markets though central clearing. Although these mechanisms are a useful step, the results of this paper highlight that the existence of

mega-sellers has persisted in the aftermath of the crisis.

One way to reduce the concentration inherent to CDS is for the central clearing party (CCP) to charge, in addition to existing margin requirements, a "concentration margin" that would penalize counterparties for outsized market shares. Of course, this would require the CCP to clear all traded swaps so that it could accurately compute market share through single name and index positions. With that said, it is not obvious what the optimal allocation of CDS should be from a welfare-maximizing perspective. At a minimum, more research is necessary to properly assess welfare implications, and in particular, whether new regulation has done enough to ensure a failure of one of these mega-sellers would not have a significant impact on financial stability. The map of true economic exposures (via CDS) created in this paper is a useful tool for this task.

Another important aspect of my stylized facts is that selling of CDS has partly moved from dealers to hedge funds and asset managers. A likely explanation for this pattern is that new regulation has made it less profitable (or even possible) for dealers to ultimately bear credit risk via CDS. Still, it is important to consider whether moving this function to the largely unregulated sphere of hedge funds and asset managers is optimal from a financial stability perspective.

Regarding CDS market design, the facts uncovered in this paper beg the question: why is there such limited participation in CDS markets? This question is particularly important given I have shown that changes in the risk bearing capacity of a small set of sellers in the market can move CDS spreads by a substantial amount. The next step is to understand how and why the concentrated market structure of CDS arises endogenously.

References

Viral Acharya, Stephen Schaefer, and Yili Zhang. Liquidity risk and correlation risk: A clinical study of the general motors and ford downgrade of 2005. *Forthcoming: Quarterly Review of Finance*, 2014.

Viral V. Acharya, Hyun Song Shin, and Tanju Yorulmazer. A theory of arbitrage capital. *Review of Corporate Finance Studies*, 2013.

Tobias Adrian and Nina Boyarchenko. Intermediary leverage cycles and financial stability. *Federal Reserve Staff Report*, (567), 2013.

Tobias Adrian and Hyun Song Shin. Procyclical leverage and value-at-risk. *Review of Financial Studies*, 2013.

Tobias Adrian, Paolo Colla, and Hyun Song Shin. Which financial frictions? Parsing the evidence from the financial crisis of 2007-09. Staff Reports 528, Federal Reserve Bank of New York, 2011.

Tobias Adrian, Erkko Etula, and Tyler Muir. Financial intermediaries and the cross-section of asset returns. *The Journal of Finance*, page Forthcoming, 2014.

Miguel Anton and Chistopher Polk. Connected stocks. *The Journal of Finance*, 69(3):1099–1127, 2014. ISSN 1540-6261.

Navneet Arora, Priyank Gandhi, and Francis A. Longstaff. Counterparty credit risk and the credit default swap market. *Journal of Financial Economics*, 103(2):280 – 293, 2012. URL http://www.sciencedirect.com/science/article/pii/S0304405X11002327.

Andrew Atkeson, Andrea Eisfeldt, and Pierre-Olivier Weill. The market for otc derivatives. *Working Paper, url: https://sites.google.com/site/andrealeisfeldt/*, 2013.

Andrew Atkeson, Andrea Eisfeldt, and Pierre-Olivier Weill. Entry and exit in otc derivatives markets. *Working Paper, url: https://sites.google.com/site/andrealeisfeldt/*, 2014a.

Andrew Atkeson, Andrea Eisfeldt, and Pierre-Olivier Weill. Measuring the financial soundness of u.s. firms, 1926-2012. *Working Paper, url: https://sites.google.com/site/andrealeisfeldt/*, 2014b.

Robert L. Axtell. Zipf distribution of u.s. firm sizes. *Science*, 293(5536):1818–1820, 2001. doi: 10.1126/science.1062081. URL http://www.sciencemag.org/content/293/5536/1818.abstract.

Jennie Bai and Pierre Collin-Dufresne. The cds-bond basis. *SSRN: http://ssrn.com/abstract=2024531*, 2013.

Jennie Bai and Liuren Wu. Anchoring credit default swap spreads to firm fundamentals. *SSRN: http://ssrn.com/abstract=2020841*, 2012.

Jack Bao and Jun Pan. Bond illiquidity and excess volatility. *Review of Financial Studies*, 2013. doi: 10.1093/rfs/hht037. URL http://rfs.oxfordjournals.org/content/early/2013/07/04/rfs.hht037.abstract.

Antje Berndt, Rohan Douglas, Darrell Duffie, Mark Ferguson, and David Schranz. Measuring default risk premia from default swap rates and edfs. *SSRN: http://ssrn.com/abstract=556080*, 2008.

Christian T. Brownlees and Robert F. Engle. Volatility, correlation and tails for systemic risk measurement. *Working Paper*, SSRN: http://ssrn.com/abstract=1611229, 2014.

Markus K Brunnermeier and Yuliy Sannikov. A macroeconomic model with a financial sector. *The American Economic Review*, 104(2):379–421, 2014.

John Y. Campbell and Robert J. Shiller. Stock prices, earnings, and expected dividends. *The Journal of Finance*, 43(3):661–676, 1988. ISSN 1540-6261. doi: 10.1111/j.1540-6261.1988.tb04598.x. URL http://dx.doi.org/10.1111/j.1540-6261.1988.tb04598.x.

Peter Carr and Liuren Wu. A simple robust link between american puts and credit protection. *Review of Financial Studies*, 24(2):473–505, 2011. doi: 10.1093/rfs/hhq129. URL http://rfs.oxfordjournals.org/content/24/2/473.abstract.

Hui Chen, Rui Cui, Zhiguo He, and Konstantin Milbradt. Quantifying liquidity and default risks of corporate bonds over the business cycle. *Working Paper, url: http://www.mit.edu/ huichen/bondliquidity.pdf*, 2014a.

Hui Chen, Scott Joslin, and Sophie X. Ni. Demand for Crash Insurance, Intermediary Constraints, and Risk Premia in Financial Markets. *SSRN*, 2014b. URL http://ssrn.com/abstract=2024416.

Kathryn Chen, Michael Fleming, John Jackson, Ada Li, and Asani Sarka. An analysis of cds transactions: Implications for public reporting. *NY Fed Working Paper*, (517), 2011.

Long Chen, Pierre Collin-Dufresne, and Robert S. Goldstein. On the relation between the credit spread puzzle and the equity premium puzzle. *Review of Financial Studies*, 22(9):3367–3409, 2009. doi: 10.1093/rfs/hhn078. URL http://rfs.oxfordjournals.org/content/22/9/3367.abstract.

Gabriel Chodorow-Reich. The employment effects of credit market disruptions: Firm-level evidence from the 2008-09 financial crisis. *Quarterly Journal of Economics*, 129:1–59, 2014.

In Choi. Unit root tests for panel data. *Journal of International Money and Finance*, 20(2):249–272, April 2001. URL http://ideas.repec.org/a/eee/jimfin/v20y2001i2p249-272.html.

Pierre Collin-Dufresne, Robert S. Goldstein, and J. Spencer Martin. The determinants of credit spread changes. *The Journal of Finance*, 56(6):2177–2207, 2001. ISSN 1540-6261. doi: 10.1111/0022-1082.00402. URL http://dx.doi.org/10.1111/0022-1082.00402.

Joshua Coval and Erik Stafford. Asset fire sales (and purchases) in equity markets. *Journal of Financial Economics*, 86(2):479–512, 2007.

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

Gregory R. Duffee. The relation between treasury yields and corporate bond yield spreads. *The Journal of Finance*, 53(6):2225–2241, 1998. ISSN 1540-6261. doi: 10.1111/0022-1082.00089. URL http://dx.doi.org/10.1111/0022-1082.00089.

D Duffie and KJ Singleton. Modeling term structures of defaultable bonds. *Review of Financial Studies*, 12(4):687–720, 1999. doi: 10.1093/rfs/12.4.687. URL http://rfs.oxfordjournals.org/content/12/4/687.abstract.

Darrell Duffie. Credit swap valuation. *Financial Analysts Journal*, pages 73–87, 1999.

Darrell Duffie. Presidential address: Asset price dynamics with slow-moving capital. *The Journal of Finance*, 65(4):1237–1267, 2010.

Darrell Duffie and Bruno Strulovici. Capital mobility and asset pricing. *Econometrica*, 80(6): 2469–2509, 2012.

Darrell Duffie, Ada Li, and Theo Lubke. Policy perspectives on otc derivatives market infrastructure. *Federal Reserve Staff Report*, (424), 2011.

Darrell Duffie, Martin Scheicher, and Guillaume Vuillemey. Central clearing and collateral demand. *Forthcoming: Journal of Financial Economics*, 2014.

Edwin J. Elton, Martin J. Gruber, Deepak Agrawal, and Christopher Mann. Explaining the rate spread on corporate bonds. *The Journal of Finance*, 56(1):247–277, 2001. ISSN 1540-6261. doi: 10.1111/0022-1082.00324. URL http://dx.doi.org/10.1111/0022-1082.00324.

Jan Ericsson, Kris Jacobs, and Rodolfo Oviedo. The determinants of credit default swap premia. *Journal of Financial and Quantitative Analysis*, 44(01):109–132, 2009.

Kenneth Froot and Paul O'Connell. On the pricing of intermediated risks: Theory and application to catastrophe reinsurance. *Special Issue on Dynamics of Insurance Markets: Structure, Conduct, and Performance in the 21st Century Journal of Banking and Finance*, 32(1):69–85, 2008.

Kenneth Froot, John Arabadjis, Sonya Cates, and Stephen Lawrence. How institutional investors frame their losses: Evidence on dynamic loss aversion from currency portfolios. *Journal of Portfolio Management*, 38(2):60–68, 2011.

Kenneth A. Froot and Paul G. J. O'Connell. The Pricing of U.S. Catastrophe Reinsurance. pages 195–232, 1999.

Kenneth A. Froot and Jeremy C. Stein. Risk management, capital budgeting, and capital structure policy for financial institutions: an integrated approach. *Journal of Financial Economics*, 47(1): 55–82, 1998.

Kenneth A. Froot, David S. Scharfstein, and Jeremy C. Stein. Risk management: Coordinating corporate investment and financing policies. *The Journal of Finance*, 48(5):1629–1658, 1993. ISSN 1540-6261. doi: 10.1111/j.1540-6261.1993.tb05123.x. URL http://dx.doi.org/10.1111/j.1540-6261.1993.tb05123.x.

Xavier Gabaix. The granular origins of aggregate fluctuations. *Econometrica*, 79(3):733–772, 2011. ISSN 1468-0262. doi: 10.3982/ECTA8769. URL http://dx.doi.org/10.3982/ECTA8769.

Xavier Gabaix and Rustam Ibragimov. Rank 1/2: A simple way to improve the ols estimation of tail exponents. *Journal of Business and Economic Statistics*, 29(1):24–39, 2011.

Xavier Gabaix, Arvind Krishnamurthy, and Olivier Vigneron. Limits of arbitrage: Theory and evidence from the mortgage-backed securities market. *The Journal of Finance*, 62(2):557–595, 2007.

Simon Gilchrist and Egon Zakrajsek. Credit spreads and business cycle fluctuations. *American Economic Review*, 102(4):1692–1720, 2012. doi: 10.1257/aer.102.4.1692. URL http://www.aeaweb.org/articles.php?doi=10.1257/aer.102.4.1692.

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

Richard C. Green, Burton Hollifield, and Norman Schurhoff. Financial intermediation and the costs of trading in an opaque market. *Review of Financial Studies*, 20(2):275–314, 2007.

Robin Greenwood and David Thesmar. Stock price fragility. *Journal of Financial Economics*, 102 (3):471 – 490, 2011. ISSN 0304-405X. doi: http://dx.doi.org/10.1016/j.jfineco.2011.06.003. URL http://www.sciencedirect.com/science/article/pii/S0304405X11001474.

Denis Gromb and Dimitri Vayanos. Equilibrium and welfare in markets with financially constrained arbitrageurs. *Journal of Financial Economics*, 66:361 – 407, 2002.

Samuel G. Hanson and Adi Sunderam. The growth and limits of arbitrage: Evidence from short interest. *Review of Financial Studies*, 2013.

Zhiguo He and Arvind Krishnamurthy. Intermediary asset pricing. *American Economic Review*, 103(2):732–70, 2013. doi: 10.1257/aer.103.2.732. URL http://www.aeaweb.org/articles.php?doi=10.1257/aer.103.2.732.

John Hull and Alan White. Valuing derivatives: Funding value adjustments and fair value. *Forthcoming: Financial Analysts Journal*, 2014.

Robert A. Jarrow and Stuart M. Turnbull. Pricing derivatives on financial securities subject to credit risk. *The Journal of Finance*, 50(1):53–85, 1995. ISSN 1540-6261. doi: 10.1111/j.1540-6261.1995.tb05167.x. URL http://dx.doi.org/10.1111/j.1540-6261.1995.tb05167.x.

Benjamin Junge and Anders B. Trolle. Liquidity Risk in Credit Default Swap Markets. *SSRN*, 2015. URL http://ssrn.com/abstract=2371374.

Bryan Kelly, Hanno Lustig, and Stijn Van Nieuwerburgh. Firm volatility in granular networks. *Working Paper*, SSRN: http://ssrn.com/abstract=2186197, 2014.

Asim Ijaz Khwaja and Atif Mian. Tracing the impact of bank liquidity shocks: Evidence from an emerging market. *American Economic Review*, 98(4):1413–42, 2008.

Albert S. Kyle and Wei Xiong. Contagion as a wealth effect. *The Journal of Finance*, 56(4): 1401–1440, 2001.

Francis A. Longstaff, Sanjay Mithal, and Eric Neis. Corporate yield spreads: Default risk or liquidity? new evidence from the credit default swap market. *The Journal of Finance*, 60(5): 2213–2253, 2005.

Matteo Maggiori. Financial intermediation, international risk sharing, and reserve currencies. *Unpublished manuscript, UC Berkeley*, 2011.

Robert C. Merton. On the pricing of corporate debt: The risk structure of interest rates. *The Journal of Finance*, 29(2):449–470, 1974.

Robert C. Merton. A simple model of capital market equilibrium with incomplete information. *The Journal of Finance*, 42(3):483–510, 1987.

Mark Mitchell, Lasse Heje Pedersen, and Todd Pulvino. Slow moving capital. *American Economic Review*, 97(2):215–220, 2007.

Tyler Muir. Financial Crises and Risk Premia. *SSRN*, 2014. URL http://ssrn.com/abstract=2379608.

Amrut Nashikkar and Lasse Pedersen. Corporate bond specialness. *Working Paper*, SSRN: http://ssrn.com/abstract=967545, 2007.

Yigal Newman and Michael Rierson. Illiquidity spillovers: theory and evidence from european telecom bond issuance. *Stanford University Working Paper*, 2004.

Martin Oehmke and Adam Zawadowski. The anatomy of the cds market. *SSRN: http://ssrn.com/abstract=2023108*, 2013.

Martin Oehmke and Adam Zawadowski. Synthetic or real? the equilibrium effects of credit default swaps on bond markets. *Working Paper, url: https://www0.gsb.columbia.edu/faculty/moehmke/papers/OehmkeZawadowskiCDSvsBonds.pdf*, 2014.

Or Schachar. Exposing the exposed: Intermediation capacity in the credit default swap market. *NY Fed Working Paper*, 2012.

Philipp Schnabl. The international transmission of bank liquidity shocks: Evidence from an emerging market. *The Journal of Finance*, 67(3):897–932, 2012.

Andrei Shleifer and Robert W Vishny. The limits of arbitrage. *The Journal of Finance*, 52(1): 35–55, 1997.

Dragon Tang and Hong Yan. What moves cds spreads? *Working Paper: http://www.sef.hku.hk/ yjtang/TangYan-CDSMove_21Jun2013.pdf*, 2013.

A Capital Fluctuations and the Pricing of Credit Risk: Complimentary Evidence

In this section, I provide empirical results that are intended to complement the main analysis in Section 4. I first show that capital losses for all sellers (not just mega-sellers) are followed by increases in the price of credit risk, Π_t. As a part of this analysis, I confirm that mega-seller capital losses affect the price of credit risk in the same way that total capital fluctuations of natural sellers do, a natural finding given the market share of mega-sellers. On the other hand, I also show that Π_t responds differently to a loss at an "average" seller compared to a loss at a mega-seller, thereby providing another piece of evidence that the distribution of capital is important for pricing. The last exercise of this section builds on this idea by showing that idiosyncratic dollar shocks at the largest sellers move default risk premiums.

A.1 The Price of Credit Risk and Aggregate Risk Bearing Capital

My first task is to link time-variation in $\pi_t = \log(\Pi_t)$ to variation in the risk-bearing capital of natural sellers of CDS. In Section 4.1, I established this link by using a panel of CDS spreads. The obvious advantage of a panel is that it takes advantage of both cross-sectional and time-series variation.

To support those findings, I now take a more "macro" approach to the same question. At each time t, I compute the total market value of all natural seller positions simply by summing across the market value of each natural seller's CDS portfolio, V_{ct}:

$$\mathcal{V}_t = \sum_{c \in A_t^s} V_{ct}$$

where, again, I identify natural sellers of CDS as the set of counterparties with a positive aggregate market share as defined in Section 3.2. To ensure that I do not include traders who switch between net buying and selling of protection, I only include counterparties who are above the 20th percentile in terms of aggregate market share. With some abuse of notation, I denote this set of sellers by A_t^s.

I view $\mathcal{V}_t$ as a good proxy for the net overall wealth of natural sellers because, again, it is difficult to directly hedge a sold CDS position.[52] The following regression then tests the extent to which variation in the price of credit risk are generated by changes in aggregate risk bearing capital:

$$\Delta\pi_t = a + \rho\Delta\pi_{t-1} + \mathbf{\Lambda}^{'}\left[\Delta\mathbf{X}_t\right] + \phi\left[\Delta\mathcal{V}_{t-1}\right] + e_t \tag{11}$$

where $\mathbf{X}_t$ is the same vector of macroeconomic controls used in my prior analysis. As a reminder, these are the log equity-to-price ratio for the S&P 500, VIX, TED, CFNAI, 10 year Treasury yield, 10 year minus 2 year treasury yield, and the CBOE Option Skew index. After first differencing these aforementioned variables, I also include the excess market return of the CRSP value-weighted index.[53]

In Section 4.1, I avoided reverse causality concerns by looking at whether losses outside of a reference entity's industry impact the pricing of that reference entity. Clearly, this strategy is only possible when using a panel of CDS spreads. Thus, to avoid any reverse causality issues in my current setup, I simply lag the change in seller capital regression (11). Based on these earlier empirical results and on theoretical models with slow-moving capital or limited financial intermediary risk bearing capacity, we would expect $\phi < 0$. When natural sellers lose capital, CDS premiums should subsequently rise. Table 6 collects the results of regression (11).

[52] As long as $\mathcal{V}_t$ is correlated with overall wealth, theoretical models with limited risk bearing capital suggest it should be important for determining risk premiums.

[53] This may seem redundant to including the log-change in the S&P 500 index, but I do so in order to account for higher frequency (weekly) equity movements. The earnings-to-price ratio is monthly and taken from Robert Shiller's website.

Table 6: The Price of Credit Risk and P&L of All Natural Sellers

Dep. Variable	$\Delta\pi_t$			
	(1)	(2)	(3)	(4)
$\Delta\mathcal{V}_{t-1}$		-0.007		-0.007
		(-2.8)**		(-2.7)**
$\Delta\mathcal{V}_{t-1} \times 1_{\Delta\mathcal{V}_{t-1}\geq 0}$			-0.005	
			(-1.43)	
$\Delta\mathcal{V}_{t-1} \times 1_{\Delta\mathcal{V}_{t-1}<0}$			-0.009	
			(-2.57)**	
$\Delta\mathcal{V}_{t-1} \times \Delta TED_{t-1}$				-0.003
				(-1.9)*
Macro Controls	Yes	Yes	Yes	Yes
Adj. R^2	21.0	24.0	24.0	23.0
N	190	190	190	190

Notes: This table presents results from estimating variants of the regression: $\Delta\pi_t = a + \rho\Delta\pi_{t-1} + \mathbf{\Lambda}^{'}[\Delta\mathbf{X}_t] + \phi[\Delta\mathcal{V}_{t-1}] + e_t$. $\Delta\mathcal{V}_{t-1}$ is the change in the mark-to-market value of the CDS positions for natural sellers of CDS protection from $t-1$ to t. Natural sellers are determined each period as those counterparties whose aggregate market share (defined in Section 3.2) is above the 20th percentile. ΔTED_t is the change TED spread over the same time span. All variables have been transformed to have a mean of zero and variance of one. **,* indicates coefficient is statistically different than zero at the 5 percent and 10 percent confidence level, respectively.

Column (1) of Table 6 shows the results of regression changes in π_t on its own lag and only the macroeconomic variables. The R^2 from this specification is 21 percent, or in other words, macroeconomic variables can explain about a fifth of the variation in the price of credit risk. Column (2) adds the lagged change in seller capital to the regression. The point estimate is statistically significant at conventional levels, and the negative sign confirms our intuition that capital losses result in increases in aggregate premiums. The incremental R^2 from this regression is 3 percent, which is large given the limited explanatory power by macroeconomic variables of changes in π_t.

I incorporate the potential for a non-linear relationship between capital and the price of credit risk into regression (11) in two ways. First, I allow lagged changes in aggregate seller capital to have differential effects on $\Delta\pi_t$, depending on whether capital increases or decreases. I accomplish this by interacting $\Delta\mathcal{V}_{t-1}$ with an indicator variable for whether it is positive or negative. As column (3) suggests, there is an asymmetrical effect of capital losses on the price of credit risk. In absolute value, the price of credit risk rises almost twice as much following a loss, relative to how much it drops following a capital gain.[54] Again, this is exactly what we would expect in lieu of the theoretical results from models like He and Krishnamurthy (2013).

The second way I incorporate non-linearity is by interacting lagged changes in aggregate seller capital with changes in the TED spread. Because the TED spread is often used as a proxy for the funding conditions of the banking sector, this interaction term captures situations where natural sellers lose capital and funding conditions tighten. Following with intuition, column (4) shows that capital losses have their strongest effect on the price of credit risk when funding conditions are simultaneously tightening. For instance, if a one standard deviation increase in the TED spread is accompanied by a one standard deviation deduction in

[54]These relative magnitudes are larger than those found in my panel analysis in Section 4.1. This is plausibly attributed to the fact that Section 4.1 examines contemporaneous changes in CDS spreads and losses, whereas the present analysis uses lags.

seller capital, the price of risk rises by about one percent in the following week.[55]

A Simple Variance Decomposition

The goal of the previous set of regressions was to establish that capital losses to natural sellers of protection are followed by increases in default risk premiums. The results from that exercise also suggest that *shocks* to the price of credit risk are likely driven by either shocks to macroeconomic fundamentals or shocks to the risk bearing capital of natural sellers. To get a better sense of what is behind variability in shocks to Π_t, I use a forecast error variance decomposition from the following first order unrestricted vector autoregression (VAR):

$$\begin{aligned} \mathbf{y}_t &= \mathbf{c} + \mathbf{\Phi}^{'} \mathbf{y}_{t-1} + \mathbf{e}_t \\ \mathbf{y}_t &:= \begin{bmatrix} \mathbf{X}_t & \mathcal{V}_t & \Pi_t \end{bmatrix}^{'} \end{aligned}$$

where the lag structure in the VAR was chosen by standard information criterion. I recursively order the variables in the VAR under the assumption that a shock to natural seller capital levels affects the price of credit risk with a lag, while macroeconomic variables can react contemporaneously to such a shock. I have further verified that the estimated VAR coefficient matrix $\mathbf{\Phi}$ is stable in the sense that all of its eigenvalues have a modulus less than one.

Figure 8 displays the amount of variation in Π_t explained by an orthogonalized one standard deviation shock to natural seller capital. An orthogonal innovation to natural seller capital accounts for about 20 percent of the variation of shocks to the price of credit risk. In other words, about one-fifth of unexpected changes in the price of credit risk can be attributed to shocks to the capital of natural CDS sellers. These results complement my earlier finding in Section 4.1 that fluctuations in mega-seller capital explain about one-ninth of the fluctuations in contemporaneous CDS spread changes, with the difference likely due to the fact that CDS spreads are (almost mechanically) noisier than Π_t.

Interpretation

One might argue that because swaps are in zero-net supply, losses to natural sellers of protection are equivalent to gains for natural buyers of protection. In turn, it is really the gains of buyers that causes an increase in aggregate CDS premiums. To see why an increase in buyer wealth may affect risk premiums, consider the following sketch of a model: existing buyers and sellers of CDS differ in their opinions about the likelihood of default, with buyers more pessimistic about default likelihoods than sellers. When a news shock occurs that increases default probabilities, equilibrium CDS spreads rise and wealth gets transferred from sellers to buyers due to mark-to-market changes. Because the wealth-weighted average expected default likelihood is now higher, the premium for bearing default risk subsequently rises even more.[56]

In this example, the capital of natural sellers is not the true determinant of variations in aggregate default risk premiums, but seller losses would nonetheless correspond to increases in risk premiums. In an alternative view of the world, it could be the case that the demand for protection is perfectly inelastic in price, maybe because buyers of CDS purchase protection in fixed quantity to hedge a constant portion of their underlying bond portfolio. In this case, a decrease in seller wealth limits their willingness or ability to bear additional risk, thereby increasing the price of credit risk.

There are few reasons why I believe my results support the notion that seller losses, and not buyer gains, economically determine CDS premiums. First, and most importantly, the identification exercises I explore in other portions of the paper support this claim. Second, if buyer wealth is really moving risk premiums,

[55] The standard deviation of log-changes in the price of credit risk is about 2 percent. So in other words, a one standard deviation event in capital losses and TED spread changes this week produces an effect on the price of credit risk on the order of 50 percent of a standard deviation next week.

[56] I thank Xavier Gabaix for a very useful discussion of this possibility.

Figure 8: Forecast Error Variance Decomposition of Seller Capital Shock

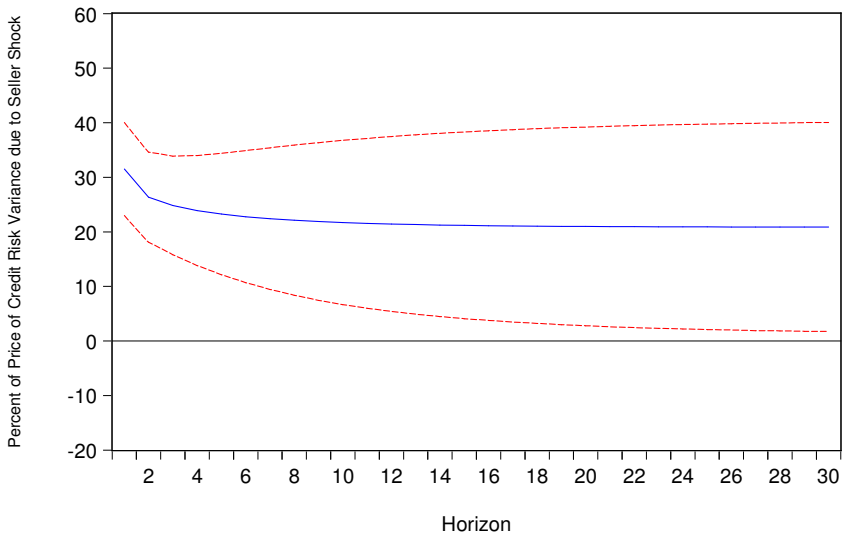

Notes: This figure plots portion of the forecast error variance of the price of credit risk Π_t explained by shocks to natural seller capital.The VAR was orthogonalized using the Cholesky decomposition with the following ordering: CFNAI index, 10Y treasury yield, 10Y-2Y treasury yield, VIX, TED, CBOE Option Skew index, earnings-to-price ratio, CRSP Value-Weighted Excess Market Return, natural seller capital $\mathcal{V}_t$, and the price of credit risk Π_t.The system was estimated using weekly data from February 2010 to June 2014. The variance decomposition is robust to alternative orderings of the macroeconomic variables.

then we would not observe an asymmetric response of risk premiums to natural seller losses. Moreover, there is no reason to believe that buyer gains should have a stronger affect on the price of risk when funding conditions for banks are worsening, as is the case in column (4) of Table 6. Finally, because I am unable to observe the joint bond-CDS portfolio of buyers, it is not entirely clear that a gain in CDS actually increases buyer wealth. As a result, I adopt the interpretation that changes in the capital of natural CDS sellers drive dynamics in the price of credit risk.

A.2 Concentration and Aggregate Risk Bearing Capital

To confirm that aggregate risk bearing capacity is largely determined by mega-sellers, in each week t, I mirror regression (11) by using changes in the capital of the mega-buyers and mega-sellers to explain subsequent changes in the price of credit risk:

$$\Delta\pi_t = a + \rho\Delta\pi_t + \mathbf{\Lambda}^{'}\left[\Delta\mathbf{X}_t\right] + \phi_S\left[\Delta\mathcal{V}_{t-1}^{MS}\right] + \varphi_B\left[\Delta\mathcal{V}_{t-1}^{MB}\right] + e_t \tag{12}$$

where $\Delta\mathcal{V}_t^{MS}$ is the change in the mark-to-market value of the entire CDS portfolio for mega-sellers from time $t-1$ to t.[57] $\Delta\mathcal{V}_t^{MB}$ is the same measure for mega-buyers. I am able to include the capital changes for the top sellers and buyers in this regression because they are not collinear. This would not be the case if these counterparties were the only sellers and buyers in the market. Table 7 presents the results of this

[57]Earlier, I used a similar version of this variable that was reference entity specific because it only included positions on reference entities outside of a specific industry.

Table 7: The Price of Credit Risk and P&L of the Largest Sellers

Dep. Variable	$\Delta\pi_t$			
	(1)	(2)	(3)	(4)
$\Delta\mathcal{V}_{t-1}^{MB}$		0.001		0.001
		(0.5)		(0.52)
$\Delta\mathcal{V}_{t-1}^{MS}$		-0.008		-0.008
		(-3.0)**		(-3.1)**
$\Delta\mathcal{V}_{t-1}^{AB}$			0.002	0.002
			(1.0)	(1.1)
$\Delta\mathcal{V}_{t-1}^{AS}$			-0.00	-0.002
			(-0.19)	(-0.96)
Macro Controls	Yes	Yes	Yes	Yes
Adj. R^2	21.0	24.1	20.7	24.1
N	190	190	190	190

Notes: This table presents results from estimating variants of the regression: $\Delta\pi_t = a + \rho\Delta\pi_t + \mathbf{\Lambda}' [\Delta\mathbf{X}_t] + \phi_S [\Delta\mathcal{V}_{t-1}^{MS}] + \varphi_B [\Delta\mathcal{V}_{t-1}^{MB}] + e_t$. $\Delta\mathcal{V}_t^{MS}$ is the change in the mark-to-market value of the entire CDS portfolio for mega-sellers from time $t-1$ to t, and $\Delta\mathcal{V}_t^{MB}$ is the same measure for mega-buyers. In columns (3) and (4), I add the change in CDS portfolio value for the "average" buyer and seller in the market, denoted by $\Delta\mathcal{V}_{t-1}^{AB}$ and $\Delta\mathcal{V}_{t-1}^{AS}$, respectively. The average seller (buyer) is determined each period as the group of sellers (buyers) whose market share is in the 45-55th percentile. All $\mathcal{V}$ variables have been transformed to have a mean of zero and variance of one. **,* indicates coefficient is statistically different than zero at the 5 percent and 10 percent confidence level, respectively.

regression.

There are at least two interesting patterns that emerge from Table 7. First, the point estimate on changes in the market value of mega-buyers' CDS in column (2) is not distinguishable from zero, indicating that this set of buyers does not drive risk price dynamics. On the other hand, a one standard deviation deduction in the capital of mega-sellers result in a 0.8 percent increase in the future price of credit risk. Moreover, the magnitude of this effect is about the same as a one standard deviation decrease in the aggregate capital of all natural sellers (compare to column (2), Table 6). This result is not surprising, given that five sellers of CDS account for almost fifty percent of the entire market. However, this is also my point about fragility: the capital level of just five protection sellers plays an outsized role in CDS pricing.

The second observation from Table 7 is that a dollar loss at a mega-seller has a different effect than a dollar loss at an "average" seller. To make this clear, in each period I define the median sellers (buyers) as counterparties whose market share of aggregate selling (buying) is in the 45-55 percentile.[58] I use medians instead of averages to avoid the undue influence of extremely large and small market shares. I denote the change in CDS portfolio value for the group of "average" sellers (buyers) by $\Delta\mathcal{V}_t^{AS}$ ($\Delta\mathcal{V}_t^{AB}$). Column (3) then includes lags of these variables in regression (12). As the point estimates suggest, there is no discernible impact on pricing from fluctuations in average seller or average buyer capital. If the distribution of capital was irrelevant and only the aggregate amount of risk bearing capital mattered for pricing, then we would not observe these results. Column (4) reinforces this idea by including in the regression changes in capital for the mega-players and the average players. A comparison of the estimated coefficient on mega-seller versus average seller capital changes indicates that the effect of a dollar loss on pricing is highly dependent on where

[58] I use a range, as opposed to a single counterparty, to avoid potential issues stemming from the actual median seller changing identities from week to week.

this loss occurs.

A.3 Idiosyncratic Losses Cause Spread Movements

A recurring theme of this paper is that the distribution of capital losses (as opposed to just the total amount of capital) for natural sellers is important because it can generate price fragility. In an environment with very concentrated sellers, idiosyncratic capital shocks to one seller translate to aggregate capital shocks. Thus, poor performance of an isolated mega-seller can still have large pricing impacts in the aggregate. Indeed, the 2012 JP Morgan London Whale is certainly one example of how the entire CDS market can move based on the actions of a single "rogue" trader. In this subsection, I formalize the logic of the "London Whale" using a slightly more complex econometric methodology that first identifies idiosyncratic shocks to large sellers and then traces out how these idiosyncratic shocks impact prices.

To start, I regress each counterparty's portfolio changes on a set of systematic factors, denoted by $\mathbf{F}_t$, and save the regression residuals:

$$\Delta V_{ct} = a_c + \boldsymbol{\beta}_c^{'}\mathbf{F}_t + \xi_{ct} \tag{13}$$

The key part of this step is choosing the appropriate set of factors, $\mathbf{F}_t$. One approach would be to use economic factors such as the excess equity market return or changes in changes in the treasury yield curve. Instead of taking a stance on the important factors that drive portfolio changes, I use principal components to choose the factors. Specifically, I extract the first five principal components from a panel of log-CDS spread changes. I use five factors because that is the minimum number needed to capture 80 percent of spread variation, but my results are not sensitive to this choice.[59]

I treat the residuals, ξ_{ct}, from this regression as idiosyncratic capital shocks to c's CDS portfolio. I do so because ξ_{ct} is the portion of c's capital change (from CDS) that cannot be explained by exposure to the systematic factors, $\mathbf{F}_t$. Next, for each reference entity r I build "buyer and seller granular residuals", defined as:

$$\begin{aligned} \Psi_{rt}^{s} &:= \sum_{c \in S_{t-1}(r)} \omega_{r,c,t-1}^{s} \xi_{ct} \\ \Psi_{rt}^{b} &:= \sum_{c \in B_{t-1}(r)} \omega_{r,c,t-1}^{b} \xi_{ct} \end{aligned} \tag{14}$$

where $S_{t-1}(r)$ and $B_{t-1}(r)$ are the set of r's sellers and buyers, respectively, at time $t-1$. Similarly, $\omega_{r,c,t-1}^{s}$ is counterparty c's share of selling in r at time $t-1$. $\omega_{r,c,t-1}^{b}$ is defined analogously for buyers.[60] In words, Ψ_{rt}^{s} is a share-weighted average of the idiosyncratic dollar shocks for r's net sellers. When r's important sellers are hit with large idiosyncratic dollar shocks, then Ψ_{rt}^{s} will also be large. Ψ_{rt}^{b} carries a similar intuition, but for the important buyers of reference entity r. I call the Ψ variables seller and buyer granular residuals in accordance with Gabaix (2011).[61]

[59] If anything, I am being too stringent in my choice of factors, given that shocks to large sellers could themselves be an important factor moving spreads.

[60] Formally, $\omega_{r,c,t-1}^{b} = -NS(c,r,t-1)/NO(r,t-1)$ conditional on $NS(c,r,t-1) < 0$. The negative sign makes buyers have positive weights, since $NS < 0$ indicates net buying.

[61] The reader may notice that the granular buyer and seller granular residuals in this section are constructed using the biggest sellers *of a particular reference entity* r. This construction is in slight contrast to my analysis of how losses transmit across mega-players' portfolios, where I focus on the top five sellers and buyers *in the aggregate*. The reason I opt for the two different approaches is for identification. In the previous section, I am able to control for unobservable macroeconomic variables (via a time fixed effect) because I have cross-sectional heterogeneity in my loss measures, the reason being that I am use at losses outside of r's industry. The easiest way to generate cross-sectional heterogeneity in the current exercise is to use r's biggest sellers/buyers. The pro is that I can include a time fixed effect to account for unobservable

Table 8: Explanatory Power of Seller and Buyer Granular Residuals

Dependent Variable	$\Delta \log(CDS_{rt})$
	(1)
Ψ^s_{rt}	-.005
	(-4.4)**
Ψ^b_{rt}	.001
	(1.4)
Firm Controls	Yes
RE FE	Yes
Macro Variables	No
Time FE	Yes
Total N	65,884
Within R^2	30.4
Cluster	(r,t)

Notes: This table presents the results of the regression: $\Delta \log(CDS_{rt}) = \alpha_r + \iota_t + \boldsymbol{\zeta}' \Delta \mathbf{Z}_{rt} + \phi_s \Psi^s_{rt} + \phi_b \Psi^b_{rt} + \varepsilon_{rt}$. Ψ^s_{rt} and Ψ^b_{rt} are reference entity r's seller and buyer granular residuals, respectively. These variables are defined formally according to Equation (15). $\mathbf{Z}_{rt}$ is a vector of reference entity specific variables: Moody's 5 year EDF and Markit's LGD. α_r and ι_t are reference entity and time fixed effects, respectively. The buyer and seller granular residuals have been standardized to have a mean of zero and variance of one. All standard errors are double-clustered by reference entity and time. **,* indicates coefficient is statistically different than zero at the 5 percent and 10 percent confidence level, respectively. Data spans March 2010 to June 2014.

In a low concentration market with many sellers, idiosyncratic shocks to any one seller should not impact total capital, and thus pricing. To determine whether whether this is true in the data, I estimate the following panel regression:

$$\Delta \log(CDS_{rt}) = \alpha_r + \iota_t + \boldsymbol{\zeta}' \Delta \mathbf{Z}_{rt} + \phi_s \Psi^s_{rt} + \phi_b \Psi^b_{rt} + \varepsilon_{rt} \tag{15}$$

where α_r is a reference entity fixed effect, ι_t is a time fixed effect, and $\mathbf{Z}_{rt}$ is a vector of reference entity controls. Again, the controls I use are Moody's expected default frequency and Markit's loss-given-default. Table 8 contains the results of this panel regression.

The results in Table 8 are supportive of the notion that the concentration of risk bearing capital is important for credit risk pricing. As column (1) shows, a one standard deviation increase in the seller granular residual leads to a 0.5 percent increase in CDS spread levels. Keep in mind that I have purged capital changes of any systematic factors, so it is not surprising that the point estimate on Ψ^s_{rt} is not particularly large, though it is still statistically significant at a five percent confidence level. The more important observation is that idiosyncratic capital losses for large sellers increase risk premiums, which is one reason that concentration in CDS creates fragility in prices.

macroeconomic variables; the con is that it is slightly inconsistent with my previous analysis.

B Appendix: Seller Capital and the CDS-Bond Basis

The conclusion drawn in the main text is losses to mega-sellers results in an increase in the price of credit risk. The price of credit risk, however, is measured from CDS spreads. An alternative explanation for my results is seller losses only push CDS spreads up, not actual bond yields. If this is the case, then losses to mega-sellers has an effect on the CDS-bond basis, not the actual price of credit risk. Loosly speaking, the CDS-bond basis is the difference between a reference entity's CDS spread and its bond yield. Standard no-arbitrage arguments like Duffie (1998) suggest that the CDS-bond basis should be equal to zero. The purpose of this appendix is to determine whether my results affect bond prices, or the basis.

B.1 Data

Z-spreads As Duffie (1998) argues, a credit default swap can be replicated by a combination of a par floating rate note issued by the reference entity and a default-free floating rate note. The reference entity's floating rate note has payments that are a fixed spread S over the default-free floating rate note. With some simplifying assumptions, buying CDS protection is equivalent to shorting the reference entity floating rate note and buying the default-free floating rate note.

In practice, though, a standard corporate bond is a fixed coupon instrument. To compare its price to a credit default swap, one must first convert the fixed payments to floating payments by appropriately layering the bond with an interest rate swap. Asset swaps are the conventional way of doing so, and the asset swap spread refers to the par spread over a benchmark (typically LIBOR) that the fixed rate payer in the asset swap receives. A closely related concept is the "Z-spread" which is uses observable bond prices and the zero-coupon yield curve in valuing the asset swap. The difference between the asset swap spread and the Z-spread is often small, and the nuances of the two are beyond the scope of this paper. For my purposes, I will compare the CDS spread of a reference entity with its Z-spread.[62] To summarize, no-arbitrage says that Z-Spread ≈CDS spread.

I obtain Z-spread data from Bloomberg. For each reference entity in my sample, I use MarkitRed to determine that reference entity's ticker. For each ticker, I then use Bloomberg to find the outstanding senior unsecured bond that is closest in maturity to 5 years. I choose 5 years to maintain consistency with the other data in the paper. The Bloomberg command to obtain the bond with maturity nearest five years is: Ticker + "GB USD SR 5Y Corp". For instance, if the company I am interested in is IBM, I obtain information on their nearest-to-five year bond using "IBM GB USD SR 5Y Corp". From within this view, I retrieve the corresponding Z-spread for each ticker using the variable "Z SPRD MID". Because Markit and Bloomberg tickers do not have a one to one mapping, I am not able to obtain Z-spread information for many of the RedIDs considered in the main text.

Bloomberg delivers Z-spreads at a daily frequency. I convert these spreads to a weekly frequency by removing non-business day observations and taking within week averages. I do so in order to smooth the otherwise noisy Z-spread data, and also to match the frequency used throughout this paper.

CDS-Bond Basis Bloomberg also provides information on the CDS-Bond basis. As argued, this is just the difference between the maturity-matched CDS spread and the Z-Spread. For each bond I consider, Bloomberg interpolates the CDS curve to determine the CDS spread to match the bond's maturity (nearest five years). I obtain the basis for each bond using the field "BLP_CDS_BASIS_MID".

B.2 Panel Regression Results

To get a visual sense of how bond implied Z-spreads and the basis move through my sample, Figure 9 plots the average Z-spread across bonds at each point in time, the average basis, and my estimate of the price

[62] This is also common practice by traders. My data on Z-spreads come from Bloomberg, who also computes the CDS-bond basis using Z-spreads.

Figure 9: Average Z-Spread, Price of Credit Risk, and CDS-Bond Basis

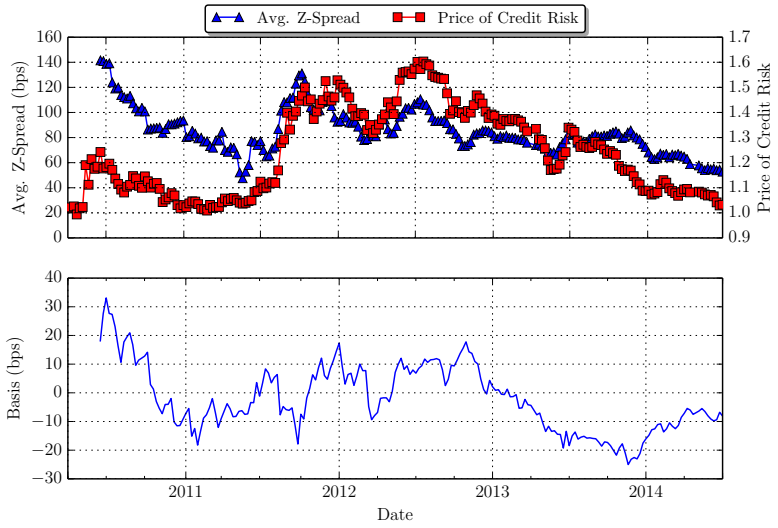

Notes: The top panel of this figures shows the average bond Z-spread and the price of credit risk implied by CDS markets. Z-spreads are obtained from Bloomberg, and are for the outstanding bond whose maturity is closes to five years. The price of credit risk derives from a panel of CDS spreads, and is estimated using the regression in Equation (22). The bottom panel shows the average CDS-bond basis, also obtained from Bloomberg.

of credit risk from CDS. The top panel of the figure reveals strong comovement between the average bond Z-spread and my price of credit risk estimate. The bottom panel of the figure suggests some comovement between the basis and price of credit risk, but also shows the average basis does not significantly depart from zero during this time period.[63] In other words, CDS spreads and bond yields moved nearly in tandem over this time period, which implies when mega-sellers lose money, CDS spreads and bond yields rise together.

To formally confirm this logic, I estimate the following sequence of panel regressions:

$$\begin{aligned} \log\left(ZSpread_{rt}\right) &= c + a_r + \phi_1 \log\left(LGD_{r,t}\right) + \phi_2 \log\left(EDF_{rt}\right) + \gamma_1 RB_t^s + \gamma_2 RB_t^b + \boldsymbol{\beta}'\mathbf{X}_t + \varepsilon_{rt} \\ Basis_{rt} &= c + a_r + \phi_1 \log\left(LGD_{r,t}\right) + \phi_2 \log\left(EDF_{rt}\right) + \gamma_1 RB_t^s + \gamma_2 RB_t^b + \boldsymbol{\beta}'\mathbf{X}_t + \varepsilon_{rt} \end{aligned}$$

where the same controls are used from Section 4.1. These include: the log of the earnings to price ratio for the aggregate stock market, the VIX index, TED spread, CFNAI index, slope of the treasury yield curve, level of the constant maturity 10 year treasury yield, and the CBOE option-implied skew index.

An important distinction with my analysis in the main text is that this regression is run in levels, not first differences. One concern is that spread levels are very persistent stochastic processes, so level regressions may suffer from the usual econometric issues that accompany near unit-root processes. Unfortunately, my sample of Z-spreads is highly uneven, meaning consecutive time-series observations are not always available. Thus, taking first differences is difficult for me to do.

I take a few steps to help mitigate the persistence of the variables in the regression. First, I cluster standard errors by reference entity and time, which should help with any autocorrelation in within-reference entity errors. Second, I do not use pure dollar losses for mega-sellers as an explanatory variable. Instead,

[63] During the financial crisis of 2007-09, Bai and Collin-Dufresne (2013) report the basis reached almost -150 basis points for some bonds.

Table 9: The Effect of Mega-Player Losses on Z-Spreads and the CDS-Bond Basis

Dependent Variable	$\log(ZSpread_{rt})$		$Basis_{rt}$	
	(1)	(2)	(3)	(4)
RB_t^s	-0.061	-0.034	-0.85	-0.29
	(-4.37)**	(-2.40)**	(-0.70)	(-0.34)
RB_t^b	0.009	0.004	1.15	0.009
	(0.89)	(0.41)	(1.35)	(0.01)
Fixed Effect	r	r	r	r
EDF and LGD	yes	yes	yes	yes
Macro Controls	no	yes	no	yes
Total N	12,485	12,328	13,620	13,451
Within R^2	10.6	13.6	1.3	6.2
Cluster	(r,t)	(r,t)	(r,t)	(r,t)

Notes: This table contains the regression results of: $Y_{rt} = a_r + \phi_1 \log(LGD_{r,t}) + \phi_2 \log(EDF_{rt}) + \gamma_1 RB_t^s + \gamma_2 RB_t^b + \boldsymbol{\beta}' \mathbf{X}_t + \varepsilon_{rt}$. Y_{rt} is either the (log) Z-spread for reference entity r or the CDS-bond basis. Z-spreads and basis are obtained from Bloomberg and correspond to the bond whose maturity is closest to five years. RB_t^j is the lagged cumulative annualized CDS portfolio return for the top five $j \in \{b, s\}$, and has been normalized to be mean zero and variance one. X_t is a set of controls: the log of the earnings to price ratio for the aggregate stock market, the VIX index, TED spread, CFNAI index, slope of the treasury yield curve, level of the constant maturity 10 year treasury yield, and the CBOE option-implied skew index. Standard errors are double clustered by reference entity and time. *, ** denote 10 and 5 percent significant, respectively.

RB^s is the annualized dollar value of mega sellers' portfolios, normalized by the total gross notional of their portfolio. Like in Section 6, this transformation yields a more stationary variable. Additionally, the CDS-bond basis is by definition stationary, as it must revert to zero by no arbitrage. Thus, when using the basis as the dependent variable, I am less worried about the aforementioned econometric issues. Table 9 collects the results from running both specifications, with and without controls.

The estimates in Table 9 confirm mega-sellers losses affect the price of credit risk, and not the CDS-bond basis. Column (1) presents the panel regression of log Z-Spreads on RB_t^s and RB_t^b, with only controls for the EDF and LGD of the underlying reference entity. These results indicate a strong and significant effect of mega-seller losses on credit spread, as a one standard deviation loss to mega-sellers results in a 6% increase in Z-Spreads. Column (2) shows the result of adding macroeconomic controls to the regression, with the effect of mega-seller losses remaining significant and large, albeit reducing by almost one half. Mega-buyer CDS losses have seemingly no effect on Z-Spreads, which is consistent with the same panel regression using log-CDS spreads as the dependent variable,

Columns (3) and (4) reveal mega-player CDS losses have no effect on the CDS-bond basis — all of the effect runs through the price of credit risk. With or without macroeconomic controls, neither RB_t^s nor RB_t^b are statistically different from zero. In lieu of the potential econometric issues with spread levels, I view the results in columns (3) and (4) as the strongest piece of evidence that capital losses impact credit risk prices and not just the basis.

We must be cautious in comparing the estimated magnitudes in columns (1) and (2) to the results in the main text. Importantly, the sample size is much smaller in the present set of regressions, mainly because there was an imperfect match between Markit tickers and Bloomberg tickers. Because the panel regression results using CDS spread changes as the dependent variable captures a much larger cross-section of reference entities, the magnitudes reported in that set of regressions — a one standard deviation loss to mega-sellers results in a roughly 2.7% CDS spread level increase — are likely more reliable.

C Appendix: Japanese Tsunami

C.1 Background and Evidence of U.S. Counterparty Exposure to Japan

On March 11, 2011, a magnitude 9.0 earthquake hit off the coast of Tohoku and was the most powerful earthquake ever recorded in Japanese history. The earthquake in turn triggered a tsunami that devastated the entire country, resulting in hundreds of billions of dollars in damages. One year after the event, the Japanese government estimated material damages could cost as much as $300 billion.[64]

A simple way to visualize the aggregate effects of the natural disaster is through the CDS spread of the entire country of Japan, which I henceforth denote by JPN. Figure 10 plots this series from March 4, 2011 to March 17, 2011. Prior to the tsunami, JPN's CDS spread was low, hovering around 80 basis points. The tsunami occurred on a Friday, and the CDS spread increased by nearly 50% on the following Monday to just over 115 basis points.

Figure 10: Five Year CDS Spread of Japan

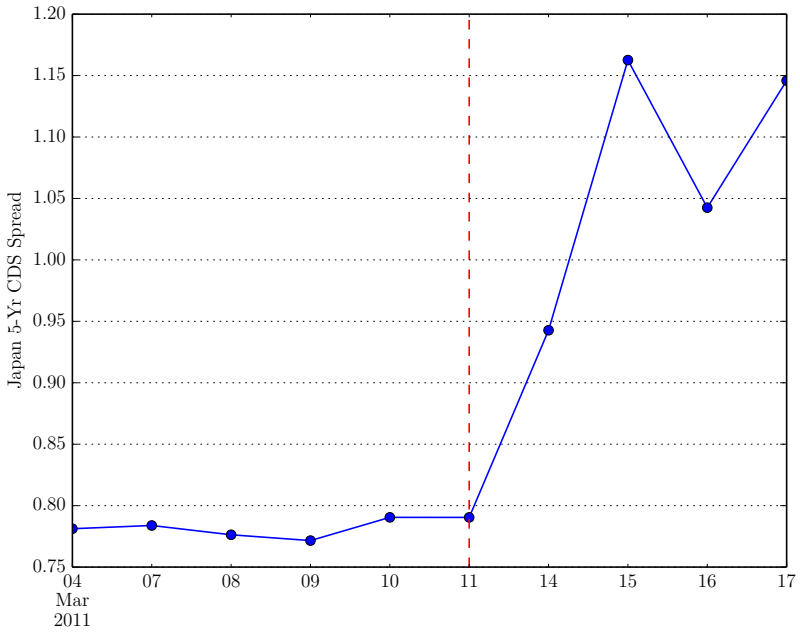

Notes: This figure plots the 5 year CDS spread for Japan from March 4, 2011 - March 17, 2011. The default event is of the clause "CR", for senior unsecured debt, and denominated in dollars.

Because of variation margin payments, a sharp rise in the credit riskiness of Japanese firms could still have a substantial impact on CDS sellers with large exposure to Japan, even without an actual default event occurring. Of course, this assumes U.S. counterparties have large exposures to Japan, which I confirm in Figure 11. To compute the exposure of a given counterparty to Japanese firms, I simply sum their net position in any reference entity who is registered in the country.[65] In the figure, the largest buyers and their

[64] Source: http://www.cnn.com/2012/02/26/world/asia/rebuilding-japan-overview/

[65] I use the country designations provided by Markit.

Figure 11: Exposures of U.S. Counterparties to Japanese Firms: 3/11/2011

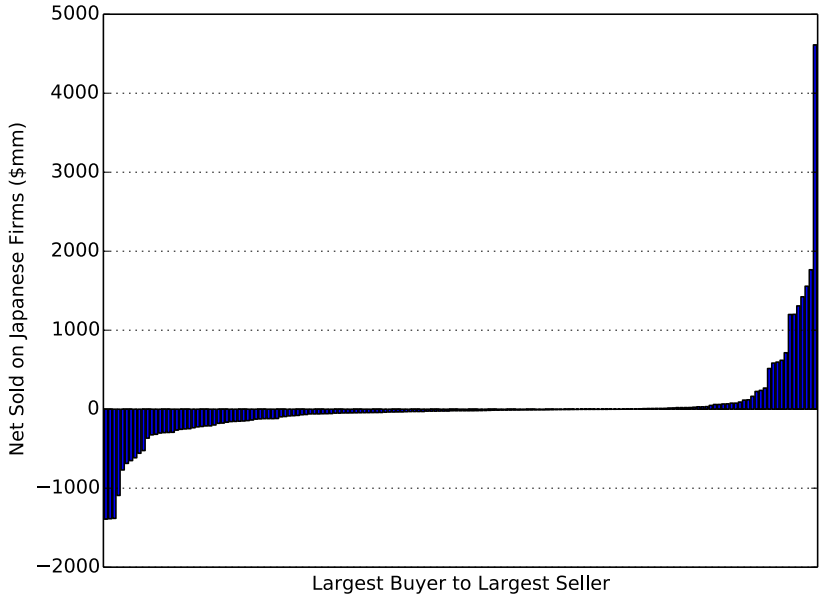

Notes: This figure shows the exposure of U.S. counterparties to Japanese firms. For each counterparty in the U.S., I compute the net amount sold on any firm whose country jurisdiction is in Japan, as classified by Markit. Positive values indicate a net seller overall, negative values indicate a net buyer.

associated positions are found on the left. The largest sellers and the size of their positions are on the right.

At the time of the tsunami, the largest seller of protection had an outstanding exposure of over $4 bn to Japanese firms. Conversely, the largest buyer of protection had bought over $1b in credit protection. Furthermore, these exposures represent a nontrivial portion of each counterparty's overall portfolio. To roughly quantify this statement, I compute the ratio of the net sold on Japanese firms to the absolute value of all net sold positions in each counterparty's portfolio. For the largest sellers of Japense firms, this ratio ranges anywhere from 4 to 106 percent.[66] Similarly, for the largest five buyers of CDS on Japanese firms, this ratio is between -2 to -90 percent. Clearly, the absolute and relative size of these positions are large, and we would therefore expect a negative shock to Japan to have a nontrivial impact on the overall portfolios of these counterparties.

To further argue that the tsunami had a non-trivial affect on risk bearing capacity, I examine the dollar gains or losses induced by the tsunami. Specifically, for each counterparty, I compute the mark-to-market loss of all Japanese exposures between March 11, 2011 and March 18, 2011. Figure 12 depicts this visually, with the largest gains on the left and the largest losses on the right. By convention, I use a negative sign to denote gains.

In the week following the tsunami, the seller who had sold the most protection on Japan lost nearly $50 million dollars due to this exposure. To put this in perspective, in the quarter leading up to the tsunami, the standard deviation of *total* portfolio value changes for this seller was $150 million. Combined, the five largest losers from the Japanese tsunami experienced a total outflow of $110 million in the week following the event.

[66]The ratio can exceed 1 in absolute value because the counterparty may, for instance, sell a lot on Japan and then be a buyer of protection in other names.

Figure 12: Mark-to-Market Losses of U.S. Counterparties from Japanese Exposure: 3/11/2011 to 3/18/2011

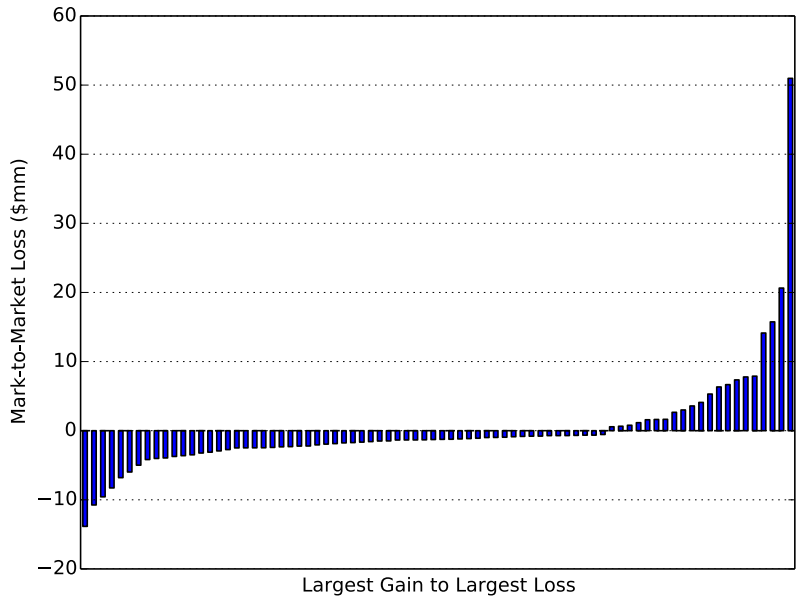

Notes: This figure mark-to-market losses of U.S. counterparties to Japanese firms in the week following the 2011 tsunami. Japanese reference entities include any firm whose country jurisdiction is in Japan, as classified by Markit. Positive values indicate a counterparty lost money due to Japanese exposure, and negative values indicate a counterparty gained money.

While these magnitudes are not huge compared to the size of large banks or the AUM of large hedge funds, a completely unexpected loss of $110 million in one week is certainly not trivial. As I discuss later, there was also substantial uncertainty in the weeks following the tsunami, so it was unclear if additional capital losses would ensue.

C.2 Data Description

Before proceeding further, let me briefly describe the data I use for the remainder of this section. On March 11, 2011, I determine the largest 500 reference entities in the U.S. CDS market as measured by net notional outstanding. I then obtain a time-series of 5-year CDS spreads for each reference entity from Markit. In particular, the CDS spreads I examine are for senior unsecured debt, with a "MR" restructuring clause. Depending on the application, I match each reference entity to Moody's EDF database and CRSP.

C.3 Potential Identification Issues

To develop the identification challenge I face, suppose the change in U.S. spreads following the tsunami takes a linear form:[67]

$$\Delta \log(CDS_{r,1}) = a + \phi_1 \Gamma_{S,r} + \phi_2 \Gamma_{B,r} + \boldsymbol{\beta}' \boldsymbol{X}_r + \gamma U_r + \varepsilon_r \tag{16}$$

[67]The following exposition could be easily generalized to more general response functions and GMM arguments.

Figure 13: Correlation with Japan Across Deciles of $\Gamma_{S,r}$ and $\Gamma_{B,r}$

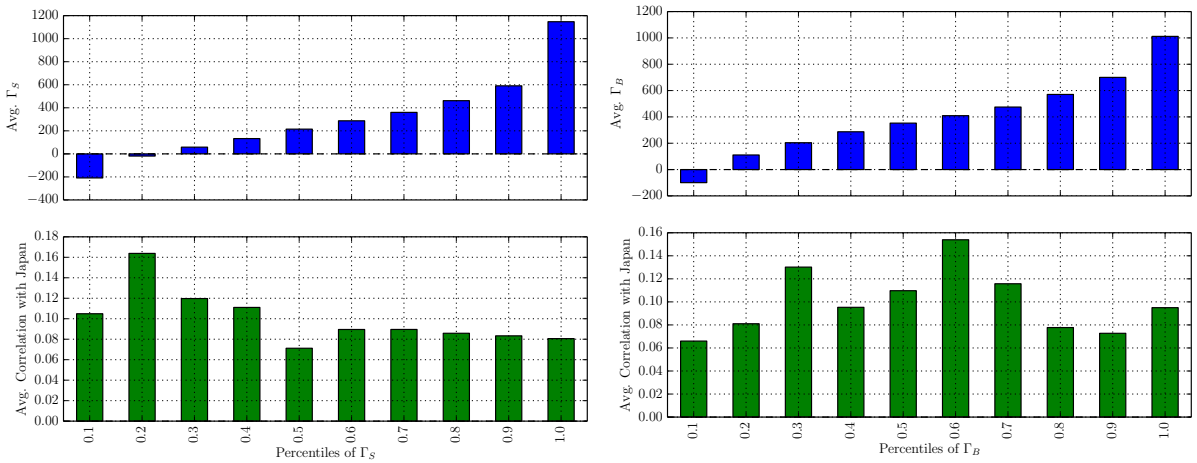

Notes: The left panel in this figure shows the correlation between U.S. reference entity CDS spreads and the country of Japan's (JPN) CDS spread, averaged within deciles of $\Gamma_{S,r}$. $\Gamma_{S,r}$ is the share-weighted average CDS exposure of r's net sellers to Japanese firms. The right panel is the same computation, but reference entities are grouped by deciles of $\Gamma_{B,r}$, which is the share-weighted average CDS exposure of r's net buyers to Japanese firms. Correlations are computed using log-changes in CDS spreads in the 90 days prior to the Japanese tsunami on March 11, 2011.

where $\Delta \log(CDS_1)$ is the log-change in the CDS spread for r following the tsunami. By convention, I take $t = 0$ to be the day of the tsunami, and $t = 1$ to be the week following the tsunami. $\boldsymbol{X}_r$ is a vector of observable characteristics for reference entity r. U_r is a vector of unobserved characteristics that are orthogonal to $\boldsymbol{X}_r$, and ε_r is an error term independent of all variables in the model. Because U_r is unobservable, we can collapse it into the error term, defined by a tilde:

$$\widetilde{\varepsilon}_r = \gamma U_r + \varepsilon_r$$

If we estimate the regression in (16), consistent estimation requires $cor(\Gamma_{S,r}, \widetilde{\varepsilon}_r) = cor\left(\Gamma_{S,r}, U_r\right) = 0$, or $\gamma = 0$. Put differently, it must be the case that r's weighted-average seller exposure to Japan is uncorrelated with unobservable characteristics that caused spread movements. The effectiveness of my identification therefore rests with my ability to argue that any possible omitted characteristics that cause spread movements are uncorrelated with the included covariates.

With that in mind, a natural objection to my identification strategy is firms with high $\Gamma_{S,r}$ are those more correlated with Japan. It could be that U.S. sellers of CDS protection specialize in U.S. reference entities who are fundamentally correlated with the state of the Japanese economy. To alleviate this concern, Figure 13 displays the average correlation between JPN's CDS spread and each U.S. reference entity in my sample, after splitting reference entities into deciles based on $\Gamma_{S,r}$ or $\Gamma_{B,r}$.[68] Each pairwise correlation between r and Japan is computed using log-changes in CDS spreads in the 90 days prior to the tsunami. The top plot in the left panel displays the average $\Gamma_{S,r}$ within each decile of $\Gamma_{S,r}$, and the bottom plot of the same panel shows the average correlation with Japan in that decile. The right panel of the figure repeats the exercise, but splits reference entities into deciles based on $\Gamma_{B,r}$.

Figure 13 provides visual evidence against the "specialization hypothesis". That is, sellers who have large exposures to Japan also have exposures to U.S. firms whose fundamentals are linked to Japan. Regardless of whether reference entities are grouped by $\Gamma_{S,r}$ or $\Gamma_{B,r}$, there is no observable pattern in terms of correlation

[68] A complete description of the data for this Section is in Appendix C.2.

with Japan. Moreover, the average correlation with Japan within any percentile is relatively low, and never reaches above 18%.

C.4 Do Observable Reference Entity Characteristics Explain $\Gamma_{i,r}$?

In addition to explicitly including the aforementioned covariates in my main regression specification, I use two sets of regressions to determine whether observable firm characteristics can explain cross-sectional variation in $\Gamma_{S,r}$ and $\Gamma_{B,r}$. My identification scheme would suggest this is not the case. Hence, to provide additional evidence that Γ is orthogonal to unobserved characteristics of U.S. reference entities that caused spread movements following the tsunami, I run the following pair of regressions:

$$\Gamma_{i,r} = a + \boldsymbol{\beta}' \boldsymbol{X}_r + \varepsilon_r, \quad i = S, B \tag{17}$$

where $\boldsymbol{X}_r$ are the full set of control variables. If U.S sellers of protection who had exposure to Japan were "randomly assigned" to U.S. reference entities, then the regression should have very little explanatory power. Table 10 collects these results.

Table 10: Regression of Γ on Reference Entity Observables

	Dependent Variable	
$\boldsymbol{X}_r$	$\Gamma_{S,r}$	$\Gamma_{B,r}$
$\rho(r, JPN)$	0.007	-0.019
	(0.012)	(-0.21)
σ_{CDS}	1.07	3.01
	(0.34)	(0.69)
$CDS_r(3/11/2011)$	0.007	-0.10
	(0.17)	(-1.71)
N	288	288
Adj R^2	0.9%	0.6%
p-value of F-test	51%	85%

Notes: The table reports the results of the following regression: $\Gamma_{i,r} = a + \boldsymbol{\beta}' \boldsymbol{X}_r + \varepsilon_r, \quad i = S, B$. For example, $\Gamma_{S,r}$ is the share-weighted average exposure of r's net sellers to Japanese firms, where exposure is the net amount of protection sold. X_r is a vector of observable characteristics of reference entity r: the 90-day trailing correlation of (changes in) r's CDS spread with the country of Japan's CDS spread, the 90-day trailing volatility of r's CDS spread, an industry fixed effect based on NAICS industry code, and the level of the CDS spread for r on the day of the tsunami.

Table 10 supports the assumption that U.S. reference entities were "randomly matched" to sellers and buyers with exposure to Japanese firms. In both regressions, none of the explanatory variables are statistically significant, and the adjusted R^2 are near zero. The F-test that all coefficients in the regression are zero cannot be rejected at a 5% confidence interval. In the regressions to come, these results will manifest themselves when I check whether the effect of Γ on CDS spread changes after the inclusion of the aforementioned control variables. Unsurprisingly, adding these controls does not alter the magnitude of my estimated effect of Γ on CDS pricing.

C.5 Does seller J just specialize?

To emphasize the importance of concentration, I explained CDS spread movements after the tsunami using seller J's share in each reference entity, denoted by $\omega_{J,r}$. Recall that seller J had a very large exposure to

Figure 14: Correlation with Japan Across Deciles of $\Gamma_{S,r}$ and $\Gamma_{B,r}$

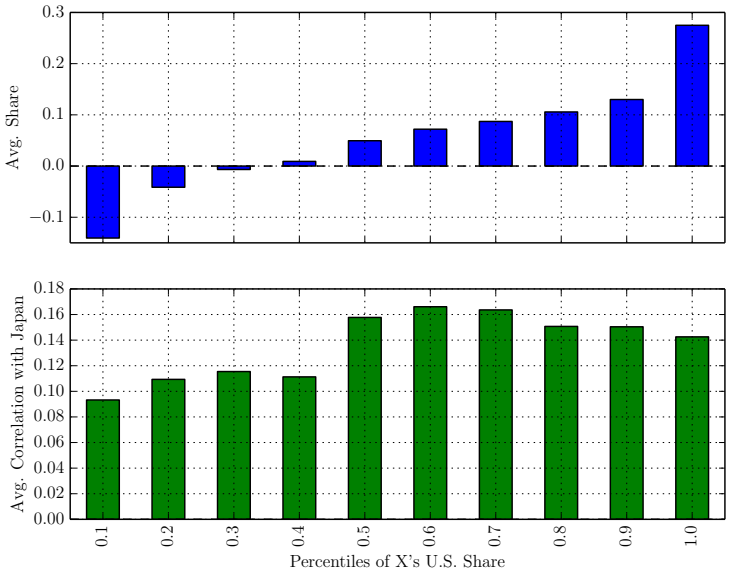

Notes: The left panel in this figure shows the correlation between U.S. reference entity CDS spreads and the country of Japan's (JPN) CDS spread, averaged within deciles of $\Gamma_{S,r}$. $\Gamma_{S,r}$ is the share-weighted average CDS exposure of r's net sellers to Japanese firms. The right panel is the same computation, but reference entities are grouped by deciles of $\Gamma_{B,r}$, which is the share-weighted average CDS exposure of r's net buyers to Japanese firms. Correlations are computed using log-changes in CDS spreads in the 90 days prior to the Japanese tsunami on March 11, 2011.

Japanese firms prior to the tsunami striking. Like with my use of the Γ variables in the previous subsection, a reasonable objection to this approach is that seller J specializes in U.S. firms who are fundamentally connected. I rule this out in the same way as before. I group U.S. reference entities into deciles based on $\omega_{J,r}$ and then compute the average correlation with Japan's CDS spread within each decile. Figure (14) displays the output of this exercise.

At first glance, there does not appear to be a pattern across the deciles of $\omega_{J,r}$. There is a slight increase in correlation with Japan when moving from the 40th the 50th percentile. However, the absolute magnitude of the correlation difference is small, 10 percent versus 16 percent. To be safe, I have estimated my regressions only using reference entities in the 50th percentile and higher. None of my main conclusions are affected. I therefore conclude that seller J is not simply specializing in U.S. firms with fundamental correlation to Japan.

D Appendix: Additional Computations Used in the Paper

D.1 Motivating the Benchmark Regression from Reduced Form Models of Credit Risk

It is standard practice in reduced-form credit risk modeling to view default events as the arrival of a Poisson process.[69] The Poisson arrival rate is most often called the default intensity or the default arrival rate. I denote this variable by $\lambda_{rt}^{\mathcal{M}}$, where the superscript $\mathcal{M} \in \{\mathbb{Q}, \mathbb{P}\}$ defines either the risk-neutral measure, $\mathbb{Q}$, or the physical measure, $\mathbb{P}$. For illustration, I assume that at each point in time the default intensity is constant for the remaining life of the CDS position. In this case, the CDS spread of a given reference entity can be decomposed as:

$$CDS_{rt}(\tau) \quad = \quad \frac{\lambda_{rt}^{\mathbb{Q}}(\tau)}{\lambda_{rt}^{\mathbb{P}}(\tau)} \times LGD_{rt}^{\mathbb{Q}} \times \lambda_{rt}^{\mathbb{P}}(\tau) \tag{18}$$

where $LGD_{rt}^{\mathbb{Q}}$ is the loss given default under the risk-neutral measure. The ratio, $\Pi_{rt} := \lambda_{rt}^{\mathbb{Q}}(\tau)/\lambda_{rt}^{\mathbb{P}}(\tau)$, can be interpreted as the default premium for reference entity r. It quantifies the risk-reward tradeoff for bearing r's default risk.[70] A basic tenet of asset pricing is that an asset's risk premium is a combination of the quantity of risk and the price of risk. As a rough approximation, suppose that the total default premium can be further decomposed by a multiplicative form:

$$\Pi_{rt} = Q_{rt}\Pi_t$$

where Q_{rt} is the quantity of risk for r and Π_t is the price of credit risk for the entire economy. Standard economic logic suggests that the quantity of risk Q_{rt} is high for reference entities that default when marginal utility is high. Suppose further that $Q_{rt} = Q_r$ is constant, so that time-variation in the total default premium is driven only by variation in the price of credit risk Π_t.

In this case, taking the log of both sides of Equation (18) gives:

$$\log\left(CDS_{rt}\right) = \log\left(LGD_{rt}^{\mathbb{Q}}\right) + \log\left(\lambda_{rt}^{\mathbb{P}}\right) + \log(Q_r) + \log\left(\Pi_{rt}\right) \tag{19}$$

where I have omitted the functional dependence of variables on the time to maturity. To make Equation (19) empirically operational I need to have estimates of $\lambda_{rt}^{\mathbb{P}}$ and $LGD_{rt}^{\mathbb{Q}}$. Like in Berndt, Douglas, Duffie, Ferguson, and Schranz (2008), I proxy for $\lambda_{rt}^{\mathbb{P}}$ using Moody's 5-year annualized EDF.[71]

Analogously, I obtain separate estimates of $LGD_{rt}^{\mathbb{Q}}$ from Markit and Moody's. Denote the choice of proxy for $LGD_{rt}^{\mathbb{Q}}$ by $\widetilde{LGD}_{rt}^{\mathbb{Q}}$. It is enough to assume that the true LGD is a scalar multiple η_r of the proxy,

[69] I use the term reduced-form in the spirit of the work by Jarrow and Turnbull (1995), Duffie (1996), and Duffie and Singleton (1999). The popular alternative to this approach are so-called structural models of credit, a la Merton (1974).

[70] Driessen (2002) and Berndt, Douglas, Duffie, Ferguson, and Schranz (2008) provide evidence that, on average, $\Pi_{rt} \approx 1.9$. In other words, for every unit of actual default risk taken, the seller of protection must be compensated as if she is taking roughly double that amount of default risk.

[71] Moody's uses observed equity values and volatility to solve for an implicit asset value process. Using observed leverage, they translate this to a distance-to-default measure as in Merton (1974). Finally, distance-to-default is mapped to a $\mathbb{P}$-likelihood of default using realized default rates.

so $LGD_{rt}^{\mathbb{Q}} = \eta_r \widetilde{LGD}_{rt}^{\mathbb{Q}}$. In logs, this means:

$$\log\left(LGD_{rt}^{\mathbb{Q}}\right) = \log(\eta_r) + \log\left(\widetilde{LGD}_{rt}^{\mathbb{Q}}\right) \tag{20}$$

When might this assumption be reasonable? For instance, the Moody's estimate of LGD is at the sectoral level. Assuming $LGD_{rt}^{\mathbb{Q}} = \eta_r \widetilde{LGD}_{rt}^{\mathbb{Q}}$ then means a firm's LGD is a time-invariant scalar transformation of the sectoral LGD. In other words, time-variation in reference entity LGD is common within a sector, which seems plausible.[72]

Substituting Equation (19) into Equation (20) yields:

$$\log\left(CDS_{rt}\right) = \log(\eta_r) + \log(Q_r) + \log\left(\widetilde{LGD}_{rt}^{\mathbb{Q}}\right) + \log\left(\lambda_{rt}^{\mathbb{P}}\right) + \log\left(\Pi_t\right) \tag{21}$$

Equation (21) is a panel regression, in logs, of CDS spreads on a reference entity fixed effect, plus proxies for the risk-neutral LGD and the physical default intensity. The reference entity fixed effects absorbs both the firm-specific component of LGD, η_r, and the comovement of r with marginal utility, Q_r. After controlling for firm specific variables, Equation (21) suggests the additional control variables that enter the regression capture the price of credit risk, Π_t.

This interpretation rests crucially on the link between CDS and bond markets. It could very well be the case that fluctuations in CDS spreads are not accompanied by variation in bond yields. That is, if I observe CDS spreads changing, it may be the CDS-bond basis — loosely speaking the difference between CDS spreads and bond spreads — is actually what is moving around. In theory, the CDS-bond basis should be zero, but there is a substantial amount of empirical evidence to suggest that this is not always the case.[73] In Appendix B, I use actual bond yields to confirm my results do indeed pertain to the price of credit risk, as opposed to the CDS-bond basis. Finally, Equation (21) suggests one can estimate Π_t using a reference entity controls plus a time fixed effect, which I implement in the next section.

D.2 A Panel Estimate of the Price of Aggregate Credit Risk

Motivated by the analysis in the previous section, I estimate the price of credit risk by running the following regression:

$$\log\left(CDS_{rt}\right) = c + a_r + \beta_1 \log\left(EDF_{rt}\right) + \beta_2 \log\left(\widetilde{LGD}_{rt}^{\mathbb{Q}}\right) + \pi_t + \varepsilon_{rt} \tag{22}$$

where a_r is a reference entity fixed effect and π_t is a time fixed effect. To recover the price of risk at any point in time, I simply exponentiate the time-fixed effect:

$$\widehat{\Pi}_t = \exp\left(\widehat{\pi}_t\right) \tag{23}$$

where the standard error of $\widehat{\Pi}_t$ is computed using the Delta method.[74] Crucially, my price of risk estimate at each point in time is a *relative* estimate. That is, $\widehat{\Pi}_t$ is the ratio of the price of risk at time t to the price of risk on the first day of the sample, which for this portion of the paper is 2/19/2014. The reason this is the case is simply because estimation via the time fixed effect cannot separately identify the price of risk for

[72] When $\widetilde{LGD}_{rt}^{\mathbb{Q}}$ comes from Markit, $\widetilde{LGD}_{rt}^{\mathbb{Q}}$ is provided for each reference entity (as opposed to each sector). In this case, the assumption says the Markit's estimate of LGD is potentially biased in a time-invariant way. For example, if $\eta_r = 1.1$, then I am assuming Markit's LGD for r are always 10% higher than reality. Of course, nothing in my approach restricts η_r from being one.

[73] For example, Bai and Collin-Dufresne (2013).

[74] i.e. the standard error for $\widehat{\Pi}_t$ is $\exp(\pi_t) \times se(\pi_t)$.

the first day in the sample from the fixed effect of the first firm.[75] Even though I am unable to pinpoint the exact level of the price of risk, it is enough for my purposes to study the relative level of the price of risk.

Next, I present the results of the simple panel regression:

$$\begin{array}{rcccccccccc} \log(CDS_{rt}) & = & 0.60 & + & 0.25 & \log(EDF_{rt}) & + & 0.36 & \log\left(\widetilde{LGD}^{\mathbb{Q}}_{rt}\right) + & \sum_r a_r 1_r & + & \sum_t \pi_t 1_t \\ & & (32.4) & & (9.9) & & & (4.8) & & & & \end{array} \tag{24}$$

where the t-statistics are listed below estimated coefficient values. All t-statistics are computed by clustering within each reference entity and time.

A similar approach is taken by Gilchrist and Zakrajsek (2012) to estimate what they call the "excess bond premium". However, instead of estimating Π_t using a time-fixed effect, they first obtain the fitted (level) of CDS spreads from the panel regression. Next, they compute "residuals" as the difference between actual credit spread levels and this fitted variable, along with a Jenson volatility correction. Finally, their excess bond premium in each period is the average residual across all bonds.

My time fixed-effect methodology is a simplification of this procedure. Recall that I motivated the time fixed effect estimator of Π_t from the stylized reduced form model of credit in the previous subsection; however, the validity of my estimate of Π_t is not anchored to the assumptions in that model. Alternatively, Π_t can be interpreted from a purely statistical perspective; it computes within-period averages of the portion of CDS spreads not captured by firm-fundamentals. This is the strictest sense in which I view Π_t as highly related to Gilchrist and Zakrajsek's (2012) excess bond premium.

D.3 Appendix: Using Mark-to-Market Losses to Compute CDS Portfolio Returns

D.3.1 Portfolio Return of a Single Counterparty

Let me start with a simpler question: what is the annualized return, from the perspective of a seller, of a single position? Suppose the position was initiated at time zero, so that the change in the mark to market value at time t is denoted by $V^p_{c,t}$. The p superscript denotes that this is the value of a single position, p. As a matter of convention, $V^p_{c,t}$ is computed from the perspective of a seller. Thus, if c is a buyer in the position, the contribution of the position to her overall portfolio is $-V^p_{c,t}$.

To answer the question, one must make an important assumption on the leverage allowed in the position. If the notional of the position is $N^p_{c,t}$, then zero leverage indicates that the seller of protection must post the full value of $N^p_{c,t}$ into an initial margin account at time zero. The zero leverage assumption has been made by previous studies such as Bao and Pan (2012), Berndt and Obreja (2010), and Junge and Trolle (2014). For generality, I assume that the seller of protection must post $l \in (0, 1]$ of the total notional as an initial margin. If at time t, the position is unwound, the annualized return is computed by looking at the cashflows:

$$\begin{aligned} R^p_{c,t} &= \frac{1}{\tau^p_t} \times \frac{\left(lN^p_{c,t} + V^p_{c,t} + AP^p_{c,t}\right) - lN^p_{c,t}}{lN^p_{c,t}} \\ &= \frac{1}{\tau^p_t} \times \frac{V^p_{c,t} + + AP^p_{c,t}}{lN^p_{c,t}} \end{aligned}$$

where τ^p_t is the time that has elapsed, in years for the position. In this example, $\tau^p_t = t$. $AP^p_{c,t}$ is the accrued

[75] This is one major advantage of the approach in BDDFS, who estimate a process for $\lambda^{\mathbb{Q}}_t$ and $\lambda^{\mathbb{P}}_t$ directly. As such, they can compute the price of risk (e.g. the ratio) at each point in time. Even in this case, they must determine the initial condition for each process in order to determine the level of the default intensities. This initialization comes from the parameters of the underlying processes.

premium, from the perspective of the seller, of the position.

At time t, the seller of protection receives back her initial margin of $lN^p_{c,t}$, plus the proceeds of $V^p_{c,t}$ from unwinding the position and any accrued premiums.[76] Her initial investment was $lN^p_{c,t}$, and scaling by $1/\tau^p_t$ puts the return in annualized terms.

The issue of the correct leverage, l, is important. It is generally difficult to observe, and for the transactions in my dataset it seems to be poorly measured.[77] To make matters worse, it seems that leverage is heterogeneous even across the positions of a single counterparty. This makes sense, as bargaining power will certainly play a role in the amount of leverage that one side of the swap can take.

In the absence of better data, I must make some additional assumptions in order to proceed. For the remainder of the paper, I assume that $l = 1$ for all counterparties. This is undoubtedly too simplified, as it seems natural that different counterparty types will be allowed to take different leverage. For robustness, I replicate all of the results in this section using dollar changes in portfolio value in the Online Appendix. This approach does not suffer from the aforementioned caveats, but does make interpretation slightly different since there is great deal of heterogeneity in the trading volumes of different counterparties in the market, not to mention the issue of annualizing gains.

Under the assumption of $l = 1$, the total annualized return of a counterparty's portfolio is just the weighted average of the returns of each position, weighted by the gross size of the position:

$$\begin{aligned} R_{c,t} &= \sum_{p\in\mathcal{P}_t(c)} \left[\frac{N^p_{c,t}}{\sum_{p\in\mathcal{P}_t(c)} N^p_{c,t}} \right] R^p_{c,t} \\ &= \frac{\sum_{p\in\mathcal{P}_t(c)} \left(V^p_{c,t} + AP^p_{c,t}\right)/\tau^p_t}{\sum_{p\in\mathcal{P}_t(c)} N^p_{c,t}} \end{aligned} \tag{25}$$

where $\mathcal{P}_t(c)$ is the set of counterparty c's positions as of time t.

D.3.2 Portfolio Return of Top Five Sellers and Buyers

Denote the aggregate top five buyers by TAB_t for "top aggregate buyers". TAS_t is the top five aggregate sellers. I treat each group as a single trader, and compute the total annualized CDS portfolio return. These are my measures of the risk bearing capacity of mega-sellers and mega-buyers. Formally, this means:

$$\begin{aligned} RB^s_t &:= \frac{\sum_{c\in TAS_t}\left[\sum_{p\in\mathcal{P}_{t-1}(c)}\left(V^p_{c,t-1} + AP^p_{c,t-1}\right)/\tau^p_{t-1}\right]}{\sum_{c\in TAS_t}\left[\sum_{p\in\mathcal{P}_{t-1}(c)} N^p_{c,t-1}\right]} \\ RB^b_t &:= \frac{\sum_{c\in TAB_t}\left[\sum_{p\in\mathcal{P}_{t-1}(c)}\left(V^p_{c,t-1} + AP^p_{c,t-1}\right)/\tau^p_{t-1}\right]}{\sum_{c\in TAB_t}\left[\sum_{p\in\mathcal{P}_{t-1}(c)} N^p_{c,t-1}\right]} \end{aligned} \tag{26}$$

D.4 Option Implied CDS Spreads

This section describes how I use American option prices to compute an implied CDS spread. For a complete theoretical treatment of this procedure, see Carr and Wu (2013), henceforth CW. In the interest of space, I

[76] Implicitly, this assumes that the initial margin account earns zero interest. This can be easily relaxed.

[77] See a recent paper by Duffie, Scheicher, and Vuillemey (2014) for a more detailed discussion of initial margining.

present only the relevant formulas and data descriptors used in the main text.

To start, Carr and Wu (2013) define what they call a "unit recovery claim" that pays a dollar if there is a default event prior to an option's expiration, and zero otherwise. CW assume that there exists a default corridor $[A, B]$ that the underlying equity price can never enter. If the equity price hits the level B, there is a default and the stock price immediately jumps to a level that is bounded above by A. In their empirical work, they set $A = 0$, which means that the equity value drops to zero upon default. I continue with this assumption for the remainder of my treatment.

Under this assumption, CW show that, regardless of the underlying asset process, there is a robust link between the unit recovery claim and CDS spreads on the underlying firm. The unit recovery claim is defined as follows:

$$U^O(t,T) = \frac{P_t(K_2,T) - P(K_1,T)}{K_2 - K_1} \tag{27}$$

where $A \leq K_1 < K_2 \leq B$. It is easy to see that, under the assumptions of the default corridor, this pays one dollar if there is default and zero otherwise.

Next, CW show that under the assumption of a constant arrival rate and constant interest rate, the CDS spread of a firm is related to the price of the unit recovery claim in the following manner:

$$U^O(t,T) = \xi k \times \frac{1 - \exp\left(-(r + \xi k)(T-t)\right)}{r(t,T) + \xi k} \tag{28}$$

where $\xi = 1/(1-R)$, R is the recovery of the bond upon default, k is the CDS spread, and $r(t,T)$ is the continously compounded interest rate between t and T. Here, T is meant to capture the expiration of both the CDS contract and the option contract. For my purposes, I will always set $T - t = 5$.

Equation (28) provides a simple way to recover a CDS spread implied by option prices. Using observed option prices, one first computes the value of the unit recovery claim. A simple numerical inversion then delivers the implied CDS spread.

To implement this procedure in practice, I merge my panel of CDS spreads with American option prices from OptionsMetrics using 6 digit CUSIPs. Furthermore, since I follow CW in assuming $A = 0$, the unit recovery claim is simple the price of a deep out of the money put option, divided by its own strike price. I use a set of filters on the options data that is similar to CW: (i) I take the option price to be the midpoint of the bid and offer; (ii) I consider options whose bid is strictly positive; (iii) I consider options whose open interest is strictly positive; (iv) the maturity of the option must be greater than 365 days; (iv) I use the put option that satisfies all of the preceding qualities, and that has the delta closest to 0 and less than -0.15.

Naturally, there is a maturity mismatch in using options that might have an expiration of 2 years to compute an implied CDS spread of 5 years. There is no real way to avoid this bias. See CW for a richer discussion. Like with other portions of the paper, the riskfree rate is obtained from interpolating the USD swap rate curve. Finally, I use the Markit reported recovery rate, which has the added advantage of maintaining consistency with the benchmark panel regression in the main text.