Concentrated Capital Losses and the Pricing of Corporate Credit Risk*

Emil N. Siriwardane†

Job Market Paper‡

November 4, 2014

Abstract

Using proprietary credit default swap (CDS) data from 2010-14, I discover that CDS markets are dominated by a handful of net protection sellers, with five sellers accounting for nearly half the market. In turn, financial frictions — in the form of capital losses — for these sellers affect the pricing of corporate credit risk. A one standard deviation loss to the CDS portfolios of the five largest sellers increases weekly credit spreads by 1.4 percent. CDS portfolio performance of the five largest sellers also appears to explain about one-eighth of variations in the price of credit risk. To alleviate identification concerns, I use heterogeneous exposure to Japanese firms following the 2011 tsunami as an exogenous source of variation in the risk bearing capacity of CDS traders. After the tsunami, spreads increased for U.S. firms whose sellers were most exposed to Japanese firms. Finally, I examine CDS transactions to exploit variation from a single buyer purchasing protection from multiple sellers. Consistent with notion that frictions impact prices, I find sellers who have experienced losses charge higher premiums.

*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, Nishani Bourmault, Matteo Crosignani, Itamar Dreschler, Greg Duffee, Darrell Duffie, Andrea Eisfeldt, Nikunj Kapadia, Patricia Mosser, Holger Mueller, Matteo Maggiori, 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 or emil.siriwardane@treasury.gov

‡The latest version of this paper can be found at: http://people.stern.nyu.edu/esiriwar/pdf/ENS_JMP.pdf

1 Introduction

In this paper, I explore whether financial frictions — in the form of capital losses — impact the pricing of corporate default risk. I study the credit default swap (CDS) market, in which buyers of CDS pay a premium to sellers for insurance against corporate default. In a transparent and frictionless environment, the losses of one seller should not affect CDS pricing. In other words, if a seller loses money and subsequently demands higher insurance prices as compensation, he should not receive future trades because another investor will offer better terms. My central result leads to the opposite conclusion. I find that there are very few net sellers of CDS protection in U.S. markets, and in turn, capital losses for these important sellers strongly affect the pricing of credit risk.

In order to conduct my analysis, I use a proprietary dataset of CDS exposures from 2010-14 that is provided by the Depository Trust & Clearing Corporation (DTCC). The DTCC supplies trade processing and registration services for all major dealers in CDS, meaning the data effectively covers the entire market.[1] Section 3 of this paper precisely characterizes the nature of credit risk transfer in U.S. CDS markets, but the broad takeaway is as follows: CDS markets transfer significant amounts of net credit risk (approximately $1.7 trillion), a large share of which can be attributed to a handful of protection buyers and sellers, with sellers more scarce than buyers. I call these large CDS players "mega-sellers", "mega-buyers", or more generally, "mega-players".

The impact of financial frictions is most likely to appear in a concentrated market, so the existence of mega-players in CDS makes the market a natural setting for investigating whether financial frictions affect asset pricing. To visualize my primary result, the left axis of Figure 1 plots the one-week lagged dollar losses on the CDS portfolio for two mega-sellers in the market. For each counterparty c, I denote the dollar loss of its CDS portfolio at time t by $V_{c,t}$, where positive values represent losses. The right axis also provides an estimate of what I refer to as the aggregate price of credit risk, denoted by $\widehat{\pi}_t$, relative to its value in February 2010. Loosely speaking, $\widehat{\pi}_t$ captures the component of credit spreads that cannot be explained by countercyclical movements in expected defaults. The full details of how I construct $\widehat{\pi}_t$ are found in Section 4.2.[2]

[1] Importantly, I also account for exposures through single name and index swaps. A single name swap is, as its name suggests, just insurance written on the default of a single reference entity. An index swap is written on a basket of reference entities. See Section 2 for more details.

[2] It derives from a panel estimator that uses a large number of CDS spreads, expected default frequencies, and expected losses given default. Gilchrist and Zakrajsek (2012) refer to their version of essentially the same quantity as the "excess bond premium."

Figure 1: Aggregate Price of Credit Risk and Lagged Losses by Two Mega-Sellers

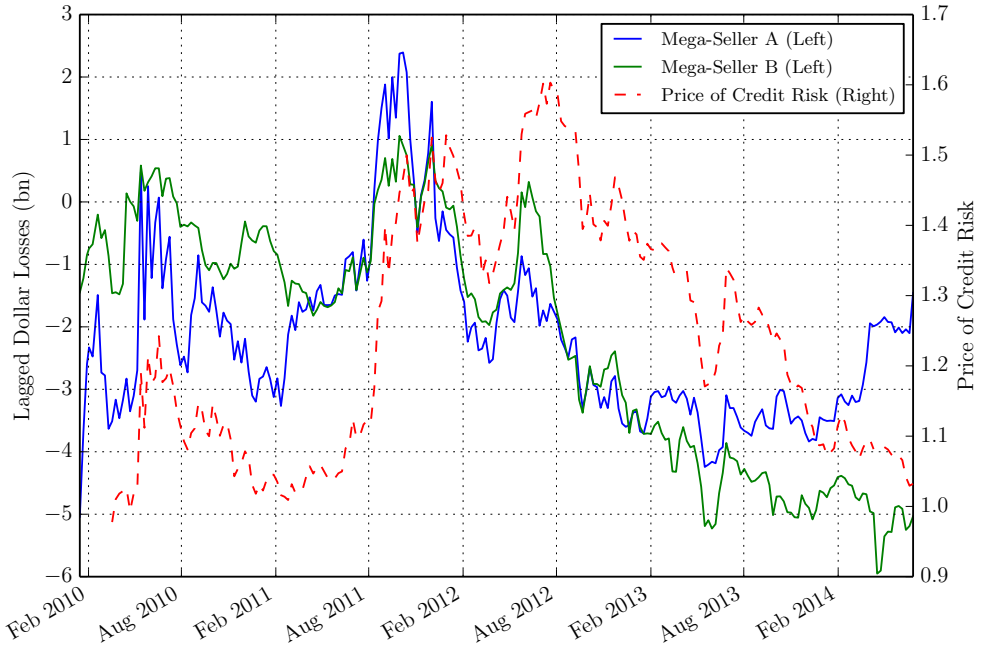

Notes: On the right axis, this figure plots the price of credit risk, $\widehat{\pi}_t$, estimated from a panel of CDS spreads using the regression in Equation (24). $\widehat{\pi}_t$ is interpreted as the price of credit risk relative to its value on February 2010, and is the portion of credit spreads that cannot be explained by countercyclical movements in expected defaults. The left axis is the lagged mark-to-market dollar loss of the CDS portfolios for two mega-sellers in the market, lagged by one week. Data is weekly and spans February 2010 to June 2014.

When these two mega-sellers experience losses, their *effective* risk aversion rises, thereby increasing the premium for bearing default risk.[3] From an institutional perspective, this is a sensible result for a number of reasons. First, trading desks at large dealers and hedge funds are subject to risk limits (e.g. value-at-risk), which may tighten with prolonged losses. Second, and more importantly, poor portfolio performance means CDS traders have less capital to make variation margin payments on mark-to-market losses. When raising new capital on short notice is costly, losses will naturally constrain the ability to take on new risk. Consequently, I treat CDS portfolio losses as a proxy for variations in risk bearing capacity.[4] Section 4.1 provides more institutional details and some preliminary empirical

[3]For instance, the correlation between mega-seller A's realized losses, $V_{A,t-1}$, and $\widehat{\pi}_t$ is 25 percent. This correlation is not being driven by a timing effect. If I increase the lag time to one month, the correlation between $V_{A,t-4}$ and $\widehat{\pi}_t$ is slightly higher at 30 percent.

[4]See Atkeson, Eisfeldt, and Weill (2014) 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

evidence motivating this interpretation.

I formally explore the relationship between mega-player losses and default risk premiums using a panel regression with CDS spreads as the dependent variable. My analysis focuses primarily on mega-sellers of CDS because I do not observe the joint bond-CDS portfolio for buyers, making it difficult to interpret what a buyer's CDS loss does to his entire portfolio. On the other hand, it is more likely that seller losses affect overall wealth because directly hedging a sold CDS position requires costly shorting of the underlying bond. Furthermore, net sellers are almost twice as concentrated as buyers, so their risk bearing capacity has a greater influence on CDS pricing. As expected, I find little evidence of a relationship between the price of credit risk and mega-buyer CDS portfolio performance.

In contrast, mega-seller CDS portfolio performance has a profound effect on risk premiums. A one standard deviation portfolio loss of the five largest sellers results in a 1.4 percent increase in spreads, after controlling for a host of macroeconomic and firm-specific factors. These elasticities are economically large. To put this in perspective, in my sample, the standard deviation of spread movements across all reference entities and times is 6 percent. Moreover, a simple regression of changes in the price of credit risk, $\widehat{\pi}_t$, on changes in mega-seller risk bearing capacity yields an incremental R^2 of 12 percent. Prima facie, this means fluctuations in mega-seller risk bearing capacity can explain about one-eighth of the variation in the price of credit risk.

Of course, movements in CDS spreads do not mechanically translate to movements in bond markets, as no-arbitrage restrictions would suggest.[5] Instead, changes in CDS spreads might be attributed to fluctuations in the CDS-bond basis, or the difference between CDS spreads and the bond-implied credit spread. However, bond spread data — as opposed to CDS spread data — confirms that mega-seller losses affect the price of credit risk and not just the CDS-bond basis.

To bolster the interpretation that mega-seller losses increase the premium for bearing default risk, the latter portion of the paper addresses the problem of identification in four ways. The first, and easiest, is to control for a rich set of observables that might cause the price of credit risk to be high. I choose these variables based on theoretical models of credit risk, variables that have been empirically shown to explain credit spreads, and measures of CDS spreads that are implied by other asset classes (e.g. options).

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.

[5] See, for instance, Duffie (1999).

The second way I treat identification is to test whether losses in one portion of mega-sellers' CDS portfolios affect pricing for unrelated reference entities. To illustrate my strategy, I check whether the CDS premium for GE is high after mega-sellers have lost money on their CDS portfolio, *excluding reference entities in the same industry as GE*. My primary conclusion that mega-seller losses affect CDS spreads is robust to this exercise, even after accounting for firm characteristics and all unobservable macroeconomic variables.[6]

The third way I address identification is by using the 2011 Japanese tsunami to study how an exogenous shock to mega-seller risk bearing 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 protection sellers were highly exposed to Japanese firms saw their CDS spreads rise 2.5 percent in the week after the tsunami, relative to reference entities whose sellers had very low exposure to Japan. Importantly, one of my controls is each firm's equity return, which rules out a large number of unobservable firm-level or macroeconomic variables that could cause U.S reference entity spreads to change in response to the tsunami.

For all of the preceding analyses, I used aggregated CDS spread quotes, but my last identification technique also finds evidence that seller losses impact pricing using executed transactions.[7] More specifically, I exploit variation in CDS premiums for a single buyer purchasing protection from multiple sellers on a fixed reference entity and date.[8] 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 within a buyer, reference entity, and trade date, sellers who have lost more money charge relatively higher premiums for CDS protection. These results are compelling because the effect should be hardest to detect in this setting — one would not necessarily expect sellers with low risk bearing capacity to transact at all.[9]

In sum, the above identification exercises support the interpretation that depressed risk

[6] Froot and O'Connell (1998) study a similar problem in catastrophe reinsurance markets.

[7] 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.

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

[9] Indeed, a simple logit regression confirms that the likelihood of a seller transacting is lower when he has lost money.

bearing capacity of mega-sellers causes increased reference entity CDS spreads. With this mechanism in the background, I employ reduced form panel regressions of CDS spreads on mega-player risk bearing capacity to determine the aggregate asset pricing consequences.

Related Literature

My primary contribution is to show that concentrated capital losses play a sizable role in generating time variation in default risk premiums. My paper therefore adds to a large literature on how credit risk is priced in the economy.[10] A short and certainly non-exhaustive list of work in this area includes Duffee (1998), Elton, Gruber, Agrawal, and Mann (2001), Collin-Dufresne, Goldstein, and Martin (2001). More recently, structural models of credit have introduced time varying risk premiums for corporate bonds. Some examples of this approach include Chen, Collin-Dufresne, and Goldstein (2009) or Chen, Cui, He, and Milbradt (2014). These papers focus on firm exposure to macroeconomic or liquidity risks, whereas my results propose that financial frictions partially account for some movement in credit risk prices.

As a result, I contribute to a rapidly growing literature that highlights the importance of financial intermediary risk bearing capacity for asset pricing. Some noteable contributions in this field include Brunnermeier and Sannikov (2014), He and Krishnamurthy (2012), Adrian and Boyarchenko (2013), and Maggiori (2014). One of my new findings is that financial intermediaries are responsible for a large portion of the aggregate selling of CDS insurance; that their risk bearing capacity matters for credit risk pricing is consistent with the aforementioned theoretical work. Moreover, empirical studies like Adrian, Colla, Shin (2011) and Adrian, Etula, Muir (2014) have used theory-based measures like leverage to examine financial frictions and asset prices. My paper augments this work by providing some of the first direct empirical evidence that frictions do indeed matter for asset pricing.[11]

In addition, Gilchrist and Zakrajsek (2012) have a closely related paper demonstrating that shocks to the price of credit risk lead to significant declines in consumption, investment,

[10] Broadly speaking, the literature on the pricing of risky debt can be split into so-called structural models of credit deriving from the contingent claims analysis of Merton (1974), and reduced form credit risk models in the vein of Duffie and Singleton (1998). This paper uses elements of both, but the primary hypotheses are tested using ideas from reduced form credit risk models.

[11] Furthermore, much of the empirical work has been based on the idea that intermediaries are the marginal investor in the asset market, so their stochastic discount factor is relevant for pricing. In turn, these studies use aggregate measures of risk bearing capacity to price assets. Because of my data, I can actually identify the marginal players in the market, which naturally affords me more power to tease out the theoretical predictions made by asset pricing models with financial frictions.

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 shocks to the CDS portfolios of mega-sellers, some of whom are financial intermediaries, cause shocks to credit risk prices. In Section 4.3, a vector autoregression showing the impulse response of the price of credit to a mega-seller shock ties the results of this paper directly to Gilchrist and Zakrajsek (2012). Consequently, my findings have important implications not just for asset pricing, but also real activity.[12]

My work is also related to a recent theoretical paper by Atkeson, Eisfeldt, and Weill (2014), 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.[13] My findings are broadly consistent with this intuition as well.

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 in this study have been used by Shachar (2012), Chen, Fleming, Jackson, Li, and Sarkar (2011), Oehmke and Zawadowski (2013), Atkeson, Eisfeldt, and Weill (2013), and Duffie, Scheicher, and Vuillemey (2014). To the best of my knowledge, I am among the first to be able to make general statements about the industrial organization of U.S. CDS markets based on individual transaction data. The end result is a map that captures full economic exposures via CDS for every counterparty and reference entity in the U.S.

As a result, I contribute to an important literature that studies how derivatives markets may affect financial stability. Duffie, Li, and Lubke (2010) and references therein provide an overview of these issues. A lot of regulatory focus has been on improving transparency and

[12] 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.

[13] More specifically, their equilibrium condition for exit is equivalent to reducing risk limits taken by traders by a constant that is generally less than one.

reducing counterparty risk in OTC markets, particularly through central clearing of OTC derivatives. My paper highlights a slightly 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 provides 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 formally explores the relationship between the price of credit risk and mega-player CDS portfolio performance. In Section 5, I conduct robustness tests that support the main conclusions drawn in the Section 4. Section 6 contains one of the ways in which I address identification using the 2011 Japanese tsunami as an exogenous shock to mega-seller risk bearing capacity. Section 7 presents an alternative way I mitigate identification concerns through actual CDS transactions. Finally, Section 8 concludes.

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, and 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 that the counterparty is a net buyer of protection. $\mathcal{C}_t$ is the set of all counterparties that are active in the CDS market as of date t. Lastly, $\mathcal{R}_t$ is the set of reference entities that are traded in the CDS market as of date t.

2 Data Description and Empirical Work

My analysis applies to the time period between January 2010 through May 2014.

2.1 Data on CDS Transactions and Positions

The DTCC data I work with is limited in the following way: I see only transactions that involve either (i) a U.S. reference entity or (ii) at least one of the counterparties in the swap is registered in the U.S.. Additionally, this dataset is enormous. 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.

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.

Index Swaps Index swaps constitute nearly half the gross notional of the entire CDS market.[14] 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.

[14] See, for example, Rajan and Siriwardane (2014).

A Comment on Bearing Credit Risk Through CDS I will often make statements about ultimate sellers of CDS protection bearing corporate credit risk. However, because I do not see the total portfolio of any counterparty, I am implicitly assuming that it is difficult to hedge the 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. At a minimum, it seems safe to assume that the entirety of a net sold position could not be hedged by a short sale of the underlying cash instrument. 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 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 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.

3 Facts About Credit Risk Sharing in CDS Markets

Before exploring whether losses at mega-players affect 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 markets. I then build a simple measure that captures the aggregate concentration of net buyers and net sellers of protection. In Appendix B, I explore alternative ways of quantifying concentration. 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 for 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, as I argue in the Online Appendix Section 1.1, the gross notional does not provide much information about the amount of risk transferred.[15]

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:

$$\begin{aligned} NO(t) &= \sum_{r \in \mathcal{R}_t} NO(r,t) \\ &= \sum_{r \in \mathcal{R}_t} \sum_{c \in \mathcal{C}_t} \max\left(NS(c,r,t), 0\right) \end{aligned} \tag{2}$$

The left axis of Figure 2 plots my estimate of the total net outstanding $NO(t)$ through time.[16] At the beginning of the sample, the size of the U.S CDS market was just under \$2 trillion, but has 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.

[15] 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. Nonetheless, I analyze the gross notional and related metrics in the Online Appendix, Section 1.

[16] A more refined look at the size of U.S. CDS markets is the Online Appendix, Section 1. I compute my estimate by taking the average size between the: (i) the entire dataset and (ii) reference entities who 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 2: Net Notional Outstanding in CDS Markets

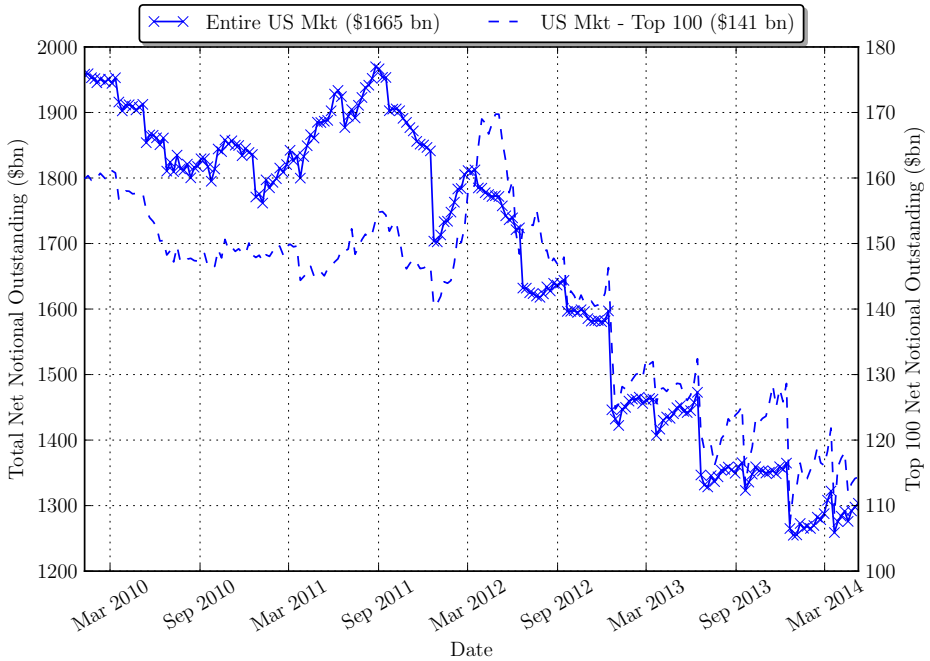

Notes: The left axis of this figure plots the net notional outstanding in U.S CDS markets, computed as $NO(t) = \sum_{r \in \mathcal{R}_t} \sum_{c \in \mathcal{C}_t} \max(NS(c,r,t), 0)$. The size of the U.S. market is the average of the net notional using all positions in the dataset and positions for 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.

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 2 also shows 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.[17]

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 Appendix B, I take a complimentary 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 influence 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.

[17] 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 3: Share of Top Five Aggregate Sellers and Buyers of CDS Protection

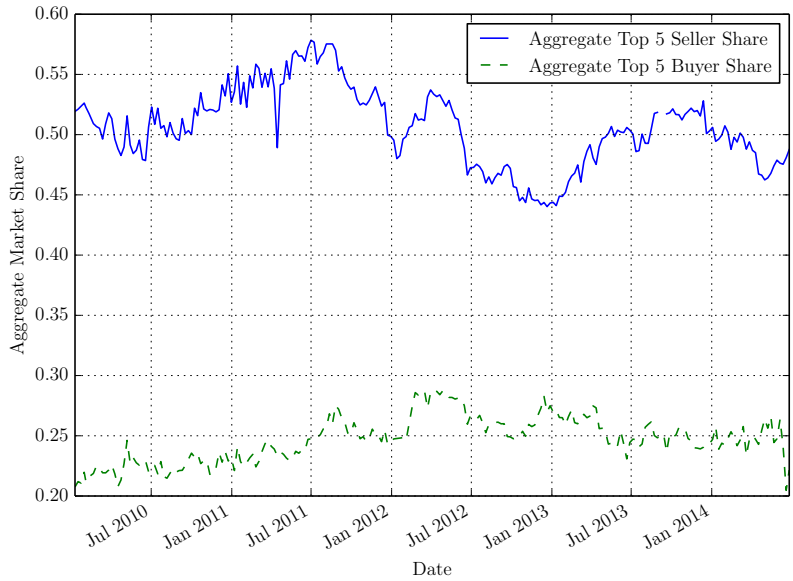

Notes: This figure plots the aggregate share of the top five sellers and buyers of CDS protection through time. The share of a single counterparty c in a given reference entity is the proportion of net selling by c in that reference entity. The aggregate share of net selling by c is the size-weighted 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.

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 3 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 1000 active counterparties in the market at any point in time, 50 percent of all selling is in the hands of only 0.5 percent of potential counterparties. Buyers are also concentrated, albeit only half 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 4 plots, for both buyers and sellers, the number of top five counterparties that remains the

Figure 4: Persistence of Top Five Aggregate Sellers and Buyers

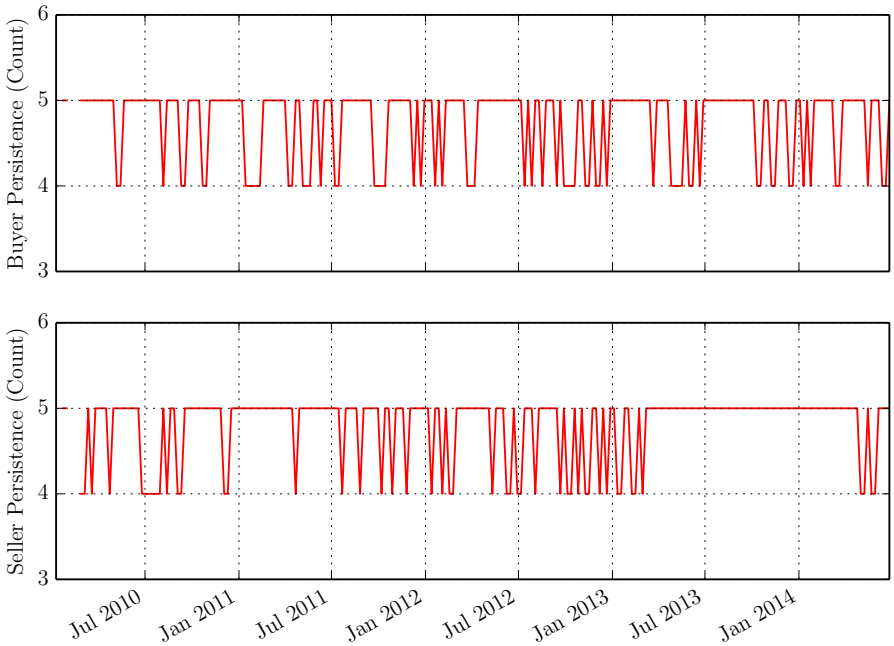

Notes: This figure plots the persistence of the aggregate share of the top five sellers and buyers of CDS protection through time. For week t, I count the number of the top five buyers who are also in the top five in week $t-1$. I do the same for the persistence of top sellers.

same from time $t-1$ to t. 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 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

Figure 5: Aggregate Proportion of Buying and Selling, by Counterparty Type

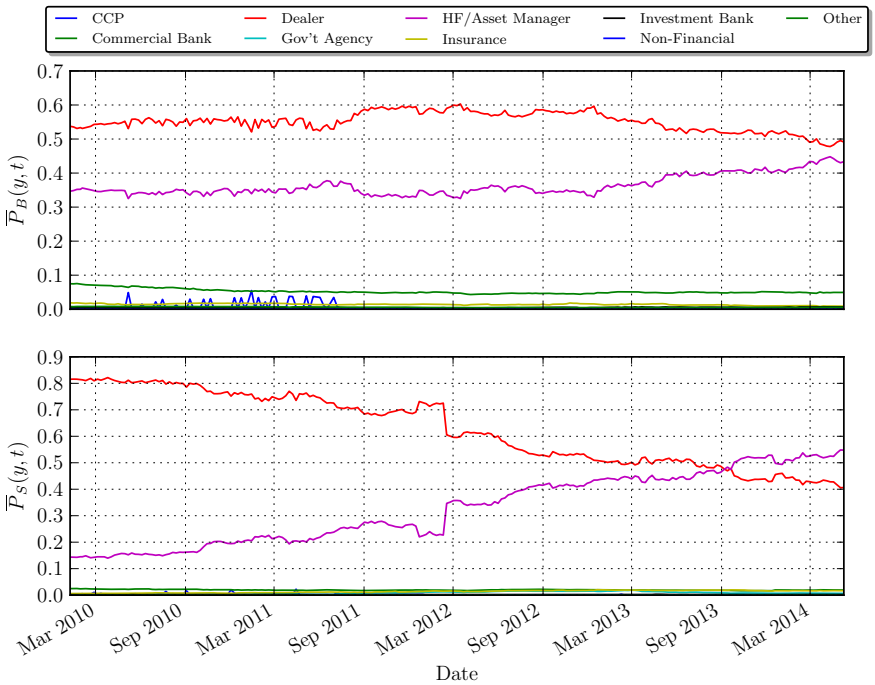

Notes: This figure plots the aggregate proportion of net selling and buying done by each counterparty type. For each reference entity, I compute the proportion of net buying and selling by each counterparty type. To aggregate, I compute the size-weighted average, across reference entities, of the proportion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

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.[18]

Figure 5 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 buying going to HFAMs.

The aggregate proportion of selling by counterparty types appears in the bottom panel of Figure 5. 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

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

dealers has declined by almost half, with dealers accounting for almost 40 percent of total selling by the end of the sample.[19] 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.

4 Capital Losses and the Pricing of Credit Risk

The hypothesis of this section is simple: because CDS markets are dominated by only a handful of sellers and buyers, changes in the risk bearing capacity of these key players should impact the pricing of default risk. I explore this hypothesis using a panel regression, which is useful for establishing a basic quantitative understanding of what drives the price of credit risk. I then address potential identification issues after presenting the benchmark empirical results.

Benchmark Regression Specification

For now, I generically define the risk bearing capacity of mega-sellers and mega-buyers as RB_t^s and RB_t^b, respectively. The benchmark panel regression I use to flesh out my main thesis is as follows:

$$\log(CDS_{rt}) = a_r + \phi_1 \log(LGD_{rt}) + \phi_2 \log(EDF_{rt}) + \gamma_S RB_t^s + \gamma_B RB_t^b + \boldsymbol{\beta}'\mathbf{X}_t + \varepsilon_{rt} \quad (4)$$

where CDS_{rt} is the 5-year CDS spread for r, LGD_{rt} is an estimate of the risk-neutral loss-given-default (LGD), and EDF_{rt} is the 5-year estimated default frequency provided by Moody's. $\mathbf{X}_t$ is a vector of control variables designed to capture the macroeconomic environment, and I will define these variables when I present the empirical results. While my benchmark regression takes a seemingly straightforward form, I motivate this specification from reduced form models of credit risk, with the details in Appendix F.1. Intuitively, because I control for firm-level fundamentals, RB_t^s and RB_t^b affect credit spreads through their impact on the premium for bearing default risk.

[19] 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.

4.1 Losses on CDS Portfolio as Risk Bearing Capacity

I now develop measures of RB_t^s and RB_t^b that are meant to capture the risk bearing capacity of the largest buyers and sellers in the market at any point in time. My eventual metrics will be a function of the CDS portfolio performance of each counterparty. I first provide motivation for using portfolio performance, and then define it properly.

4.1.1 How Do Losses Connect to Risk Bearing Capacity?

It is natural that the performance of a credit derivatives desk (or portfolio) should affect risk bearing capacity. Based on conversations with practitioners, credit derivatives traders will become effectively more risk-averse after a poor month or quarter. One plausible psychological channel is that when traders have lost significant amounts of money, they may become less aggressive in their strategies.[20] Moreover, there are institutional details that suggest risk bearing capacity and portfolio performance 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).[21] 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.

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 (2010). Theoretical treatments of VaR constraints use the wealth of the trader as a state variable. Naturally, this wealth depends on

[20] 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."

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

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.[22]

Funding Constaints

Perhaps the most likely way for portfolio performance 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.[23] Froot and Stein (1996) 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.[24]

Preliminary Evidence

A complete investigation of how exactly losses impact risk bearing capacity is outside the scope of this paper, but there is high-level support of a link in the data. For illustration, suppose that the funding friction channel is what connects losses and risk bearing capacity.

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

[23]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.

[24]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 (2012) provide a detailed overview of the potential issues.

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. Thus, for losses to affect risk bearing capacity, it must be that the CDS desk is 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.[25] I restrict myself to the subset of dealers in my sample because I can only compute leverage measures for these counterparties. I outline how I compute the mark-to-market value of the CDS portfolio shortly in Section 4.1.2.

Presumably dealers with high measured leverage face tighter borrowing constraints. The funding friction channel then implies that negative changes in the CDS portfolio depress risk bearing capacity because they are also accompanied by increases in external funding constraints for the dealer. 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.[26] Figure 6 collects the results, with each pair of bars in the plot corresponding to a single dealer.

The leftmost counterparties in the figure 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 will also bear out in my benchmark regression results in Section 4.2.

The important portion of the graph 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 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.

[25] i.e. the ratio of book value of debt to market value of equity.

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

Figure 6: Correlation Between CDS Portfolio Changes and Market Leverage

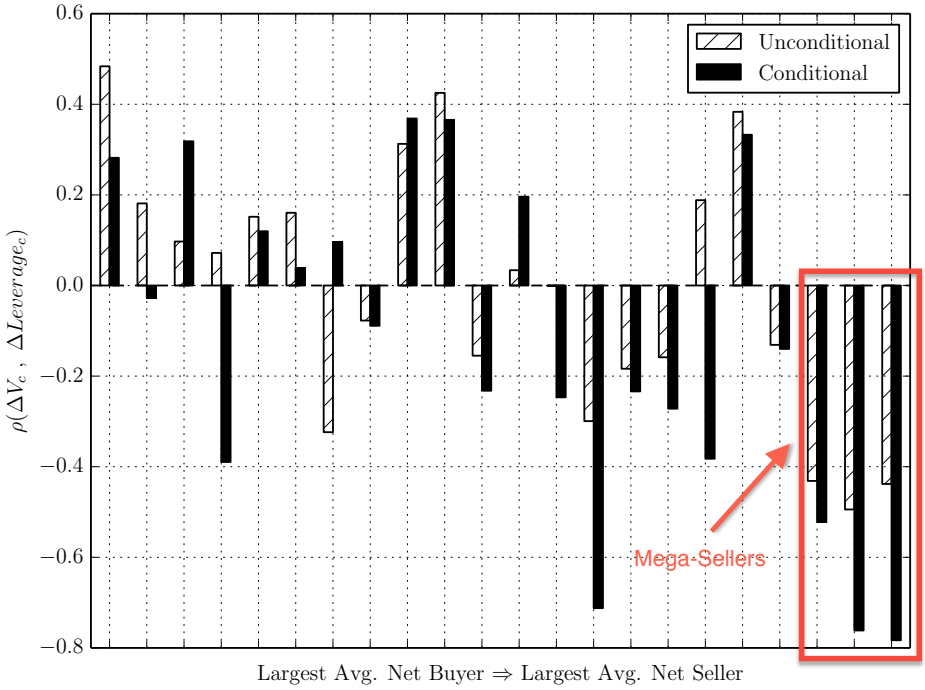

Notes: The figure computes, for each counterparty c who is a dealer, the correlation between log-changes in market leverage, $\Delta Leverage_c$, and the change in the mark-to-market dollar value of the portfolio, ΔV_c. In addition, I compute the conditional correlation between these two series. The 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.

The correlations jump substantially when conditioning 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 limits and VaR story. It is entirely possibly (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.

A second important fact is the CDS portfolios of these three mega-sellers are highly correlated, meaning losses for these key players occur simultaneously. Unconditionally, their pairwise correlations are 56, 58, and 83 percent. Moreover, if I condition their combined portfolio on being in its 6-month 10 percent tail, the pairwise correlations all rise substantially

to 88, 89, and 92 percent. For there to be an aggregate effect of losses on the price of credit risk, it must be the case that mega-sellers are constrained at the same time; otherwise, credit risk would flow from constrained to unconstrained mega-sellers because the unconstrained sellers can offer better prices.

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

In light of this discussion, I assume the risk bearing capacity of a counterparty is an increasing function of performance of its entire CDS portfolio.[27] Using CDS portfolio performance, as opposed to measured leverage, allows me to compute risk bearing capacity measures of all counterparties in my sample (e.g. hedge funds don't have readily available leverage measures). This is important because a handful of non-dealers are mega-sellers of CDS protection. My remaining task is to quantify CDS portfolio performance for mega-players in the overall market.

4.1.2 Marking-to-Market and CDS Portfolio Returns

To quantify the portfolio performance of a counterparty, I first compute the dollar gain of its CDS positions. Gains on the total CDS portfolio are the sum of mark-to-market changes and accrued premiums across all positions. In Section 2 of the Online Appendix, I document a straightforward way to use observed CDS spreads at time t to mark each position to market. Because of the computational challenge of valuing a majority of the positions in the U.S. market, I opt for the simplest possible methodology. I denote the total value of the CDS portfolio by V_{ct}, where the subscript c designates a counterparty. $V_{ct} > 0$ means counterparty c has made money on its outstanding positions as of time t.

Next, I translate dollar gains V_{ct} to portfolio returns. $V_{c,t}$ is computed by summing across positions that were initiated at different points in time, so computing a return from $V_{c,t}$ is

[27] This assumption is also motivated by my analysis of the 2011 Japanese tsunami in Section 6. In that case, I use the net amount sold on Japan to measure the effect of the tsunami on risk bearing capacity. A likely channel through which the net amount sold on Japanese firms relates to risk bearing capacity is via mark-to-market losses from increased Japanese CDS spreads. I provide some additional evidence of this claim by confirming that the magnitudes of mark-to-market losses induced by the tsunami were reasonably large.

not straightforward. In Appendix F.3, I show how to convert the dollar gain of each position in a portfolio to the annualized return of the entire portfolio. Intuitively, it is the sum of all annualized changes in market value (e.g. the capital gain) plus any accrued premiums, divided by the total initial investment amount.

My proxies for RB_t^s and RB_t^b are the total annualized CDS portfolio return of the top five sellers and buyers in the market.[28] The top five buyers and sellers are determined using the aggregate market share measure from Section 3.2. I treat the top five sellers as if they are a single trader, and compute the return of their combined portfolio. I do the same for the top five buyers. See Appendix F.3.2 for a formal definition. In this appendix, I also explain why this definition inherently assumes the leverage taken by the top five sellers (or top five buyers) is effectively the same. My analysis of the counterparty types of the largest sellers and buyers in Section 3.3 substantiates this assumption, because the mega-players are relatively homogenous in type and size.

Finally, with some abuse of notation, I compute RB_t using one-week *lagged* annualized portfolio returns. This is one way in which I try to avoid concerns of simultaneity (i.e. losses are high because credit risk prices are high, not vice versa).

4.2 Benchmark Panel Regression Results

With measures of RB_t^s and RB_t^b properly defined, I am now poised to test the benchmark specification in Equation (4), which I repeat for convenience:

$$\log(CDS_{rt}) = a_r + \phi_1 \log(LGD_{rt}) + \phi_2 \log(EDF_{rt}) + \gamma_s RB_t^s + \gamma_b RB_t^b + \boldsymbol{\beta}'\mathbf{X}_t + \varepsilon_{rt}$$

The two coefficients of interest are γ_s and γ_b. γ_s tells us how the premium for bearing default risk changes in response to portfolio performance at the mega-sellers in the economy. We therefore expect $\gamma_s < 0$ because if mega-sellers have lost money, then their depressed risk bearing capacity should increase the premium they charge for selling CDS protection.

γ_b is the elasticity of CDS spreads with respect to losses at mega-buyers of CDS in the economy. The expected sign of γ_b is less clear. When a buyer of protection makes money on his CDS portfolio, he simultaneously loses money on this bond portfolio. The overall magnitude and direction of the portfolio change is determined by the degree to which the bond portfolio is hedged.

[28] In the Online Appendix, I also consider alternative definitions of risk bearing capacity that derive from $V_{C,t}$ that qualitatively deliver similar results.

Table 1 contains the results of running variants of my benchmark regression. For this set of regressions, I present the results using Markit's estimate of LGD. In the Online Appendix, I present the same set of results using Moody's sectoral LGD. The results are not dependent on either measure. All regressions have a reference entity fixed effect to absorb any unobservable time-variant firm characteristics.

Column (1) is the regression with only reference entity fixed effects, Moody's EDF, and Markit's LGD. Naturally, both variables are positive and statistically significant. The coefficient on EDF is lower than its expected value of one, but this is plausibly explained by measurement error.[29] Markit's LGD always has a coefficient near one, which is expected based on the motivating examples in Appendix F.1.

Column (2) is the same regression as column (1), with time fixed effects. This is also the specification that delivers my estimated price of credit risk, $\widehat{\pi}_t$. The within-group R^2 indicates that aggregate variables (e.g. common in the cross-section of r) can explain at most 51 percent of time series variation in log-credit spreads within a reference entity.

Column (3) is the regression that includes only the (lagged) portfolio performance of the top five sellers and buyers. These regression results are the first piece of evidence that losses at mega-sellers are important for the pricing of credit risk. Without other macroeconomic controls, a one standard deviation portfolio loss to mega-sellers results in a 5.0 percent increase in CDS spreads. Keep in mind the frequency of these regressions is weekly, so a 5.0 percent spread movement is large. On the other hand, mega-buyer losses do not appear to affect default risk premiums, a finding that is also consistent with the natural experiment I analyze later in Section 6.

Column (4) adds a set of control variables that may also drive the default risk premiums, and if ignored, would contaminate identification of $\gamma_{s,b}$ through an omitted variable bias. To be clear, I am not asserting that the inclusion of these variables perfectly addresses other identification issues. At the very least though, it should improve the accuracy of the point estimates in the regression. The list of variables I consider is as follows: the log of the earnings to price ratio for the aggregate stock market captures the price of equity risk; the VIX index encompasses economy-wide attitudes towards risk; the TED spread captures the credit riskiness of banks; and the CFNAI index represents general macroeconomic conditions. Based on the previous literature on credit spreads, I also include: the slope of the treasury yield curve to capture expectations of future interest rates; the level of the constant maturity 10 year treasury yield; and the CBOE option-implied skew index to capture jump risks in

[29] Gilchrist and Zakrajsek (2012) obtain similar coefficient estimates.

Table 1: What Determines Default Risk Premiums?

Dependent Variable	$\log(CDS_{rt})$						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
RB_t^s			-0.050	-0.014	-0.011	-0.014	-0.009
			(-4.3)**	(-3.1)**	(-2.2)**	(-2.3)**	(-2.0)**
RB_t^b			0.006	0.004	0.002	0.008	0.005
			(0.7)	(0.9)	(0.5)	(1.6)	(1.4)
$\log(EDF_{rt})$	0.239	0.243	0.215	0.229			
	(15.6)**	(9.6)**	(14.1)**	(11.4)**			
$\log(LGD_{rt})$	0.994	0.895	1.016	0.940			
	(2.4)**	(2.2)**	(2.4)**	(2.3)**			
$\log(E_t/P_t)$				0.327	1.007	0.955	0.590
				(2.0)**	(7.8)**	(7.0)**	(4.5)**
VIX_t				-0.000	0.028	0.043	0.032
				(-0.0)	(4.9)**	(6.1)**	(5.6)**
TED_t				0.046	0.036	0.044	0.035
				(10.1)**	(7.5)**	(7.7)**	(6.8)**
10yr Treasury Yield				-0.057	-0.003	-0.038	-0.022
				(-2.8)**	(-0.1)	(-1.6)	(-1.1)
10yr-2yr Treasury				-0.028	-0.003	0.019	0.016
				(-1.4)	(-0.1)	(0.7)	(0.8)
CFNAI				-0.002	-0.011	-0.004	-0.003
				(0.6)	(3.2)**	(-1.0)	(-1.0)
Option-Skew				-0.004	-0.003	0.010	0.011
				(-1.0)	(-1.0)	(1.9)*	(2.6)*
log (EDF-I-CDS)					0.398		0.397
					(11.3)**		(8.7)**
log (Option-I-CDS)						0.094	0.053
						(4.8)**	(4.4)**
Time Fixed Effect	No	Yes	No	No	No	No	No
Total N	66,191	66,191	66,191	66,191	65,868	31,482	31,465
Within R^2	40.0	51.0	34.4	48.6	49.9	29.6	46.2
Cluster		(r,t)	(r,t)	(r,t)	(r,t)	(r,t)	(r,t)

Notes: The table above contains the results from the regression in Equation (4). LGD_{rt} and EDF_{rt} are proxies for the risk-neutral loss-given default and expected default frequency for firm r, respectively; they come from Markit and Moody's. The control variables include: the aggregate log E-P ratio, VIX, TED spread, 10 yr Treasury yield, slope of the Treasury yield curve, CFNAI index, CBOE's Option Skew index. Except for the E-P ratio, all aggregate controls have been standardized to be mean zero and variance one. Additional controls are Moody's EDF-implied-CDS measure and the Carr and Wu (2012) option-implied-CDS measure. Observations are weekly and span March 2010 - June 2014. Standard errors are double clustered by reference entity and time. **,* indicates 5 percent and 10 percent significance, respectively.

the economy.

Most of the estimated coefficients behave as expected. The positive coefficient on the (log) earnings-to-price ratio indicates the credit risk prices moves in tandem with equity risk prices. The estimated coefficient on the VIX suggests it is not important in determining default risk premia, but the other specifications point to the contrary. The likely culprit, then, is the VIX is highly correlated with the other regressors and its statistical significance is dampened in this particular specification. The TED spread has a strong effect on CDS spreads, consistent with my finding that dealers play a principal role as ultimate sellers of CDS protection.

The coefficients on the remaining macroeconomic variables are also broadly in line with previous research. Like in Duffee (1998), column (4) indicates when the risk free rate rises, credit spreads fall. Increases in the slope of the Treasury yield curve also point to decreases in credit spreads, echoing the findings of Collin-Dufresne, Goldstein, and Martin (2001). Though not statistically significant, increases in the CFNAI index result in lower CDS spreads. As general economic conditions improve, one expects credit spreads to fall. The last macroeconomic control, the S&P 500 skewness measure, has a coefficient that is effectively zero and non-significant.

The first four rows of column (4) contain the primary takeaway from my benchmark panel regression. Unsurprisingly, losses to mega-buyers continue to have no effect on CDS premiums. However, losses to mega-sellers remain important for the pricing of credit risk. Even after the inclusion of the control variables, the coefficient γ_s on RB_t^s is economically large and highly significant. Relative to column (4), the absolute magnitude of γ_s drops to 1.4 percent, but this not surprising because the macroeconomic controls are correlated with mega-seller portfolio performance. To this end, the estimated coefficient in column (4) likely understates the true magnitude because some of the variables, like the TED spread, are more likely to be caused by losses to mega-sellers, not vice versa. If the specification in column (4) omits the TED spread, the estimated γ_s rises to about 2 percent in absolute value.

Model Implied CDS Spreads Columns (5) through (7) of Table 1 replace the LGD and EDF variables with direct estimates of CDS spreads. I include alternative measures of CDS spreads mainly because they should also contain any potential omitted variables. For instance, I employ a measure akin to Carr and Wu (2011), who use deep out of the money options to derive an implied CDS spread. I describe the entire process in Appendix F.4, and

ultimately I compute an "option-implied CDS spread" for each firm.[30] The use of options to infer credit spreads is useful for my purposes because the assumptions necessary to do so are less restrictive than computing implied CDS spreads through a specific structural model. Moreover, the options market and CDS market should, in theory, embed the same information regarding the firm and macroeconomy. Any differences between the two is then specific to either options market or CDS markets. The specific feature of the CDS market that I am trying to highlight is mega-seller losses.

The second measure is also provided by Moody's and is an "EDF implied CDS spread" that combines the EDF and LGD variables into a single credit spread. Because Moody's LGD and EDF measure are in the $\mathbb{P}$-measure, Moody's changes measure to arrive at an implied CDS spread by incorporating their estimate of the market price of risk.[31] The Moody's EDF implied CDS spread comes with potential issues. Broadly speaking, the market price of risk estimate is fit using observed bond prices. To the extent that the mechanism I am trying to highlight is driving this market price of risk, then Moody's implied EDF will, in some ways, subsume the effect.

Column (5) of Table 1 confirms my conjecture that the EDF-implied CDS spread might wash out some of the effects of seller losses. Still, the coefficient on RB_t^s remains reasonably large in economic terms, and is still statistically significant at a 5 percent confidence level. The estimated coefficient on the EDF-implied-CDS measure is highly significant, and it is interesting to note this specification delivers a within-group R^2 that is the closest to its maximum of column (2).

Column (6) uses the option-implied CDS spread instead of the EDF-implied measure. The first thing to notice is the sample size drops by over a factor of two, owing to an imperfect match between Markit CDS spreads and OptionMetrics put prices. The option-implied CDS spread coefficient is nonetheless significant, though much less than one. Carr and Wu (2011) find a similar downward bias, which they attribute to measurement error.[32] The effect of seller losses remains stable and significant under this specification as well.

Finally, column (7) is a regression that includes all control variables, as well as both

[30] The approach of Carr and Wu (2011) assumes only the existence of a "default corridor" for the underlying equity price process of the firm. When the equity price hits the upper bound of the default corridor, the firm is assumed to default, after which the equity price never rises above the lower bound of the default corridor. Under these assumptions, the authors demonstrate how one can use deep out of the money put options to recover the implied CDS spread, and they find their method to be quite accurate for their sample of firms.

[31] Moody's EDF Implied CDS

[32] They do a standard measurement error correction to obtain a better point estimate. I do not follow suit, since I am less interested in the magnitude of the coefficient.

measures of implied CDS spreads. Interesting, both the EDF and option implied CDS spreads display explanatory power, thereby implying the two contain different information. Naturally, the coefficient on RB_t^s dips in magnitude, but it is still significant at a 5 percent level.

To summarize, the results in Table 1 suggest losses to the mega-sellers in CDS markets are an important determinant of the price of credit risk. A one standard deviation decrease in the annualized CDS portfolio of these mega-sellers results in a increase in CDS spreads of 1.4-2 percent. As mentioned, Appendix C also uses bond yield data to verify that mega-seller losses are affecting the price of credit risk, and not the CDS-bond basis.

4.3 Shocks to Mega-Sellers and the Price of Credit Risk: A Simple VAR Analysis

To compliment the findings from the previous subsection, I now turn my attention to isolating the impact of a mega-seller portfolio shock on the price of credit risk. A simple way to trace out the impact of a mega-seller shock is through a vector autoregression (VAR). Section 6 explores a better identified shock in an isolated setting, whereas the VAR I use here allows for a more general interpretation. I specify the VAR as follows:

$$\begin{aligned} \boldsymbol{y}_t &= \boldsymbol{c} + \boldsymbol{\Phi}\boldsymbol{y}_{t-1} + \boldsymbol{\varepsilon}_t \\ \boldsymbol{y}_t &:= \begin{bmatrix} \mathbf{X}_t & RB_t^b & RB_t^s & \widehat{\pi}_t \end{bmatrix}' \end{aligned}$$

where $\widehat{\pi}_t$ is an estimate of the price of credit risk from the panel of CDS spreads used in my benchmark regression in Section 4.2. As previously mentioned, Appendix F.2 provides more details of how I estimate $\widehat{\pi}_t$, which is closely related to the excess bond premium of Gilchrist and Zakrajsek (2012). Loosely speaking, it captures the component of credit spreads that is not explainable by cyclical movements in the likelihood of default; hence, I refer to it as the price of credit risk. $\mathbf{X}_t$ are the following macroeconomic controls (in this order for the VAR): 10 year Treasury yield, 10 year minus 2 year Treasury yield, and the CRSP value-weighted excess market return.[33] I choose the lag length of the VAR based on standard information

[33]The reader might notice that I use slightly different controls from the benchmark panel regression. I use the CRSP value-weighted market return instead of the aggregate price-divided ratio in order to make the results more comparable to Gilchrist and Zakrajšek (2012). I exclude the VIX and TED spread because it is unclear how exogenous these variables are. In unreported results, I rerun the VAR with all control variables and find the main conclusions to be largely unchanged. Also, the ordering of the macroeconomic controls does not materially affect the results.

Figure 7: Impulse Response of $\widehat{\pi}_t$ to a Negative Shock to RB_t^s

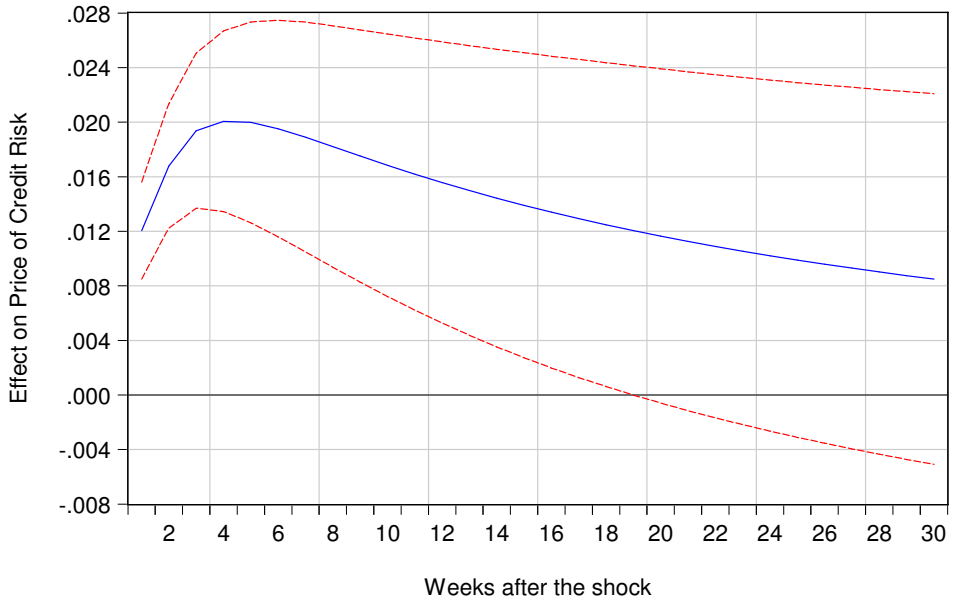

Notes: This figure plots the impulse response of a negative shock to RB_t^s on the price of credit risk. The VAR was orthogonalized using the Cholesky decomposition with the following ordering: (i) 10 year Treasury yield, (ii) 10 year minus 2 year Treasury yield, (iii) the CRSP value-weighted excess market return, (iv) RB_t^b, (v) RB_t^s, and (vi) the price of credit risk $\widehat{\pi}_t$. $RB^{s,b}$ are the risk-bearing capacity of mega-sellers and mega-buyers, respectively. Mega-sellers are the top five aggregate sellers by market share, which is defined in Equation (3). Buyers are analogously defined. RB is the total annualized return of the mega-players' CDS portfolios. The system was estimated using weekly data from February 2010 to June 2014.

criterion.

Next, I compute the impulse response of the price of credit risk to a negative shock to the mega-sellers' CDS portfolio. Because the VAR residuals are correlated, the ordering of the system is important. I treat the shocks to the macroeconomic controls as the most exogenous, and the price of credit risk as the most endogenous. Figure 7 displays the estimated impulse response, along with 95 percent confidence bands.

The impulse response function captures the essence of this paper: a shock to the risk bearing capacity — measured by the CDS portfolio return — of mega-sellers pushes up the price of credit risk. The magnitude of the impulse response is also in line with the results from Section 4.2. A one standard deviation orthogonalized shock to RB_t^s immediately results in a 1.2 percent increase in the price of credit risk. Due to the persistence of the variables, a negative mega-seller shock takes almost 19 weeks to dissipate. Additionally, in Appendix D

I use the VAR to show that RB_t^s Granger causes the price of credit risk, but not vice versa. Given the findings of Gilchrist and Zakrajšek (2012) that link the price of credit risk to real economic outcomes, the effect of mega-seller shocks on the price of credit risk is all the more substantial.

Finally, I regress — in first differences – the price of credit risk on all macroeconomic controls, RB_t^b, and RB_t^s. The incremental R^2 in the regression from including RB_t^s is 12 percent. Prima facie, this means about one-eighth of fluctuations in the aggregate price of credit risk can be explained by five mega-sellers of CDS protection.[34]

5 Robustness

5.1 Additional Econometric Considerations

In the Online Appendix, I address a number of econometric issues that may affect the panel regression results of Section 4.2. None of these exercises affect the conclusions drawn in the main text. In particular, I explore the following situations.

CDS Spreads Are Unit Roots I consider the possibility that CDS spreads are unit roots, which I test formally. Because the typical statistical tests for unit roots have lower power when the true process is a near unit-root, I consider near unit-root cases separately below. In the case of unit roots, the classic spurious regression problem of Granger and Newbold (1974) may distort parameter estimates and statistical inference. I address this issue by recasting my hypothesis in first differences. I find negative portfolio returns by mega-sellers are followed by subsequent increases in CDS spreads. The magnitude and statistical significance of the estimated effect of mega-seller losses on CDS spreads is also quantitatively similar. This is true even if I add lagged CDS spread changes.

CDS Spreads Are Highly Persistent Indeed, CDS spreads are very persistent. It is quite common for first order autocorrelations to be above 0.9. In this case, OLS provide consistent parameter estimates, but the standard errors will be severely biased. I tried to address this issue using double clustered standard errors in the main regression specification of the paper, which is designed to account for autocorrelation within a reference entity.

[34] This regression now includes the VIX and the TED spread. A variance decomposition delivers similar magnitudes in terms of the amount of forecast error variance explained by the RB_t^s.

Nonetheless, with near unit root processes, this form of clustering has been met with some degree of skepticism. To address this issue, I use GLS to estimate the panel by explicitly specifying an autocorrelated error term. Naturally this alters the point estimates, but qualitatively, the significance of variables remains unchanged.

5.2 Are Losses at the Mega-Players Really so Special?

To emphasize it is really the large players in CDS markets that drive pricing, I conduct a placebo test. The benchmark panel regression in Section 4 examined how, after controlling for macroeconomic and firm-specific variables, CDS portfolio performance of the top five aggregate buyers and sellers affected CDS spreads. In this section, I create two additional variables: the annualized CDS portfolio return of the median seller and median buyer. If the large traders in CDS markets are really the relevant ones for pricing, then the inclusion of the portfolio performance of the median buyer and seller should not alter the results.

At each point in time, I identify the median seller and buyer using the market share measure, $MS_S(c,t)$, of Section 4. The median seller is chosen from the set of aggregate sellers, or those with $MS_S(c,t) > 0$. Similarly, the median buyer is chosen from the set of aggregate buyers, or those with $MS_S(c,t) < 0$.[35] I denote the lagged portfolio performance (my proxy for risk bearing capacity) of the median seller and buyer by RB_t^{ms} and RB_t^{mb}, respectively. Table 2 contains the results from including these variables in the regression in equation (4), plus the usual set of macroeconomic controls.

From column (1) it is clear the risk bearing capacity of the median buyer and seller has no discernible affect on the pricing of CDS contracts. Column (2) uses the risk bearing capacity of the median buyer and seller, along with the risk bearing capacity of mega-buyers and sellers. The effect of mega-seller losses is still roughly 1.4 percent, and the CDS portfolio performance of the mega-buyers remains insignificant. These findings corroborate the assertion that it is really the large players who are important for credit risk pricing.

5.3 Identification: Losses Transmit Across Mega-Sellers' Portfolios

I now design a specification similar in spirit to Froot and O'Connell (1998), who show that limited risk bearing capacity of mega-sellers of catastrophe reinsurance affects the pricing

[35] Because the leverage and gross amounts traded are vastly different between large and small players, I exclude observations where the absolute value of the median players' annualized portfolio return is less than 30 percent in order to avoid outliers.

Table 2: Placebo Test Using the Median Seller and Buyer in CDS

Dependent Variable	$\log(CDS_{r,t})$	
	(1)	(2)
RB_t^{ms}	-0.005	-0.003
	(-1.0)	(-0.5)
RB_t^{mb}	-0.003	-0.007
	(-0.6)	(-1.3)
RB_t^{s}		-0.014
		(-2.7)**
RB_t^{b}		0.002
		(0.6)
Control Variables	Yes	Yes
RE Fixed Effect	Yes	Yes
Total N	60,066	60,066
Cluster	(r,t)	(r,t)

Notes: This table presents the results of regressing log-CDS spreads on: (i) control variables, found in Section (4.2); (ii) the lagged total annualized portfolio return of the median seller and buyer in the CDS market (RB_t^{ms} and RB_t^{mb}); (iii) the total annualized portfolio return of the top five sellers and buyers in the CDS market (RB_t^{s} and RB_t^{b}). The top five sellers and median seller (or buyers) are determined using the market share measure in Equation (3). Furthermore, the median seller is chosen from the set of sellers with a non-negative market share. Similarly, buyers are the set with a non-positive market share. 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.

of this insurance. For example, their identification technique exploits how a hurricane in Florida might causing prices to rise for freeze damage insurance in New England. The following specification implements a similar concept in the context of CDS markets:[36]

$$\log(CDS_{rt}) = a_r + \phi_1 \log(LGD_{r,t}) + \phi_2 \log(EDF_{rt}) + \gamma_s RB^s_{-r,t} + \gamma_b RB^b_{-r,t} + \alpha_t + \varepsilon_{rt} \quad (5)$$

where $RB^s_{-r,t}$ is the portfolio performance of r's top five sellers, *excluding reference entities in the same industry as* r. The inclusion of this variable captures the idea that when r's sellers lose money on other reference entities, their limited risk bearing capacity results in an increase in the price of insurance for r. The same logic applies to $RB^b_{-r,t}$. α_t is a time fixed-effect, and controls for any macroeconomic factors that may affect CDS spreads. Table

[36]The two approaches are not, however, directly comparable. Continuing with the hurricane example, Froot and O'Connell (1998) 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 run my specification by including the losses of the top five buyers of r's CDS specifically. When r's CDS buyers have lost money, they shift their demand. I find the effect of mega-seller risk bearing capacity to be basically the same.

Table 3: Do Losses to Mega-Players Transmit Across Their Portfolio?

Dependent Variable	$\log(CDS_{r,t})$	
	(1)	(2)
$RB^s_{-r,t}$	-0.014	-0.015
	(-3.1)**	(-1.7)*
$RB^b_{-r,t}$	0.004	0.008
	(1.1)	(1.8)*
EDF and LGD	Yes	Yes
RE FE	Yes	Yes
Macro Controls	Yes	No
Time FE	No	Yes
Total N	64,647	64,647
Cluster	(r,t)	(r,t)

Notes: This table presents the results of the regression: $\log(CDS_{rt}) = a_r + \phi_1 \log(LGD_{r,t}) + \phi_2 \log(EDF_{rt}) + \gamma_s RB^s_{-r,t} + \gamma_b RB^b_{-r,t} + \alpha_t + \varepsilon_{rt}$. $RB^s_{-r,t}$ is the total annualized return of the mega-sellers' CDS portfolio, excluding reference entities in r's industry. $RB^b_{-r,t}$ 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.

3 contains the results.

The important variable is γ_s, the coefficient on $RB^s_{-r,t}$. A one standard deviation loss to the of mega-sellers' portfolio, excluding reference entities in the same industry as r, results in a 1.5% increase in the CDS spread for reference entity r. The magnitude of the effect is robust across two specifications. In column (1), I use the same macroeconomic controls described in Section 4 instead of a time fixed effect. In this case, the coefficient on $RB^s_{-r,t}$ is also significant at a 1 percent confidence level. In column (2), I replace the macroeconomic controls and use a time fixed-effect. This controls for all unobservable macroeconomic variables that may affect credit risk pricing. The point estimate is now statistically significant only at a 10 percent confidence level, but this could be driven by the fact that there is strong correlation between the time fixed-effect and $RB^s_{-r,t}$. Still, these results are very much in line with the benchmark regression in Section 4.

Interestingly, in column (2), the coefficient on $RB^b_{-r,t}$ is positive and significant at a 10 percent confidence level. These results suggest when mega-buyers make money on other parts of their portfolio, they are willing to pay a higher premium for reference entity r. A potential explanation for this result is they are simultaneously losing money on their bond portfolio (excluding reference entities in the same industry as r), and wish to hedge

themselves against future increases in the credit risk of r. These results are, however, to be interpreted with some caution because: (i) the coefficient is not significant in column (1), and (ii) mega-buyer CDS portfolio performance did not appear to matter for the price of credit risk in the analysis of Section 4.

5.4 Idiosyncratic Shocks and CDS Spread Movements: An Alternative Approach

I chose to present the benchmark specification of Section 4.2 first because it is the simplest way to gain a quantitative sense of how mega-player portfolio performance relates to CDS spreads. Moreover, it allows me to make easily interpretable statements such as how a one standard deviation loss to five players affects default risk premiums. The downside of my benchmark regression is that, despite controlling for a large number of observables, it is does not provide great identification. The analysis of how losses transmit across mega-player portfolios in Section 5.3 and the natural experiment of the 2011 Japanese tsunami in Section 6 provide better identification that support the interpretation of my benchmark results. In this subsection, I consider a more complicated identification strategy that also validates my benchmark results.

My approach in this subsection is related to Gabaix (2012), who examines how idiosyncratic shocks to large firms relates to macroeconomic fluctuations. Similarly, I am going to argue that idiosyncratic shocks to mega-sellers generates CDS spread movements. To start, I assume a simple factor structure for changes in CDS spreads:

$$\Delta \log(CDS_{rt}) = a + \mathbf{\Lambda}_r^{'}\mathbf{F}_t + e_{rt} \tag{6}$$

From an economic standpoint, one can view the factors $\mathbf{F}_t$ as macroeconomic variables important for credit spread movements (e.g. the return on the aggregate stock market index). Instead of taking a stance on what these factors should be, I extract the first five principal components from a large panel of CDS spreads.[37] I choose five principal components because this the number of factors necessary to capture over 80% of variation in spread changes, but the number of factors is not particularly important for my analysis.

[37] In unreported results, I use economic factors instead of principal components. The factors are the Fama-French factors, Fama-French 5 industry portfolio returns, and changes in the investment grade and high yield CDX indices. The set of CDS spreads I use for principal component regression is a subset of those used in my benchmark regression, where I remove names that do not have a full time series.

Once I extract the factors $\mathbf{F}_t$, I use them to compute idiosyncratic shocks to the dollar value V_{ct} of each counterparty's CDS portfolio via simple time-series regressions:

$$\Delta V_{ct} = c + \boldsymbol{\beta}_c^{'}\mathbf{F}_t + \xi_{ct} \tag{7}$$

I define the idiosyncratic shocks to counterparty c by the regression residuals, ξ_{ct}. It is the portion of c's portfolio changes that cannot be explained by exposure to common factors driving all CDS spread movements.

Next, for each reference entity r I build "buyer and seller granular residuals", defined as:

$$\begin{aligned} \Psi_{rt}^s &:= \sum_{c \in \mathcal{S}_{t-1}(r)} \omega_{r,c,t-1}^s \xi_{ct} \\ \Psi_{rt}^b &:= \sum_{c \in \mathcal{B}_{t-1}(r)} \omega_{r,c,t-1}^b \xi_{ct} \end{aligned} \tag{8}$$

where $\mathcal{S}_{t-1}(r)$ and $\mathcal{B}_{t-1}(r)$ 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.[38] 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 (2012).

Finally, to determine whether idiosyncratic shocks to mega-players impact spread movements, I run the following panel regression:

$$\Delta \log(CDS_{rt}) = \alpha_r + \iota_t + \boldsymbol{\zeta}^{'}\Delta\mathbf{Z}_{rt} + \phi_s\Psi_{rt}^s + \phi_b\Psi_{rt}^b + \varepsilon_{rt} \tag{9}$$

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. The controls I use are Moody's expected default frequency, Markit's loss-given-default, and the lagged log-change in r's CDS spread, $\Delta \log(CDS_{r,t-1})$. Table 4 contains the results of this panel regression.

Consistent with my benchmark panel regression results, the buyer granular residual has no explanatory power for CDS spread movements. On the other hand, negative seller granular residuals correspond to a significant increase — both statistically and economically — in

[38] 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.

Table 4: Explanatory Power of Seller and Buyer Granular Residuals

Dependent Variable	$\Delta\log(CDS_{r,t})$
Ψ^s_{rt}	-.006
	(-4.5)**
Ψ^b_{rt}	.001
	(1.3)
Firm Controls	Yes
RE FE	Yes
Time FE	Yes
Total N	65,577
Within R^2	30.5
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 (9). $\mathbf{Z}_{rt}$ is a vector of reference entity specific variables: Moody's 5 year EDF, Markit's LGD, and the lagged log-change in CDS spread. 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.

CDS spreads. A negative seller granular residual represents adverse idiosyncratic shocks for important protection sellers, so the sign of the point estimate on Ψ^s_{rt} is in line with the other results in my paper. A simple way to understand the magnitudes in Table 4 is to compare reference entities whose seller granular residual is in the 10th percentile and the 90th percentile. Reference entities whose seller granular residual is in the 10th percentile experience a 1.4% spread increase, relative to the 90th percentile.

More importantly, because I control for reference entity fundamentals and the time fixed effect ι_t absorbs all macroeconomic factors, this regression highlights the impact of idiosyncratic seller losses on default risk premiums. Furthermore, the seller granular residual gives added weight to the biggest sellers, so these results also substantiate my claim that it is specifically mega-seller losses that increase default risk premiums.

6 A Natural Experiment: The 2011 Japanese Tsunami

Sections 4 and 5 took a reduced-form approach in showing that the risk bearing capacity of mega players affects the dynamics of the price of credit risk. Because these results are continuous over a long time series, they are useful for forming an aggregate view of what

drives the price of credit risk. For the remainder of the paper, I will provide more refined identification techniques that are simply intended to aid the interpretation that mega-seller losses increase the price of credit risk.

To start, I use a natural experiment to establish a causal link between reduced mega-seller risk bearing capacity and CDS spreads. 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 went from 80 to 115 basis points on the following Monday, a nearly 50 percent increase. Appendix E.1 contains additional background information on the tsunami, and its after-effects.

6.1 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, the expected payouts facing Hedge Fund A increase, and presumably their risk bearing capacity decreases; however, the same does not hold true for Hedge Fund B.[39] My hypothesized mechanism then suggests the U.S. firms for whom Hedge Fund A is large seller will experience increases in their CDS risk premia. On the other hand, U.S firms where Hedge Fund B is a large seller will not see their spreads rise.

In Appendix E.1, I verify that U.S. counterparties had large CDS exposures to Japanese firms. I also show that the tsunami caused non-neglible 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.

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

6.1.1 Measurement

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

$$\begin{aligned} \Gamma_{S,r} &:= \sum_{c \in \mathcal{S}(r)} \left[\frac{NS(c,r)}{NO(r)} \right] \times NS(c, Japan) \\ \Gamma_{B,r} &:= \sum_{c \in \mathcal{B}(r)} \left[-\frac{NS(c,r)}{NO(r)} \right] \times NS(c, Japan) \end{aligned} \tag{10}$$

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.[40] 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 $\Gamma_{S,r}$ 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.[41]

6.1.2 Observable Characteristics of Reference Entities

To develop the identification challenge I face, suppose the change in U.S spreads following the tsunami takes a linear form:[42]

$$\Delta \log(CDS_{r,1}) = a + \phi_1 \Gamma_{S,r} + \phi_2 \Gamma_{B,r} + \boldsymbol{\beta}' \boldsymbol{X}_r + \gamma U_r + \varepsilon_r \tag{11}$$

[40] 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.

[41] 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.

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

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 (11), consistent estimation requires $cor(\Gamma_{S,r}, \widetilde{\varepsilon}_r) = cor\,(\Gamma_{S,r}, U_r) = 0$, 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 8 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 $\Gamma_{B,r}$.[43] 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 8 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 with Japan. Moreover, the average correlation with Japan within any percentile is relatively low, and never reaches above 18%.

What other variables might be latent in U_r? Since certain industries may have been more exposed to Japanese firms, I include each reference entity's NAICS code. It may also be the case that sellers/buyers of Japanese CDS possess certain expertise, and therefore specialize

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

Figure 8: Correlation with Japan Across Deciles of $\Gamma_{S,r}$ and $\Gamma_{B,r}$

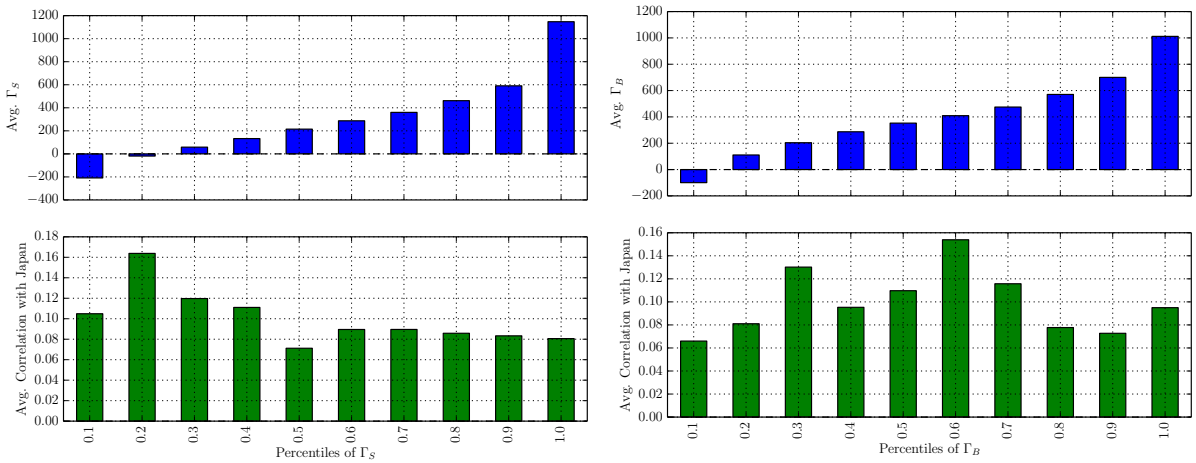

Notes: The left panel in this figure shows the correlation between U.S. reference entity CDS spreads and the country of Japan's (JPN) CDS spread, averaged within deciles of $\Gamma_{S,r}$. $\Gamma_{S,r}$ is the share-weighted average CDS exposure of r's net sellers to Japanese firms. The right panel 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.

in a particular segment of U.S. markets. While the NAICS code controls for this in one dimension, I also include level of CDS spreads for each reference entity on 3/11/2011. That way, I allow for the possibility that Γ captures sellers/buyers who specialize in riskier credits. Moreover, I include the 90-day running 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.

Furthermore, to control for changes in 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. In terms of the regression in Equation (11), it is difficult to imagine an (additional) unobserved characteristic U_r that is (i) orthogonal to the equity returns of the firm and (ii) still causal in terms of spread movements following the tsunami.

In Appendix E.3, 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.

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

Table 5 summarizes the results of running variations of the following regression:

$$\Delta \log(CDS_{r,1}) = a + \phi_1 \Gamma_{S,r} + \phi_2 \Gamma_{B,r} + \boldsymbol{\beta}' X_r + \varepsilon_r \tag{12}$$

Table 5: 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	0.017		0.016
	(2.57)**		(2.63)**	(3.35)**		(2.95)**
$\Gamma_{B,r}$		0.002	0.003	0.002		0.0043
		(0.35)	(0.70)	(0.53)		(0.93)
$\Gamma_{S,r}^{placebo}$					0.01	0.00
					(1.49)	(0.63)
$\Gamma_{B,r}^{placebo}$					-0.00	-0.00
					(-0.81)	(-0.83)
Control Variables	No	No	No	Yes	Yes	Yes
Total N	288	288	288	288	288	288
Adj. R^2	1.9%	0%	2.4%	26.6%	24.4%	26.9%

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 ar (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, the (log) 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.

Consistent with the results in Section 4, 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 spreads increase 2.5% more than 10th percentile in the week following the tsunami.

One of the important stylized facts I uncovered in Section 3 was a common set of sellers dominate U.S. CDS markets. To tie the current exercise to these results, I construct placebo versions of $\Gamma_{S,r}$ ($\Gamma_{B,r}$) by taking simple averages of seller (buyer) exposures to Japan:

$$\Gamma_{S,r}^{placebo} := \sum_{c \in \mathcal{S}(r)} \left[\frac{1}{\|\mathcal{S}(r)\|} \right] \times NS(c, Japan) \tag{13}$$

$$\Gamma_{B,r}^{placebo} := \sum_{c \in \mathcal{B}(r)} \left[\frac{1}{\|\mathcal{B}(r)\|} \right] \times NS(c, Japan) \tag{14}$$

where the $\|\cdot\|$ operator denotes the size of a set. It is important to note that, for instance, $\Gamma_{S,r}$ is computed as a *share-weighted* average exposure of r's sellers to Japanese firms; thus, differences between the effects of $\Gamma_{S,r}^{placebo}$ and $\Gamma_{S,r}$ speak to the role of large sellers in shock transmission.

Columns (5) and (6) in Table 5 confirm the role that large sellers play in the pricing of U.S. CDS. Column (5) is the specification using the placebo versions of Γ, with the full set of controls. Both $\Gamma_{S,r}^{placebo}$ and $\Gamma_{B,r}^{placebo}$ are insignificant in this regression. In column (6), I run the regression with Γ and its placebo versions. The coefficient on $\Gamma_{S,r}$ remains essentially unchanged in magnitude, and is still highly significant. Because the specification in column (6) controls for the average exposure of r's sellers to Japanese firms, it provides strong evidence that it is the shocks to the risk bearing capacity of large sellers that really matter for credit spreads.[44] Furthermore, because equity returns control for changes in the quantity of risk, this identification exercise lends credence to the interpretation that mega-seller losses affect the price of bearing default risk.

[44] As another placebo test of my hypothesis, I run the same set of regressions (unreported) using data spanning between two weeks and one week prior to the tsunami. The main takeaway is none of the Γvariables have explanatory power for the cross-section of spread movements. This finding further supports my hypothesized shock propagation mechanism, given we wouldn't expect exposure of Japanese firms in the week prior to the tsunami to cause a shock to anyone's portfolio.

6.3 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. Within the 288 reference entities I analyze, I split reference entities into two buckets. Reference entities in the "low Γ_S" bucket 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" bucket are those reference entities whose sellers were most exposed to Japan (above the median Γ_S). Within each bucket, I then compute a weighted-average CDS spread. The weights assigned to each reference entity within each bucket are proportional to $\Gamma_{S,r}$. Thus, in the high Γ_S bucket, reference entities whose sellers were most exposed to Japan get the most weight; in the low Γ_S bucket, 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 9 plots the difference in the time series of each bucket'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 occur. Before Japanese officials were able to stabilize the plant, many experts opined a full meltdown would be akin to dropping a nuclear bond 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 E.1, but loosely speaking, for mega-sellers of U.S. CDS, Japanese exposures were roughly 4-6 percent of their overall

Figure 9: How Long Did the Shock of the Tsunami Last?

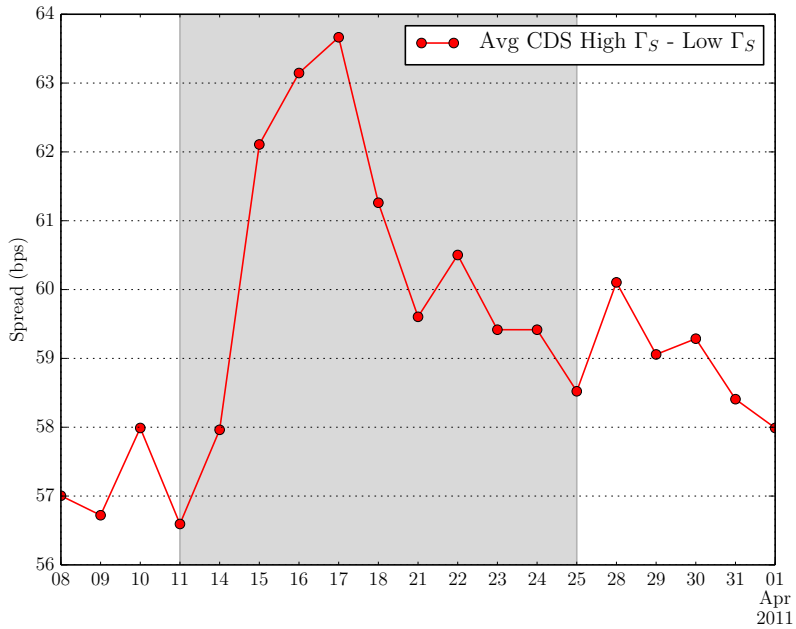

Notes: This figure shows the difference between the weighted-average CDS spread of high $\Gamma_{S,r}$ reference entities and low $\Gamma_{S,r}$ reference entities. High (low) $\Gamma_{S,r}$ reference entities are those whose sellers were most (least) exposed to Japanese firms. The weights assigned to 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.

portfolio. While this certainly non-trivial, one would also not expect the shock of the tsunami 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. The analysis of this section was used to identify and provide evidence a shock to risk bearing capacity moves CDS spreads. The effect of the totality of these shocks on the aggregate pricing of credit risk was examined in Section 4. The results of this section complement — and motivate — the more macroeconomic view taken in that section.

7 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 and mega-buyer 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 or mega-buyers is transmitted to CDS spreads.

7.1 Empirical Strategy

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

$$\log\left(CDS_{it}\right) = g\left(H_t^r, H_t^b, RB_{it}^s; X_{it}\right) \tag{15}$$

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; 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 for the same reasons as in Section 4

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:

$$\log\left(CDS_{it}\right) = \alpha_{rtb} + \eta RB_{it}^s + \varepsilon_{it} \tag{16}$$

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 CDS spreads paid 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.

Note that I can reverse the thought experiment by fixing on the seller, reference entity, and date:

$$\log\left(CDS_{it}\right) = \alpha_{rts} + \eta RB_{it}^b + \varepsilon_{it}$$

where α_{rts} is a reference entity-time-seller fixed effect and RB_{it}^b is the risk bearing capacity of the buyer in the transaction. This regression allows me to isolate the impact of buyer risk bearing capacity on offered CDS spreads. In my empirical work, I run both sets of

regressions.

7.2 Data Description

All my analysis thus far has relied on positional data. The data used in this section come from actual transactions. Transactions represent flows, whereas total CDS positions represent stocks. I still use positional data to compute RB_{it}^{s} as before, whereas I use transactional data to define the dependent variable in my regressions. The transaction information provided by the DTCC contains information on the upfront 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 the 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 (16).

I consider transactions on all types of CDS: single name, index, and tranche. The reason being that my within buyer-reference entity-time estimator already distinguishes product types. However, this is one reason the results I will present are not comparable to those in Section 3.2. The empirical exercises in Section 4 crucially focused only on ultimate single name exposures (i.e. ignoring tranches, and after disaggregating indices). This is because I was interested in: (i) determining ultimate buyers and sellers of CDS protection and (ii) explaining determinants of the aggregate price of credit risk. In the current section, my goal is different, and is designed to highlight my hypothesized causal mechanism.

7.3 Results

Table 6 present the results of the regression of Equation (16). The first two columns run the regression with a buyer-reference entity-time fixed effect. This isolates the effect of seller risk

bearing capacity on offered CDS spreads. I call this the "seller-channel" regression. The last three columns repeat the exercise, but use a seller-reference entity-time fixed effect. This specification isolates the effect of buyer risk bearing capacity on CDS spreads. Similarly, I refer to this as the "buyer-channel" regression.

Table 6: CDS Transaction Prices After Seller and Buyer Losses

Dependent Variable	$\log(CDS_{it})$			
	(1)	(2)	(3)	(4)
RB^s_{it}	-0.019	-0.011		
	(-7.7)**	(-6.9)**		
RB^b_{it}			-0.014	-0.009
			(-7.4)**	(-7.8)**
Transaction Controls	No	Yes	No	Yes
FE	(b, r, t)	(b, r, t)	(s, r, t)	(s, r, t)
Total N	31,872	31,872	33,744	33,744
Within R^2	.1%	20.8%	.08%	23.0%

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 group. *, ** represent statistical significance at a 10 and 5 percent level, respectively.

Column (1) of the table is seller-channel regression with only the risk bearing measure for sellers. These results indicate a one standard deviation increase in seller losses results in an implied CDS spread increase of roughly 2%. Column (2) adds other observable features of each transaction such as the notional, maturity, and fixed coupon. Importantly, the coefficient on RB^s_{it} is still large, negative, and significant after the inclusion of these observable transaction characteristics.[45]

Columns (3) and (4) report the results when fixing the seller-reference entity-trade date. These findings indicate when a seller faces multiple buyers (for the same reference entity and date), buyers whose CDS portfolio performance has been relatively poor are willing to pay a higher premium for protection. The estimated coefficients are also stable with the addition of other transaction features. Lastly, it appears the magnitude of the buyer-channel is (very) slightly less than the seller-channel. These findings might appear to contradict the other results in my paper, which pointed to either no effect or a slightly positive relationship

[45]In unreported results, I instead cluster standard errors by seller. The t-statistic on RB^s_{it} drops to -3.1, but is still significant at a 1% confidence level.

between buyer risk bearing capacity and prices. A subtle point is that these regression results are not a one-to-one comparison with my earlier findings, which I discuss further now.

Interpretation

Relation to Section 4

As previously mentioned, there should be some caution in comparing these results to those of Section 4. For one, notice that the current identification exercise examines situations where a seller faces two buyers (or for the second regression, when a buyer faces two sellers), and asks how risk bearing capacity affects pricing in this isolated setting. My earlier results examine how the risk bearing capacity of the *ultimate* buyers and sellers of CDS (e.g. the mega-players) affects pricing in the aggregate. My transactional study does not require the buyers or sellers themselves to be mega-players, and hence is not directly comparable to my analysis in Section 4.

Furthermore, the current set of results also takes advantage of within buyer-reference entity-date variation. However, pure theoretical asset pricing models with a singular aggregate price of credit risk would never predict this variation to exist in the first place. The Granger causality analysis in Appendix D takes this asset pricing consideration more seriously, at the cost of weaker identification.[46]

In lieu of this discussion, the proper way to view my analysis in this section is that I relax the asset pricing restrictions for the sake of stronger identification. In the current setting, I am able to show that fluctuations in risk bearing capacity do in fact affect prices. Presumably, this more granular effect manifests itself in the aggregate via a higher overall price of credit risk. Of course, a formal theoretical model would be needed to flesh this idea out fully.

An additional wrinkle is when the risk bearing capacity of a seller is low, he may choose to not participate in the market. In this case, we won't observe any transactions with this counterparty, but that does not mean Markit-reported CDS spreads can't be high. This is because the Markit-reported CDS spread is a function of transactions *and* quotes. 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 2011 Japanese tsunami results from Section 6 are more useful for analyzing how a shock to risk-bearing capacity of mega-players affects quotes (i.e. those reported by Markit).

[46] Indeed, it is difficult to construct a perfectly identified test of how mega-seller losses affect the *aggregate* price of credit risk.

Is It Just Counterparty Risk?

A natural set of objections to the results in Table 6 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_{it}^s 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_{it}^s are good counterparties, and thus can earn a higher premium from buyers. Clearly, the results of Table 6 contradict this story.

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

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

where CP_{it}^s 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 $corr(RB_{it}^s, CP_{it}^s) < 0$ because sellers with high risk bearing capacity are not likely to be bad counterparties. Denote the estimated coefficient on RB_{it}^s when omitting CP_{it}^s by $\hat{\eta}$, and the true parameter value by η. Classical OLS theory says if $\gamma < 0$ and $corr(RB_{it}^s, CP_{it}^s) < 0$, then $\eta < \hat{\eta}$. In other words, the true coefficient on RB_{it}^s is even more negative, which suggests sellers with low risk bearing capacity charge even higher premiums than the estimates I've reported.

8 Conclusion

The major contribution of this paper is to demonstrate how financial frictions — in the form capital losses at mega-sellers — affects the pricing of credit default risk. I find mega-seller CDS portfolio losses have a strong and positive effect on default risk premiums, a conclusion that is further supported by corporate bond data. I then conduct a variety of empirical exercises that reinforce the relationship between default risk premiums and mega-seller losses.

My results bring out some important issues for the asset pricing literature on credit risk going forward. To match the time-series dynamics of CDS spreads, it appears necessary to account for fluctuations in risk prices that are driven by frictions, or limited risk bearing capacity, at a small set of sellers. A natural extension of the empirical conclusions in this paper would be to use models like Brunnermeier and Sannikov (2014) and He and Krishnamurthy (2012) to link credit risk pricing with financial frictions. To this end, more empirical work is also needed flesh out the link between losses and risk bearing capacity. In other words, what specific agency problem must be present for losses to affect risk bearing capacity?

A second important contribution of this paper is a new set of facts that characterize the distribution of credit risk in U.S. CDS markets. Put simply, CDS markets are large, concentrated, and common; a few protection buyers, and even fewer sellers, account for large share of the risk transferred by CDS.

These facts raise potential red flags in terms of financial stability and market design. Take, for instance, 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. While 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 possible 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.

A more subtle aspect of my stylized facts is that selling of CDS has partly moved from dealers to hedge funds and asset managers (HFAMs). A possible explanation for this pattern is 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 HFAMs 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 only five sellers in the market can move CDS spreads by a substantial amount. In some ways, I have taken the structure of the CDS market as given, and studied the effect of this structure on pricing. The next step is to understand how and why this market structure arises endogenously.

References

[1] Adrian, T., and Boyarchenko, N. Intermediary leverage cycles and financial stability. *Federal Reserve Staff Report*, 567 (2013).

[2] Adrian, T., Colla, P., and Shin, H. S. Which financial frictions? Parsing the evidence from the financial crisis of 2007-09. Staff Reports 528, Federal Reserve Bank of New York, 2011.

[3] Adrian, T., Etula, E., and Muir, T. Financial intermediaries and the cross-section of asset returns. *The Journal of Finance* (2014), n/a–n/a.

[4] Anton, M., and Polk, C. Connected stocks. *The Journal of Finance 69*, 3 (2014), 1099–1127.

[5] Arora, N., Gandhi, P., and Longstaff, F. A. Counterparty credit risk and the credit default swap market. *Journal of Financial Economics 103*, 2 (2012), 280 – 293.

[6] Atkeson, A., Eisfeldt, A., and Weill, P.-O. The market for otc derivatives. *Working Paper, url: https://sites.google.com/site/andrealeisfeldt/* (2013).

[7] Atkeson, A., Eisfeldt, A., and Weill, P.-O. Entry and exit in otc derivatives markets. *Working Paper, url: https://sites.google.com/site/andrealeisfeldt/* (2014).

[8] Axtell, R. L. Zipf distribution of u.s. firm sizes. *Science 293*, 5536 (2001), 1818–1820.

[9] Bai, J., and Collin-Dufresne, P. The cds-bond basis. *SSRN: http://ssrn.com/abstract=2024531* (2013).

[10] Bai, J., and Wu, L. Anchoring credit default swap spreads to firm fundamentals. *SSRN: http://ssrn.com/abstract=2020841* (2012).

[11] Bao, J., and Pan, J. Bond illiquidity and excess volatility. *Review of Financial Studies* (2013).

[12] Berndt, A., Douglas, R., Duffie, D., Ferguson, M., and Schranz, D. Measuring default risk premia from default swap rates and edfs. *SSRN: http://ssrn.com/abstract=556080* (2008).

[13] Berndt, A., and Obreja, I. Decomposing european cds returns*. *Review of Finance 14*, 2 (2010), 189–233.

[14] Brunnermeier, M. K., and Sannikov, Y. A macroeconomic model with a financial sector. *The American Economic Review 104*, 2 (2014), 379–421.

[15] Campbell, J. Y., and Shiller, R. J. Stock prices, earnings, and expected dividends. *The Journal of Finance 43*, 3 (1988), 661–676.

[16] Carr, P., and Wu, L. A simple robust link between american puts and credit protection. *Review of Financial Studies 24*, 2 (2011), 473–505.

[17] Chen, H., Cui, R., He, Z., and Milbradt, K. Quantifying liquidity and default risks of corporate bonds over the business cycle. *Working Paper, url: http://www.mit.edu/ huichen/bondliquidity.pdf* (2014).

[18] Chen, K., Fleming, M., Jackson, J., Li, A., and Sarka, A. An analysis of cds transactions: Implications for public reporting. *NY Fed Working Paper*, 517 (2011).

[19] Chen, L., Collin-Dufresne, P., and Goldstein, R. S. On the relation between the credit spread puzzle and the equity premium puzzle. *Review of Financial Studies 22*, 9 (2009), 3367–3409.

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

[21] Choi, I. Unit root tests for panel data. *Journal of International Money and Finance 20*, 2 (April 2001), 249–272.

[22] Collin-Dufresne, P., Goldstein, R. S., and Martin, J. S. The determinants of credit spread changes. *The Journal of Finance 56*, 6 (2001), 2177–2207.

[23] Driessen, J. Is default event risk priced in corporate bonds? *Review of Financial Studies 18* (2005), 165–195.

[24] Duffee, G. R. The relation between treasury yields and corporate bond yield spreads. *The Journal of Finance 53*, 6 (1998), 2225–2241.

[25] Duffie, D. Credit swap valuation. *Financial Analysts Journal* (1999), 73–87.

[26] Duffie, D., Li, A., and Lubke, T. Policy perspectives on otc derivatives market infrastructure. *Federal Reserve Staff Report*, 424 (2011).

[27] Duffie, D., Scheicher, M., and Vuillemey, G. Central clearing and collateral demand. *Forthcoming: Journal of Financial Economics* (2014).

[28] Duffie, D., and Singleton, K. Modeling term structures of defaultable bonds. *Review of Financial Studies 12*, 4 (1999), 687–720.

[29] Elton, E. J., Gruber, M. J., Agrawal, D., and Mann, C. Explaining the rate spread on corporate bonds. *The Journal of Finance 56*, 1 (2001), 247–277.

[30] Ericsson, J., Jacobs, K., and Oviedo, R. The determinants of credit default swap premia. *Journal of Financial and Quantitative Analysis 44*, 01 (2009), 109–132.

[31] Froot, K. A., and O'Connell, P. G. J. The Pricing of U.S. Catastrophe Reinsurance. 195–232.

[32] Froot, K. A., Scharfstein, D. S., and Stein, J. C. Risk management: Coordinating corporate investment and financing policies. *The Journal of Finance 48*, 5 (1993), 1629–1658.

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

[34] Gabaix, X. The granular origins of aggregate fluctuations. *Econometrica 79*, 3 (2011), 733–772.

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

[36] Greenwood, R., and Thesmar, D. Stock price fragility. *Journal of Financial Economics 102*, 3 (2011), 471 – 490.

[37] He, Z., and Krishnamurthy, A. Intermediary asset pricing. *American Economic Review 103*, 2 (2013), 732–70.

[38] Jarrow, R. A., and Turnbull, S. M. Pricing derivatives on financial securities subject to credit risk. *The Journal of Finance 50*, 1 (1995), 53–85.

[39] Kahneman, D., and Tversky, A. Choices, values, and frames. *American psychologist 39*, 4 (1984), 341.

[40] Khwaja, A. I., and Mian, A. Tracing the impact of bank liquidity shocks: Evidence from an emerging market. *American Economic Review 98*, 4 (2008), 1413–42.

[41] Ljungqvist, L., and Sargent, T. Recursive macroeconomic theory. *The MIT Press; 2nd edition* (2004).

[42] Maggiori, M. Financial intermediation, international risk sharing, and reserve currencies. *Unpublished manuscript, UC Berkeley* (2011).

[43] Oehmke, M., and Zawadowski, A. The anatomy of the cds market. *SSRN: http://ssrn.com/abstract=2023108* (2013).

[44] Oehmke, M., and Zawadowski, A. 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).

[45] Schachar, O. Exposing the exposed: Intermediation capacity in the credit default swap market. *NY Fed Working Paper* (2012).

[46] Tang, D., and Yan, H. What moves cds spreads? *Working Paper: http://www.sef.hku.hk/ yjtang/TangYan-CDSMove_21Jun2013.pdf* (2013).

A Appendix: Additional Tables

Table 7: Cross-Sectional Distribution of CDS Overlap

		Buyer Overlap					*Seller Overlap*				
			Percentile					Percentile			
Year	**Avg. # Pairs**	**Mean**	**25**	**50**	**75**	**95**	**Mean**	**25**	**50**	**75**	**95**
ALL	651,845	0.223	0.010	0.169	0.371	0.667	0.269	0.008	0.201	0.467	0.774
2010	777,113	0.201	0	0.137	0.339	0.623	0.252	0	0.165	0.448	0.764
2011	693,022	0.219	0	0.159	0.363	0.663	0.270	0	0.200	0.475	0.779
2012	620,263	0.237	0.009	0.187	0.391	0.685	0.278	0.008	0.216	0.478	0.775
2013	526,439	0.251	0.027	0.191	0.387	0.694	0.278	0.023	0.220	0.466	0.777
2014	487,907	0.145	0.044	0.208	0.396	0.697	0.279	0.027	0.201	0.459	0.775

Notes: This table presents summary statistics for the distribution of buyer and seller overlap measures. The pairwise seller (buyer) commonality between reference entities i and j is the proportion net sellers (buyers) who are common to i and j. Observations are grouped together by year and exclude CDS written on Mortgage Backed Securities (see Data Appendix 1.4.2).

B Appendix: Complimentary Approach to Quantifying Concentration

In this appendix, I use an alternative set of metrics that quantify the concentration of CDS buyers and sellers of protection. Namely, I look at concentration within a reference entity, and then commonality of buyers and sellers across reference entities.

B.1 Credit Risk is Highly Concentrated within a Reference Entity

To quantify the concentration of buyers and sellers within a reference entity, I use concentration ratios, or the share of total net selling (buying) by the largest sellers (buyers). To fix ideas, I define the top five sellers of a reference entity r on date t as the set $TS(r,t)$. Similarly, I define the top five buyers as $TB(r,t)$. For each reference entity and date, seller and buyer concentration ratios are then given by:

$$\begin{aligned} CR_S(r,t) &= \sum_{c \in TS(r,t)} \frac{NS(c,r,t)}{NO(r,t)} \\ CR_B(r,t) &= \sum_{c \in TB(r,t)} \frac{|NS(c,r,t)|}{NO(r,t)} \end{aligned} \tag{17}$$

In the second line, I use the absolute value operator in the numerator because, by my definition, a net buyer overall has a negative net exposure.

Because there are thousands of reference entities traded, I aggregate $CR_S(r,t)$ or $CR_B(r,t)$ in two ways: (i) an equally weighted average of concentration across reference entities, and (ii) a size-weighted average of concentration across reference entities, where size is measured in the amount of net notional outstanding, $NO(r,t)$. The reason for the latter is very small reference entities usually have only one or two sellers and buyers, which means the concentration indices for these reference entities are close to 100 percent. A size-weighted average properly offsets this effect. Formally, I define size-weighted aggregate seller concentration index as follows:

$$CR_S^{SW}(t) = \frac{NO(r,t)}{NO(t)} \sum_{r \in \mathcal{R}_t} CR_S(r,t)$$

Size-weighted aggregate buyer concentration ratios are defined analogously. The equal-weighted counterparts are also self-explanatory. Figure 10 plots these various aggregate concentration ratio indices through time. In addition, this plot also displays the equal-weighted average buyer and seller concentration for the top 100 largest reference entities by net notional outstanding.

Figure 10 shows credit risk in CDS is highly concentrated for both buyers and sellers. Through the sample and for the average-sized reference entity, five sellers account for nearly 73 percent of the total amount of selling in CDS markets. Similarly, five buyers account for almost 57 percent of all buying. Sellers are also always more concentrated than buyers. This is true when considering reference entities equally, weighting them by size, or focusing on the 100 largest reference entities. Using the size-weighted indices as a benchmark, sellers are nearly 1.3 times as concentrated as buyers in CDS markets.

Figure 10: How Much Risk is Bought/Sold by the Top 5 Counterparties?

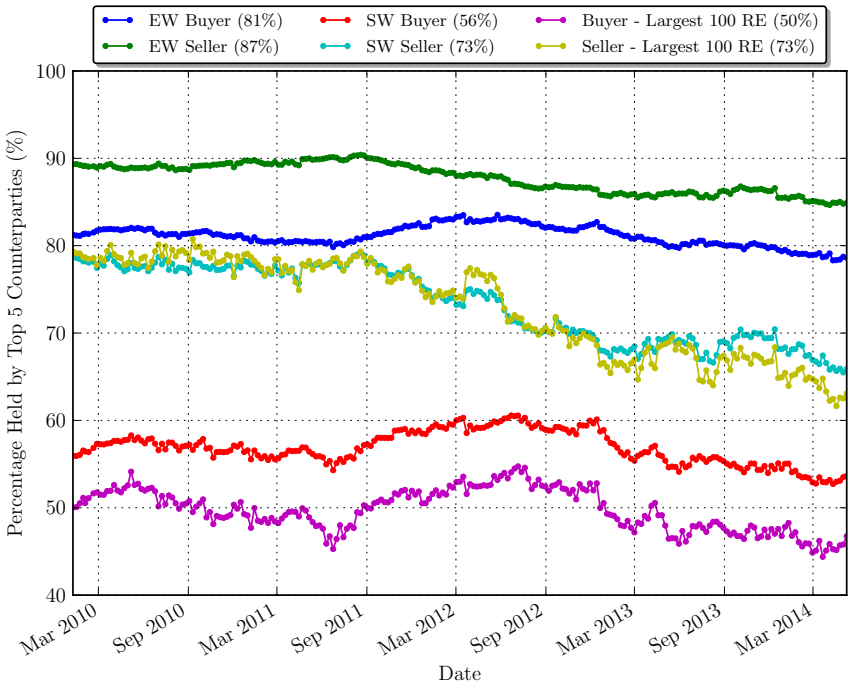

Notes: This figure plots the aggregate concentration of buyers and sellers in CDS markets. Seller (buyer) concentration is defined as the proportion of net selling (buying) by the top five largest sellers (buyers). This statistic is computed for each reference entity and date, and then aggregated by taking either an equal-weighted average or a size-weighted average. Finally, I compute the (equal-weighted) average for the largest 100 reference entities by net notional. In the legend, time-series averages are shown in parenthesis.

The concentration of sellers has nonetheless decreased over time. At the beginning of the sample, the top five sellers account for nearly 80 percent of selling within a reference entity. By the end of the sample, their share drops to 65 percent, still quite high. The reduction in seller concentration is robust to equally weighing reference entities, or considering only the largest 100. On the other hand, the concentration of buyers has remained relatively constant throughout the sample period. As of March 2014, the largest five buyers accounted for an average of 54 percent of all buying within a reference entity

The picture I have painted thus far is still incomplete because I have investigated only how concentrated buyers or sellers are within a particular reference entity. It is equally important to consider whether buyers or sellers are the same *across* reference entities, which I refer to as "commonality." .

B.2 Credit Risk Is Highly Common Across Counterparties

I now propose simple measures of commonality of buyer ownership and seller ownership. For instance, two reference entities have high buyer commonality if there is a large overlap of net

buyers of their CDS. First, define the common buyer measure at time t, denoted by CB_{ijt}, as:

$$CB(i,j,t) := -\frac{\sum_{c=1}^{O_B}(NS(c,i,t) + NS(c,j,t))}{NO(i,t) + NO(j,t)} \tag{18}$$

where O_B is the number of parties who are net buyers of both i and j at time t. The negative sign in the definition comes from, again, the fact that I define net buyers as negative net sellers. The common seller measure, $CB(i,j,t)$, is similarly defined but uses the set of net sellers who are common to both reference entities.

Both buyer and seller overlap measures are bounded between zero and one. Naturally, if there is no overlap between two reference entities, the corresponding overlap measure will be zero. Similarly, if one counterparty is the only buyer (seller) of protection in both reference entities, then the buyer (seller) overlap measure will be one.

Empirical Results

Table 7 in Appendix A summarizes my commonality statistics across reference entities and time. Over the whole sample, the average pair of reference entities has 22 percent of their buyers in common and 27 percent of their sellers in common. Anton and Polk (2013) compute a similar set of computations for mutual fund ownership in equities. Their sample ranges from 1980 to 2010, and they find the *maximum* average overlap for stock ownership to be 1.6 percent. Clearly, credit risk in CDS markets has a much more common ownership base compared to equities.

Some reference entity pairs have a great deal of overlap in both their sellers and buyers of protection. The top 5 percent percentile of reference entities pairs have 67 percent of their buyers in common, and 77 percent of their sellers in common. The reference entities with the most common buyers and sellers are also the largest. To highlight this fact, Figure 11 plots the equal-weighted average overlap measure over all pairs for the top 100 reference entities.

For the largest reference entities, an average of 80 percent of the credit risk in CDS is borne by a common set of counterparties. Is this problematic? It seems reasonable to think a large number of counterparties might equallysell protection on many of the largest reference entities.[47] This would give rise to a large average pairwise seller overlap in the largest reference entities, but also a small seller concentration. As we saw in Figure 10, this is not the case. In order for the seller overlap measure in Figure 11 to simultaneously match the concentration analysis in Section B.1, it must be that only a few counterparties are responsible for credit risk transfer in the U.S. credit default swap market. The extremely small number of net sellers and net buyers overall is further confirmed in the Online Appendix Section 1.4, when I simultaneously study how and where credit risk flows through the CDS network.

[47] In the extreme, consider the simplest CAPM environment from equities. In this case, everyone holds the market portfolio. This would indicate a very large commonality across equities, but also a very small ownership Herfindahl within any given name.

Figure 11: Buyer and Seller Commonality for CDS Markets

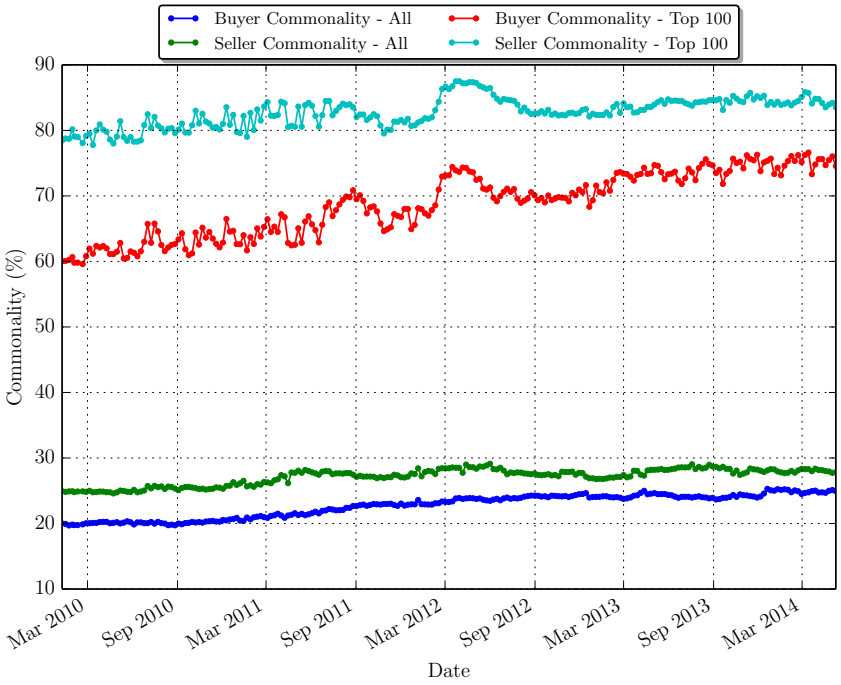

Notes: This figure plots the average pairwise buyer and seller commonality measures through time. The pairwise seller (buyer) commonality between reference entities i and j is the proportion net sellers (buyers) who are common to i and j. The series corresponding to the top 100 refers to the average buyer and seller commonality across all pairwise combinations of the 100 reference entities with largest net notional outstanding.

C Appendix: Mega-Seller Risk Bearing Capacity 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.

C.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.[48] To summarize, no-arbitrage says that Z-Spread $\approx$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

[48] 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.

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".

C.2 Panel Regression Results

To get a visual sense of how bond implied Z-spreads and the basis move through my sample, Figure 12 plots the average Z-spread across bonds at each point in time, the average basis, and my estimate of the price 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.[49] 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.

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

Figure 12: Average Z-Spread, Price of Credit Risk, and CDS-Bond Basis

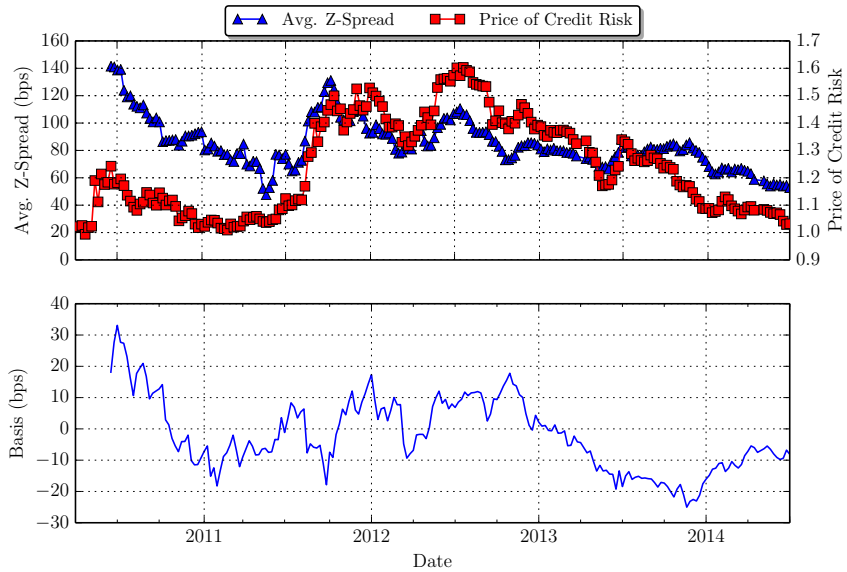

Notes: The top panel of this figures shows the average bond Z-spread and the price of credit risk implied by CDS markets. Z-spreads are obtained from Bloomberg, and are for the outstanding bond whose maturity is closes to five years. The price of credit risk derives from a panel of CDS spreads, and is estimated using the regression in Equation (24). The bottom panel shows the average CDS-bond basis, also obtained from Bloomberg.

To formally confirm this logic, I rerun the benchmark panel regression from Section 4.2, but replace log-CDS spreads with either (log) Z-spreads or the basis:

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

where the same controls are used from Section 4.2. 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. Table 8 collects the results from running both specifications, with and without controls.

The estimates in Table 8 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

Table 8: 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.

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.

We must be cautious in comparing the estimated magnitudes in columns (1) and (2) to the results from Section 4.2. 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 spreads 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% CDS spread increase — are likely more reliable.

D Granger Causality Tests from a VAR

A potential objection to the conclusions in the main text is a high price of credit risk drives losses to mega-sellers, not vice versa. To alleviate this "reverse causality" concern, I check the direction of Granger causality between the price of credit risk and mega-seller losses. Sections 6 and 7 present further evidence against reverse causality.

If it is really the losses at key players that drives the price of credit risk, then RB_t^s should Granger cause the price of credit risk $\widehat{\pi}_t$, but not the other way around. As previously mentioned, Appendix F.2 provide details of how I estimate $\widehat{\pi}_t$ using the benchmark regression of the preceding subsection. To implement the appropriate Granger causality tests, I employ an unrestricted first order vector autoregression (VAR):

$$\begin{aligned} \boldsymbol{y}_t &= \boldsymbol{c} + \boldsymbol{\Phi}\boldsymbol{y}_{t-1} + \boldsymbol{\varepsilon}_t \\ \boldsymbol{y}_t &:= \begin{bmatrix} X_t & RB_t^b & RB_t^s & \widehat{\pi}_t \end{bmatrix}' \end{aligned}$$

where X_t is a control variable designed to capture macroeconomic conditions at time t. To keep the size of the VAR manageable and ease the interpretation of the tests, I use same macroeconomic controls from before (e.g. the VIX), but insert them one at a time. It is important to include some type of control for general macroeconomic conditions. If I omitted X_t, I might find RB_t^s Granger causes the price of credit risk, when in reality X_t is really driving both variables.

Based on standard information criterion (AIC), I use of a first-order VAR. I then vary the control variable X_t and test for Granger causality in the system using the traditional F-test. If there is no direct causality between a pair of variables, I continue to test whether there is indirect causality via the other variables.[50] Table 9 reports the relevant results of this exercise.

Table 9: Granger Causality Tests for the Price of Credit Risk

	Control Variable, X_t				
Direction	**T10Y-2Y**	**T10Y**	**TED**	**VIX**	**CFNAI**
$\widehat{\pi}_t \to RB_t^s$	No	No	No	No	No
$RB_t^s \to \widehat{\pi}_t$	Yes	Yes	Yes	Yes*	Yes
$\widehat{\pi}_t \to RB_t^b$	No	No	No	No	No
$RB_t^b \to \widehat{\pi}_t$	No	No	No	No	No

Notes: This table shows the results of Granger causality tests for the price of risk, RB_t^s, RB_t^b, and a control variable X_t. $a \to b$ tests whether a Granger causes b. All tests were carried out via a first-order VAR and a subsequent F-test. When there is no direct Granger causality, I test for indirect causality via all possible variables. If all tests fail, then the null of no causality cannot be rejected, and I say one variable does not Granger cause another. If any pass, I say the opposite. All tests are conducted at a 5% confidence interval. T10Y-2Y is the 10 year minus 2 year constant maturity Treasury yield; T10Y is the level of the 10 year treasury yield; the remaining variables are self explanatory. * denotes indirect Granger causality.

Put simply, I find that mega-seller losses Granger cause the price of credit risk, but not vice versa. Regardless of the control variable, RB_t^s has significant forecasting ability for the price of risk,

[50] This turns out to be quite easy for all of my estimated VAR coefficients (results omitted), since RB_t^b never Granger causes the price of risk and very rarely do the variables considered Granger cause RB_t^s.

even after controlling for its own lag, losses to buyers, and general macroeconomic conditions. The Granger causality is direct for all but one of the control variables.[51]

Unsurprisingly, the effect is not present when considering Granger causality between the price of risk and losses to large buyers, RB_t^b. This finding is also in line with the panel regression results from Section 4.2.

Another piece of supportive evidence is found in the correlation matrix of the residuals from the VAR, which I present in Table 10. For illustration, I computed the correlations when using the slope of the Treasury term structure as the conditioning variable, but the results are qualitatively similar when using the other control variables.

Table 10: Correlation Matrix of VAR Residuals

	$\widehat{\pi}_t$	$\mathbf{RB_t^s}$	$\mathbf{RB_t^b}$	**T10Y2Y**
$\widehat{\pi}_t$	1	-0.64	0.30	0.07
$\mathbf{RB_t^s}$	-	1	-0.51	-0.02
$\mathbf{RB_t^b}$	-	-	1	0.05
T10Y2Y	-	-	-	1

Notes: This table displays the correlation matrix of the VAR of: (i) the price of credit risk estimate; (ii) lagged annualized CDS portfolio returns of mega-sellers and (iii) lagged annualized CDS portfolio returns of buyers; and (iv) the slope of the constant maturity Treasury yield curve (10 year minus 2 year).

As one would expect, shocks to the portfolios of mega-buyers and mega-sellers are quite negatively correlated (approximately -50 percent). Swaps are a zero sum product, so when sellers make money, buyers lose money. Interestingly, shocks to the slope of the term structure are uncorrelated with the other shocks in the model. As is clear from the first row of the table, shocks to the price of risk are very negative correlated — about negative 64 percent — with shocks to the portfolio of sellers. Taken alone, this is not proof of a causal relationship between seller portfolio returns and the price of credit risk, but in the context of the other findings of the paper, it at the very least suggestive.

[51] When $X_t = VIX_t$, I find RB_t^s Granger causes the VIX, and the VIX in turn Granger causes the price of credit risk.

E Appendix: Japanese Tsunami

E.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.[52]

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 13 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 13: Five Year CDS Spread of Japan

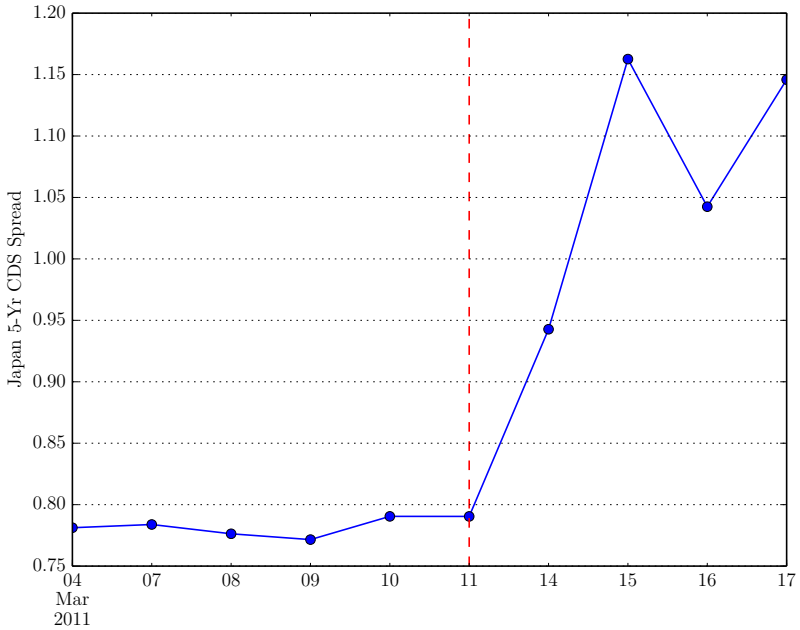

Notes: This figure plots the 5 year CDS spread for Japan from March 4, 2011 - March 17, 2011. The default event is of the clause "CR", for senior unsecured debt, and denominated in dollars.

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

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

Figure 14: Exposures of U.S. Counterparties to Japanese Firms: 3/11/2011

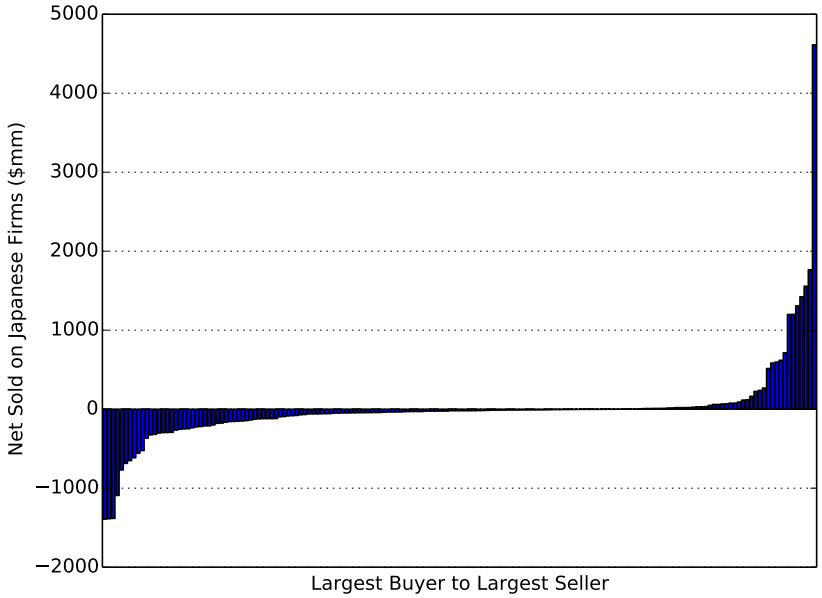

Notes: This figure shows the exposure of U.S. counterparties to Japanese firms. For each counterparty in the U.S., I compute the net amount sold on any firm whose country jurisdiction is in Japan, as classified by Markit. Positive values indicate a net seller overall, negative values indicate a net buyer.

Japan, which I confirm in Figure 14. 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.[53] In the figure, the largest buyers and their 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.[54] 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

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

[54] 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 15: Mark-to-Market Losses of U.S. Counterparties from Japanese Exposure: 3/11/2011 to 3/18/2011

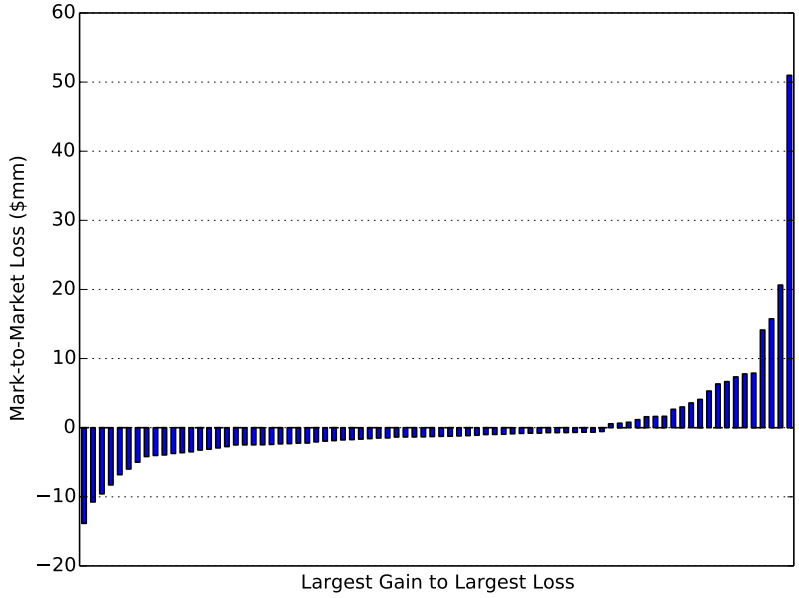

Notes: This figure mark-to-market losses of U.S. counterparties to Japanese firms in the week following the 2011 tsunami. Japanese reference entities include any firm whose country jurisdiction is in Japan, as classified by Markit. Positive values indicate a counterparty lost money due to Japanese exposure,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

15 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. 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.

E.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.

E.3 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:

$$\Gamma_{i,r} = a + \boldsymbol{\beta}' \boldsymbol{X}_r + \varepsilon_r, \quad i = S, B \tag{19}$$

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 11 collects these results.

Table 11: Regression of Γ on Reference Entity Observables

	Dependent Variable	
$\boldsymbol{X}_r$	$\Gamma_{S,r}$	$\Gamma_{B,r}$
$\rho(r, JPN)$	-0.28	0.22
	(-1.46)	(1.17)
σ_{CDS}	7.57	1.64
	(0.93)	(0.21)
$\log(NAICS)$	0.76	-0.12
	(1.33)	(-0.23)
$CDS_r(3/11/2011)$	0.17	-0.14
	(1.46)	(-1.24)
N	288	288
Adj R^2	1.9%	-0.1%
p-value of F-test	0.051	0.45

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, the (log) NAICS industry code, and the level of the CDS spread for r on the day of the tsunami.

Table 11 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 (barely) 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.

F Appendix: Additional Computations Used in the Paper

F.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.[55] 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:

$$CDS_{rt}(\tau) \quad = \quad \frac{\lambda_{rt}^{\mathbb{Q}}(\tau)}{\lambda_{rt}^{\mathbb{P}}(\tau)} \times LGD_{rt}^{\mathbb{Q}} \times \lambda_{rt}^{\mathbb{P}}(\tau) \tag{20}$$

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 trade-off for bearing r's default risk.[56] 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 (20) gives:

$$\log\left(CDS_{rt}\right) = \log\left(LGD_{rt}^{\mathbb{Q}}\right) + \log\left(\lambda_{rt}^{\mathbb{P}}\right) + \log(Q_r) + \log\left(\pi_{rt}\right) \tag{21}$$

where I've omitted the functional dependence of variables on the time to maturity. To make Equation (21) 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.[57]

[55] 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).

[56] 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.

[57] 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.

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

$$\log\left(LGD_{rt}^{\mathbb{Q}}\right) = \log(\eta_r) + \log\left(\widetilde{LGD}_{rt}^{\mathbb{Q}}\right) \tag{22}$$

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 plausbile.[58]

Substituting Equation (21) into Equation (22) yields:

$$\log\left(CDS_{rt}\right) = \log(\eta_r) + \log(Q_r) + \log\left(\widetilde{LGD}_{rt}^{\mathbb{Q}}\right) + \log\left(\lambda_{rt}^{\mathbb{P}}\right) + \log\left(\pi_t\right) \tag{23}$$

Equation (23) 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 (23) suggests the additional control variables that enter the regression capture the price of credit risk, π_t. This logic underlies my empirical approach, and is why I interpret the regression in Equation (4) as testing how mega-player risk bearing capacity affects the price of credit risk.

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.[59] In Appendix C, 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 (23) suggests one can estimate $\widehat{\pi}_t$ using a reference entity controls plus a time fixed effect, which I implement in the next section.

F.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:

$$\log\left(CDS_{rt}\right) = c + a_r + \beta_1 \log\left(EDF_{rt}\right) + \beta_2 \log\left(\widetilde{LGD}_{rt}^{\mathbb{Q}}\right) + \alpha_t + \varepsilon_{rt} \tag{24}$$

[58] 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.

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

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

$$\widehat{\pi}_t = \exp\left(\alpha_t\right) \tag{25}$$

where the standard error of $\widehat{\pi}_{rt}$ is computed using the Delta method.[60] 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 the first day in the sample from the fixed effect of the first firm.[61] 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. This point will be become more apparent in Section 4.3.

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

$$\begin{array}{ccccccccccc} \log(CDS_{rt}) & = & 0.60 & + & 0.25 & \log(EDF_{rt}) & + & 0.36 & \log\left(\widetilde{LGD}^{\mathbb{Q}}_{rt}\right) & + & \sum_r a_r 1_r \quad + \quad \sum_t \alpha_t 1_t \\ & & (32.4) & & (9.9) & & & (4.8) & & & \end{array} \tag{26}$$

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

Figure 16 plots my estimate for the price of risk, $\widehat{\pi}_t$. It is clear that there is significant time-variation in the aggregate price of credit risk. The two major peaks in this series occur during the summer of 2012, and the late fall/winter of 2011. At these two points in time, the price of risk was nearly 160 percent higher relative to January 2010. 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 so-called "Fiscal Cliff"). All of these events contributed to the rise in the price of risk during this time period. As is clear from the trend at the end of the time series, the price of risk basically returned to early 2010 levels by mid-2014.

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

[60] i.e. the standard error for $\widehat{\pi}_t$ is $\exp(\alpha_t) \times se(\alpha_t)$.

[61] This is one major advantage of the approach in BDDFS, who estimate a process for $\lambda_t^{\mathbb{Q}}$ and $\lambda_t^{\mathbb{P}}$ 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.

Figure 16: Aggregate Price of Credit Risk

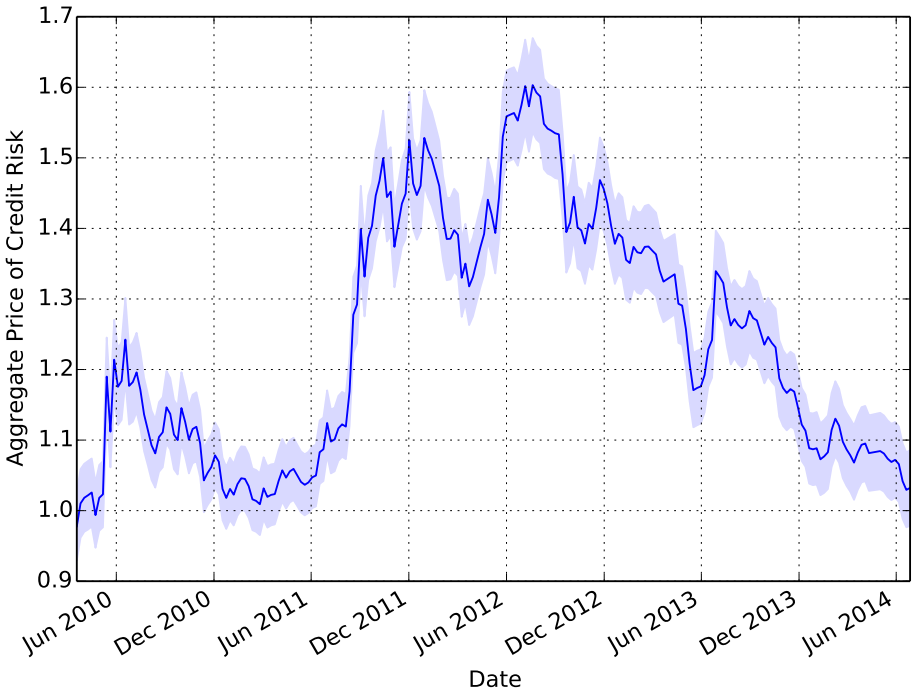

Notes: This figure plots the price of risk, $\widehat{\pi}_t$, estimated from the following regression: $\log(CDS_{rt}) = c + a_r + \beta_1 \log(EDF_{rt}) + \beta_2 \log\left(\widetilde{LGD}^{\mathbb{Q}}_{rt}\right) + \alpha_t + \varepsilon_{rt}$. EDF_{rt} and LGD_{rt} are the Moody's 5 year expected default frequency and industry loss-given-default, respectively, for reference entity r. a_r is a firm fixed-effect andα_t is a time fixed-effect. The price of credit risk $\widehat{\pi}_t = \exp(\alpha_t)$. 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.

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.

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

F.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 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.[62] 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.[63] 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

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

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

the weighted average of the returns of each position, weighted by the gross size of the position:

$$
\begin{aligned}
R_{c,t} &= \sum_{p \in \mathcal{P}_t(c)} \left[\frac{N^p_{c,t}}{\sum\limits_{p \in \mathcal{P}_t(c)} N^p_{c,t}} \right] R^p_{c,t} \\
&= \frac{\sum\limits_{p \in \mathcal{P}_t(c)} \left(V^p_{c,t} + AP^p_{c,t} \right) / \tau^p_t}{\sum\limits_{p \in \mathcal{P}_t(c)} N^p_{c,t}}
\end{aligned} \tag{27}
$$

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

F.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:

$$
\begin{aligned}
RB^s_t &:= \frac{\sum_{c \in TAS_t} \left[\sum_{p \in \mathcal{P}_{t-1}(c)} \left(V^p_{c,t-1} + AP^p_{c,t-1} \right) / \tau^p_{t-1} \right]}{\sum_{c \in TAS_t} \left[\sum_{p \in \mathcal{P}_{t-1}(c)} N^p_{c,t-1} \right]} \\
RB^b_t &:= \frac{\sum_{c \in TAB_t} \left[\sum_{p \in \mathcal{P}_{t-1}(c)} \left(V^p_{c,t-1} + AP^p_{c,t-1} \right) / \tau^p_{t-1} \right]}{\sum_{c \in TAB_t} \left[\sum_{p \in \mathcal{P}_{t-1}(c)} N^p_{c,t-1} \right]}
\end{aligned} \tag{28}
$$

F.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 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:

$$
U^O(t, T) = \frac{P_t(K_2, T) - P(K_1, T)}{K_2 - K_1} \tag{29}
$$

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:

$$U^O(t,T) = \xi k \times \frac{1 - \exp\left(-(r + \xi k)(T - t)\right)}{r(t,T) + \xi k} \tag{30}$$

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 (30) 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.