HARVARD | BUSINESS | SCHOOL

Concentrated Capital Losses and the Pricing of Corporate Credit Risk

Emil N. Siriwardane

Working Paper 16-007

Concentrated Capital Losses and the Pricing of Corporate Credit Risk

Emil N. Siriwardane
Harvard Business School

Working Paper 16-007

Concentrated Capital Losses and the Pricing of Corporate Credit Risk*

Emil N. Siriwardane†

Abstract

Using proprietary data on all credit default swap (CDS) transactions in the U.S. from 2010 to 2014, I show that a firm's CDS spreads are driven by capital fluctuations of that firm's net protection sellers. Capital fluctuations of sellers account for 10 percent of the time-series variation in spread changes, a significant amount given that observable firm and macroeconomic factors account for less than 16 percent of variation during this span. Sellers of protection are also highly concentrated, with five sellers responsible for nearly half of net selling. This concentration leads to market fragility — losses at the largest sellers have an outsized impact on CDS pricing. These findings suggest a high degree of short-run market segmentation, and support theories where capital market frictions play a first-order role in determining market prices.

---

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

I am incredibly grateful to my dissertation committee: Robert Engle (chair), Viral Acharya, Xavier Gabaix, and Stijn Van Nieuwerburgh. In addition, I would like to thank: Nina Boyarchenko (discussant), Itamar Drechsler, Darrell Duffie, Jan Ericsson (discussant), Sam Hanson, Alan Moreira (discussant), Holger Mueller, Thomas Philippon, Sriram Rajan, Alexi Savov, Philipp Schnabl, Andrei Shleifer, Erik Stafford, Adi Sunderam, and various seminar participants. 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: Harvard Business School and the Office of Financial Research, U.S. Department of Treasury. E-mail: esiriwardane@hbs.edu.

1 Introduction

A core assumption in neoclassical asset pricing theories is that capital can always flow frictionlessly to investment opportunities. For many asset classes, however, there are barriers to capital entry because investment requires specialized knowledge and technology, or capital itself may be scarce due to agency problems. For instance, financial institutions that participate in derivatives markets must have access to a steady source of funding, employ traders that have the requisite knowledge to properly evaluate risk, and possess the trading technology to execute and process trades.

Even within these financial institutions, internal capital market frictions can play a considerable role in impeding the flow of investment resources in the short run. Many financial institutions are active in a number of different asset classes, but trading desks within these institutions often focus on a specific market or a specific firm within a market. In turn, specialized trading desks are allocated a pool of capital to finance trading activity, but this capital is not easily replenished on short notice due to, for instance, agency problems within the firm.

In the presence of these capital market frictions, asset prices may behave quite differently than what neoclassical theory would predict, at least at high frequencies. Instead of price dynamics depending solely on exposure to fundamental risk factors, price movements can also reflect changes in the capital position of a small subset of trading desks in the market.

The spirit of this idea is at the heart of theories of limits to arbitrage (Shleifer and Vishny (1997); Kyle and Xiong (2001)), slow moving capital (Duffie (2010)), and financial intermediary-based asset pricing (He and Krishnamurthy (2013)). A major challenge for these models though is on the empirical side. Identifying the causal impact that limited capital has on asset price behavior requires detailed knowledge of market participants and their portfolio positions. Moreover, without granular data, it is difficult to know the degree to which markets are segmented. In the absence of such data, much of the existing empirical work in the field has linked high level measures of capital, such as financial intermediary leverage, to asset price dynamics (Gabaix et al. (2007) and Adrian et al. (2014)).[1]

[1] Other examples include Froot and O'Connell (2008), Mitchell, Pedersen, and Pulvino (2007), Coval and Stafford (2007), He, Kelly, and Manela (2015), and Chen, Joslin, and Ni (2014b).

In this paper, I take steps to overcome these hurdles using a proprietary dataset of credit default swap (CDS) positions that covers the entire U.S. market from 2010 to 2014.[2] I begin by showing that net buyers and net sellers of CDS protection are both highly concentrated, with sellers twice as concentrated as buyers. The top five sellers account for nearly half of all net selling — 50 percent of net selling is in the hands of less than 0.1 percent of the total number of CDS traders. Because limited capital is most likely to impact pricing in concentrated asset classes that require specialization, CDS markets (and derivative markets in general) are therefore an especially attractive venue to test how capital market frictions impact pricing.

To properly identify the causal link between capital and price dynamics, I examine how the CDS spread of a firm responds when that firm's default insurance providers suffer capital losses on their positions taken on unrelated firms.[3] I find that changes in seller capital account for a substantial amount of CDS spread movements. Capital fluctuations for protection sellers account for nearly 10 percent of the variation in weekly CDS spread movements. To put this in perspective, observable firm-level and macroeconomic factors explain only 16 percent of spread variation over the same time period.

Intuitively, capital losses raise the effective risk aversion of sellers, thereby increasing the premium they require for bearing default risk. When a firm's sellers experience a one billion dollar capital loss (roughly 1.4 standard deviations), the level of CDS spreads rises by 3.1 percent per week. This elasticity is economically large, as the standard deviation of weekly spread movements is 6 percent for the average firm in my sample. Moreover, these results suggest that CDS markets are partially segmented at the firm level. Put differently, in the short run, one can view the CDS market for a particular firm as a standalone market, thereby implying a fairly high degree of capital market frictions. Consistent with a story

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

[3] As an example of my strategy, I examine how seller losses on positions taken outside of the auto industry impact changes in the CDS spread for Ford Motor Company. Importantly, I also control for a large number of reference entity characteristics and macroeconomic variables that may drive movements in CDS spreads. In my most stringent tests, I also use an industry-by-time fixed effect to compare, for instance, the spread response of Ford relative to GM when Ford's sellers experience larger losses (from outside the auto industry) than GM's sellers.

of slow moving capital, I also find that the impact of seller losses on CDS spreads dissipates fairly quickly and is reversed after about 8 weeks.

Next, I show that institutional-level measures of capital constraints — namely financial leverage — have negligible explanatory power for high-frequency CDS spread dynamics. Instead, it is the specialized capital of the CDS desks at the largest sellers that is important for explaining weekly changes in CDS spreads. This finding supports the view that internal capital market frictions within financial institutions play a key role in preventing capital from flowing into the market at shorter horizons.

In a segmented market like CDS, it is also natural to think that price dynamics might depend on the types of active financial institutions in the market at a given point in time. Indeed, a notable trend in the data is that the since the 2008 financial crisis, hedge funds and asset managers have steadily replaced dealers (e.g. financial intermediaries) as the largest net sellers of CDS protection. Motivated by this trend, I also test whether capital losses impact prices differently depending on whether dealers or hedge funds have a large market share. My results suggest that capital market frictions are larger for hedge funds than dealers, as evidenced by the fact that losses at hedge funds have a stronger impact on prices than do losses at dealers. This is an important result given that friction-based asset pricing models that utilize a representative investor would not predict the heterogeneity of market participants to play a first-order role in pricing.

Another aspect of the market structure that I examine is the extent to which the concentration of net sellers generates fragility, a particularly relevant question given the high level of concentration in CDS markets. High concentration creates fragility because an idiosyncratic capital shock to a large seller can have a sizable impact on the total amount of risk bearing capital in the market. In line with this thinking, I show that capital losses at sellers with a large market share have more of a pricing impact than losses at sellers with a small market share. This finding also implies that the distribution of risk bearing capital within a segmented market is an important consideration for pricing.

One limitation of my analysis is that I do not observe market players' holding of underlying bonds. This is primarily an issue for net buyers of credit protection, who typically use CDS to hedge an underlying corporate bond portfolio. Thus, losses at the CDS desk for net buyers may be offset by gains in their corporate bond portfolio. This is less likely to be the case for the CDS desks at large net sellers because directly hedging a sold CDS position requires costly shorting of the underlying bond. Furthermore, net sellers are much more concentrated than buyers, so their risk bearing capacity should have a greater influence on CDS pricing. As expected, I find little evidence of a relationship between credit spreads and buyer capital movements.

Finally, to reinforce the causal link between capital and pricing, I use the 2011 Japanese tsunami to study how an exogenous shock to seller risk bearing capacity affected CDS spreads on U.S. firms. To trace out the impact of the tsunami, I exploit the fact that U.S. counterparties had large and heterogenous CDS exposures to Japanese firms prior to the tsunami. I then compare U.S. firms whose sellers had large Japanese CDS exposures to U.S. firms whose sellers had low Japanese exposure. Firms whose primary protection sellers were highly exposed to Japan saw their CDS spreads rise 2.5 percent in the week after the tsunami, relative to reference entities whose main sellers had very low exposure to Japan. In addition, I find no evidence of buyers transmitting the shock of the tsunami to U.S. reference entities. To emphasize the importance of concentrated positions, I make use of the fact that one counterparty had a particularly outsized exposure to Japanese firms. I then compare U.S. firms based only on this counterparty's share of their selling (or buying). I show that U.S. firms where this counterparty had a larger share of selling also experienced larger spread increases after the tsunami struck.

The remainder of the paper proceeds as follows. Section 2 gives a brief description of the data and methods used in this paper, with details found in a separate Data Appendix. Section 3 presents the main stylized facts that form the basis of the rest of the paper. Section 4 establishes how seller capital losses impact CDS pricing, the degree of segmentation in the market, and the role of internal capital market frictions in preventing capital from flowing into the CDS market. Section 5 demonstrates how institution type and concentration augment the link between capital and pricing. Section 6 provides more empirical support of my main results by using the 2011 Japanese tsunami as an exogenous shock to the risk bearing capacity of CDS market participants. Finally, Section 7 concludes.

2 Data Description

The primary data I work with comes from the DTCC, who provides trade processing services for every major dealer in the credit default swap market. I have access to two complementary subsets of the DTCC's database: transactions and positions. Transactions represent flows in CDS, and positions represent stocks. In practice, computing positions from transactions is quite complicated and is done using the DTCC's own internal algorithms.

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

In addition, I merge my transaction and position data with Markit and Moody's expected default frequency (EDF) database. Markit reports a daily CDS spread term structure for a large number of reference entities. Markit CDS spreads represent a composite spread that is computed using quotes and transaction information from 30 major market participants. Moody's EDF is a standard database of expected firm default probabilities that is derived from structural models of credit risk (Merton (1974)).

The DTCC dataset contains over 40 million index positions and 600 million single name positions. To be precise as possible, I document 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 analyses contained in the main text.

The Treatment of Index Swaps Index swaps constitute nearly half the gross notional of the entire CDS market, so accounting for exposures via index swaps is crucial for understanding true credit risk exposures. CDS index products contain a basket of single name swaps. For example, suppose a trader sells \$100 of notional on an index swap that contains 100 different single names. Like a single name swap, if one of the names defaults, the trader must pay out up to \$1 in notional to the buyer of the index swap, depending on the recovery rate of the underlying bond. After this payment, there are 99 names in the index remaining. Writing \$100 in protection via an index is therefore equivalent to writing 100 different single name swaps, each worth \$1 in notional. When considering the amount of credit risk exposure to a single reference entity, I am careful to account for exposures via single name swaps and index swaps. Full details of this procedure are contained in the Data Appendix, Section 1.3.

3 How Is Credit Risk Shared in the CDS Market?

In this section, I document four facts that form the basis of main empirical analyses: (i) the U.S. CDS market is large in terms of net notional credit risk transferred, with a conservative lower bound of around \$1 trillion; (ii) net sellers and buyers of CDS protection are very concentrated, with sellers more concentrated than buyers; (iii) the identities of the largest buyers and sellers is persistent through time; and (iv) dealers were the primary net sellers of protection until 2012, after which hedge funds have become the largest net sellers.

To establish these facts I will repeatedly use the following notation. r defines the underlying firm on which a credit default swap is written. c represents a counterparty (e.g. each transaction has two counterparties, the buyer and seller).[4] $NS(c, r, t)$ denotes the net amount of protection sold by c on reference entity r as of date t. Positive values of NS indicate a net selling position and negative values indicate a net buying position. For instance, if trader c sells \$100 of protection on reference entity r in one trade, and buys \$25 protection on r in a different trade, then $NS(c, r, t) = \$75$. Lastly, C_t is the set of all counterparties with open positions on date t, and R_t is the set of reference entities traded in the CDS market as of date t.

[4] I define c at the financial institution level, which inherently assumes that CDS trading within an institution is to some extent coordinated. When a majority of CDS trading at an institution is done at through one desk, as it often is, then this is a reasonable assumption.

3.1 The Size of the CDS Market

To quantify the size of the CDS market, I consider two alternative measures. The first measure is the gross notional size of the market, which is just the sum of the notional amount of all outstanding positions. Gross notional is thus a measure of volume, but importantly, it does not speak to the net amount of credit risk transfer for a given reference entity or for the entire market.

The second measure of market size that I use is the net notional outstanding of all positions. I define the net notional amount of credit risk outstanding for a *given reference entity* as:

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

$NO(r,t)$ is analogous to the face value of debt outstanding in bond markets — it captures the net amount of protection sold (or equivalently bought) on a particular reference entity. I then measure the total net outstanding for the whole market by summing $NO(r,t)$ over all reference entities:

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

Panels A and B in Figure 1 plot the gross notional outstanding and the net notional outstanding of the entire U.S. CDS market through time. Both measures provide a conservative lower bound on the size of the U.S. market because I only include reference entities that I can definitively classify as being based in the United States. The Data Appendix contains details of this classification procedure, and the Online Appendix provides less conservative bounds for the net outstanding in the market.

According to both measures, the size of the CDS market has steadily declined since the beginning of 2010. In January 2010, the gross notional size of the U.S. market was roughly \$17 trillion, but by May 2014 had fallen nearly 24 percent to \$13 trillion.[5] Similarly, the net notional outstanding of the CDS market declined

[5]These statistics roughly accord with aggregate data provided by the Bank for International Settlements (BIS): http://www.bis.org/publ/qtrpdf/r_qt1512_charts.pdf. The gross notional that I report is slightly less than the BIS's estimate, which is not surprising since I focus only on the U.S. CDS market. The net notional outstanding is new to this paper, given that one must

about 38 percent over this same time period. The average net notional outstanding in the market is about \$1 trillion for the entire sample.

Despite the downward trend in the size of the CDS market, the overall amount of notional credit risk transferred is still large. As a rough comparison of magnitude, the face value of debt outstanding in U.S. corporate bond markets is approximately \$9 trillion. Conservatively speaking, the size of the CDS market is thus about one-ninth of the size of the corporate bond market.

3.2 Concentration of Net Buyers and Sellers of Protection

3.2.1 Notional Concentration

In this subsection, I develop a simple measure of concentration for the CDS market. In the Online Appendix, I explore a number of complimentary measures that all deliver a similar set of implications. To quantify the concentration of net buyers and sellers of a *given* reference entity r, I use the following measure of market share for each counterparty and reference entity:

$$MS(c,r,t) := \frac{NS(c,r,t)}{NO(r,t)} \tag{3}$$

$MS(c,r,t)$ measures counterparty c's share of net selling in reference entity r. Positive market shares indicate that c is a net seller of r, and negative market shares indicate that c is a net buyer. For instance, if $MS(c,r,t) = 0.2$, then c is responsible for 20% of the total net notional protection sold on r.

Next, to compute an *aggregate* market share measure for each counterparty, I take a size-weighted average across all reference entities:

$$\begin{aligned} MS(c,t) : &= \sum_{r \in R_t} \omega_{rt} \times MS(c,r,t) \\ \omega_{rt} &:= NO(r,t) / \sum_{r \in R_t} NO(c,r,t) \end{aligned} \tag{4}$$

where I use a size-weighted average instead of an equal-weighted average to offset the influence of extremely small reference entities that typically have only one net buyer and one net seller.

net single name exposures against index exposures in order to compute this statistic.

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

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

Net sellers of CDS are highly concentrated. According to my definition of market share, the top five sellers account for 50 percent of all protection sold. Put differently, because there are about 1700 counterparties in the market, 50 percent of all selling is in the hands of less than 0.1 percent of potential counterparties. Buyers are also concentrated, albeit only half 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. Panel B of Figure 2 plots, for both buyers and sellers, the count of top five counterparties that remains the same from time $t-1$ to t. For example, in week t, I count the number of top five sellers who were also in the top five in the previous week. On average, 94 percent of the top five buyers and 96 percent of the top five sellers remain constant from week to week. Thus, not only are CDS markets highly concentrated with a handful of buyers and sellers, but this organizational feature of the market is also fairly static through time.

3.2.2 Market Beta Concentration

An alternative way to quantify concentration is in terms of risk exposure. In turn, a natural risk factor to consider is exposure to market risk, as measured by a traditional CAPM beta. More specifically, for each counterparty c, I compute the dollar value of their CAPM beta via the following time-series regression:

$$\Delta V_{ct} = a + \beta_c Ret_t^{Mkt}$$

where ΔV_{ct} is the dollar change in the market value of c's portfolio from week $t-1$ to t, and Ret_t^{Mkt} is the excess return of the CRSP value-weighted stock market index over the same horizon. I provide additional details of how I compute ΔV_{ct} for each counterparty in Section 4. There are of course many other potential risk factors that a counterparty might be exposed to, but I focus exposure to the aggregate market as a first pass.

If $\beta_c > 0$, it means that counterparty c is in some sense selling insurance on the market — if the market crashes, then these counterparties will incur losses. I denote this group of counterparties as "sellers". Conversely, if $\beta_c < 0$, it means that counterparty c is buying protection against a market downturn, and I denote this group of counterparties as "buyers". To quantify the concentration of sellers, I then compute a weight $w_{ct}^s = \beta_c / \sum_{i \in sellers} \beta_i$ for each counterparty. The weights for buyers, w_{ct}^b, are computed in an analogous fashion.

Based on this simple measure, the concentration of sellers is striking. The top five sellers by w_{ct}^s account for nearly 75 percent of all market risk exposure. The concentration of buyers is also high, though not nearly as large as the concentration of sellers. According to w_{ct}^b, the top five buyers account for roughly 23 percent of protection bought against the aggregate stock market. Both of these findings echo the level of concentration found when looking at net notional exposures, though sellers are certainly more concentrated by this market-risk-based metric.

3.3 Who Are the Big Players in the Market?

Given the size and concentration of the CDS market, it is natural to ask: who are the major buyers and sellers 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, dealers, etc.

Next, for each reference entity and date, I compute the proportion of net buying and selling done by each type according to the measure in Section 3.2.1. For instance, I compute what proportion of GE's net outstanding is sold by insurance companies. The computation is analogous to calculating the market share of an individual counterparty in a reference entity, except I do so for a counterparty type. Finally, I create an *aggregate* index of the proportion bought and sold by each type y, which I denote by $\overline{P}_B(y,t)$ and $\overline{P}_S(y,t)$, respectively. Each aggregate index is simply type y's size-weighted average market share across all reference entities. More details of these computations are also contained in the Online Appendix.

Panel A and Panel B of Figure 3 plot $\overline{P}_B(y,t)$ and $\overline{P}_S(y,t)$, respectively, for dealers and hedge funds/asset managers (HFAMs) through time. I focus on these two counterparty types because they are by far the largest two types for both buyers and sellers. As seen in Panel A of Figure 3, dealers have consistently purchased approximately 55 percent of protection, with the most of the remaining buying going to HFAMs.

The aggregate proportion of selling by counterparty types appears in Panel B of Figure 3. In contrast to buyers, the composition of sellers has dramatically changed since 2010. At the beginning of the sample, dealers accounted for 80 percent of all protection sold in U.S. CDS markets, with this share heavily skewed towards less than five dealers (approximately 50 percent of aggregate selling). Nonetheless, the total proportion sold by dealers has declined by almost half, with dealers accounting for almost 40 percent of total selling by the end of the sample. 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. Instead, HFAMs have grown into a much larger role in providing default insurance for the U.S. market. More specifically, less than five HFAMs account for nearly 30 percent of all net selling of protection as of the first quarter of 2014.

Why Are Markets So Concentrated? Intuitively, concentration develops naturally in any market with high fixed entry costs. CDS markets are costly to enter for a few reasons. First, trading CDS requires back-office processing of trades and risk management to manage existing positions. To this point, many smaller hedge funds will pay their dealer an additional fee in return for the dealer handling the oversight of trades. Moreover, establishing a CDS desk requires substantial information acquisition (Merton (1987)), not only in terms of hiring traders and managers with expertise in credit risk, but also specifically in credit derivatives. Second, CDS trading is similar to banking in the sense that relationships are "sticky" (Chodorow-Reich (2014)). For example, in the Online Appendix I show that the average non-dealer trades with only three counterparties. In lieu of the costs of building new trading relationships, it is no surprise that trading activity in all over-the-counter markets is dominated by a handful of dealers who can use their existing relationships from many lines of business. Third, operating a CDS desk is costly from a funding standpoint. Since the 2007-09 crisis, it has been common practice for CDS positions to be marked-to-market every day. Consequently, CDS desks need a stable source of funding in order to survive daily fluctuations in mark-to-market values. There are large economies to scale in terms of funding, and as a result, large dealers and hedge funds naturally emerge as key players in the market.

Of course, there are a multitude of additional reasons why CDS markets are concentrated. The purpose of this paper is not to answer this question, but rather to understand how limited capital in the market ultimately affects pricing in CDS. Nonetheless, my findings do shed some light on the question of concentration. For instance, because I find that HFAMs have become a dominant seller of CDS protection, it is unlikely that relationships play a first order role in concentration; if relationships were primarily driving concentration, dealers would always be the largest net buyers and sellers. On the other hand, the fact that limited capital does seem to impact prices suggests that funding frictions may be important for explaining the concentration of the CDS market. While concentration is certainly an interesting and important topic, further inquiry is outside of the scope of this paper.

4 The Impact of Seller Capital on CDS Spreads

As just discussed, CDS markets — and derivative markets in general — require stable funding, specific trading technology, and expertise. For these reasons, it is also natural to expect the capital of CDS market participants to play a first-order role in CDS pricing. The argument is a straightforward application of asset pricing theories with capital market frictions (e.g. Shleifer and Vishny (1997) and Duffie and Strulovici (2012)). These theories predict that when there are impediments to investment capital flowing into an asset market, shocks to capital of a small subset of agents play an important role in asset price dynamics. In this section, I provide direct empirical evidence of this phenomenon in the context of CDS markets.

Measuring Capital

I define the risk bearing capital of a given counterparty as the capital available to CDS traders for the purposes of initiating and maintaining new investments. At the trading desk, capital is required to initiate new trades because of initial margin payments and upfront payments that make the swap NPV zero. Maintaining an existing trade requires capital to make payments on net bought positions, variation margin payments, and in the case of net sellers, potential default payments. I measure changes in risk bearing capital using changes in the mark-to-market value of each counterparty's CDS positions, which captures CDS profits and losses (P&L).

I focus on the risk bearing capital of CDS traders, as opposed to the entire trading entity, for a few important reasons. From an institutional perspective, it is reasonable to view CDS trader-specific capital as the correct state variable for pricing. Trading desks at large dealers and hedge funds are subject to risk limits (e.g. value-at-risk), which may tighten with prolonged losses. Poor portfolio performance also means CDS traders have less capital to make variation margin payments on mark-to-market losses.[6]

This argument relies heavily on the existence of *internal* capital market frictions, at least in the short run. For example, if the CDS desk at a hedge fund suffers significant losses, it is not easy (or even optimal) to transfer capital from another desk to the CDS desk. Mitchell, Pedersen, and Pulvino (2007) provide evidence consistent with this story by showing that information barriers within a firm can lead to capital constraints for specific trading desks who have experienced mark-to-market losses. In Section 4.2, I show that a similar mechanism is

[6]Another potential channel for losses to affect pricing follows from Froot, Arabadjis, Cates, and Lawrence (2011), who demonstrate that loss aversion for institutional investors affects future trading. Froot and O'Connell (2008) also develop a model where costly external financing of intermediaries leads to above-fair pricing of catastrophe insurance.

at play in CDS markets.

It is also important to recognize that for net protection buyers, the capital of the CDS desk alone is not likely to capture the true dynamics of risk bearing capital. This is because net buyers often purchase CDS protection to hedge underlying corporate bond positions. However, the wealth of the CDS desk should capture the risk bearing capital of large protection sellers because it is unlikely their positions are hedged with other securities. Ultimately, this issue can be resolved empirically. If the P&L of the CDS trading desk tracks changes in risk bearing capital well, then P&L should also help explain price movements. Consistent with this reasoning, I find that seller capital, not buyer capital, impacts prices.

The final reason I use P&L as a measure of risk bearing capital is practical. Recent empirical research on limited intermediary risk bearing capacity has used leverage as a measure of risk bearing capacity (e.g. Adrian and Boyarchenko (2013)). The theoretical underpinnings of this work suggest leverage is a sensible metric because it is a proxy for the wealth available for bearing risk. In some sense, I have a more direct measure of this wealth because I have proprietary data on actual positional holdings, which means I can compute the dollar value of each counterparty's CDS portfolio. Moreover, I find that hedge funds are a large player in CDS, but leverage measures for these entities are either non-existent or poorly measured.

Computing the mark-to-market value of each counterparty's CDS portfolio is itself a computationally challenging task. It requires me to mark over 600 million CDS positions to market for each day in my sample period. To keep the problem manageable, I choose the simplest possible methodology, with the details found in Section 3 of the Data Appendix.

4.1 Capital Fluctuations and CDS Pricing

In any theory of slow moving capital, it is critical to define the level of segmentation. Segmentation may occur at an asset class level, like equities versus fixed income, or could be more granular, such as within an industry. Most models then imply that asset prices within a segment are driven by the total capital of market participants in that same segment.[7]

With this in mind, I examine how CDS spreads for a particular reference entity respond to changes in the capital of net buyers and sellers of that same reference entity (e.g. what happens to Ford's CDS spreads when Ford's net buyers and sellers lose capital). Thus, my empirical approach implicitly tests whether CDS markets are segmented at the reference entity level, at least in the short run.

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

The following regression, run at a weekly frequency, implements a similar concept in the context of CDS markets:

$$\Delta \log\left(CDS_{rt}\right) = a_r + \boldsymbol{\beta}_1^{'}\Delta \mathbf{Z}_{rt} + \boldsymbol{\beta}_2^{'}\Delta \mathbf{X}_t + \zeta_s OCF_{rt}^s + \zeta_b OCF_{rt}^b + \varepsilon_{rt} \qquad (5)$$

where CDS_{rt} is the 5-year CDS spread of reference entity r at time t. The most important variables in regression (5) are the OCF measures, which stand for "outside capital fluctuations".[8] For example, OCF_{rt}^s captures the total capital fluctuations of r's net sellers, with the caveat that these fluctuations are due to changes in the market value of *positions on reference entities outside of r's industry*. Formally, OCF_{rt}^s is computed as:

$$OCF_{rt}^s = \sum_{c\in S_{r,t-1}} \Delta V_{c,-r,t}$$

where $V_{c,-r,t}$ is the mark-to-market value of counterparty c's portfolio for all ref-

[7]If markets are not fully segmented, then this argument holds over a certain time horizon. Over longer time horizons, capital will slowly move between partially segmented markets. See Duffie (2010) for a more detailed discussion.

[8]Because I run the regression in log-differences, I interpret the OCF variables as impacting the risk premium embedded in the spread, as opposed to actual default risk. See Appendix A.1 for a more detailed discussion.

erence entities outside of the same industry as r and $S_{r,t-1}$ is the set of net sellers of protection on r at time $t-1$. OCF^b_{rt} is the same variable, but for r's net buyers of protection.

I obtain CDS spreads from the data vendor Markit, and defer further details of the underlying data to the Data Appendix. a_r is a reference entity fixed effect that absorbs any time invariant firm characteristics, though my results are nearly identical if I do not include a reference entity fixed effect. $\mathbf{Z}_{rt}$ is a vector containing the log of Moody's 5-year expected default frequency (EDF) and Markit's expected loss-given-default (LGD); I choose these firm-level controls based on reduced form models of credit risk. In some versions of regression (5), $\mathbf{Z}_{rt}$ also includes the log-CDS spread implied by options markets and the at-the-money implied volatility of options. To compute an option-implied CDS spread for reference entity r, I translate the price of out-of-the-money put options to CDS spreads using the methodology of Carr and Wu (2011). The details of this procedure are contained in Appendix A.2. The important advantage of using firm controls that derive from options markets is they control for a large number of unobservable firm-level and macroeconomic factors that may drive credit spreads.

$\mathbf{X}_t$ is a set of observable macroeconomic variables that may also cause CDS spread movements. I choose these controls based on theoretical models of credit risk and previous research on credit spread variation (e.g. e.g. Collin-Dufresne et al. (2001)). These variables are the log earning-to-price ratio for the S&P 500, VIX, TED, CFNAI, 10 year Treasury yield, 10-year-minus-2-year Treasury yield, and the CBOE Option Skew index. After first differencing these aforementioned controls, I also include the excess market return of the CRSP value-weighted index.[9] In my most stringent tests, I replace the vector $\mathbf{X}_t$ with an industry-by-time fixed effect to ensure that the point estimates in the regression are not biased by any unobservable industry factors at a given point in time. For each reference entity, I use the definition of industry as provided by Markit. Table 1 contains the results of regression (5).

Column (1) of Table 1 provides a benchmark in terms of how much CDS spread variation is explained by traditional firm and macroeconomic risk factors. The R^2 from column (1) indicates traditional credit spread determinants

[9] I include S&P 500 returns in order to account for higher frequency (weekly) equity movements. The earnings-to-price ratio is monthly and taken from Robert Shiller's website.

only capture 15.7 percent of spread variation on their own. The relatively low R^2 in the regression echoes previous research (e.g. Collin-Dufresne et al. (2001)) on the determinants of credit spread changes.

Column (2) adds the outside capital variables to the baseline regression with firm and macroeconomic controls. As is clear from the point estimates and their standard errors, outside capital fluctuations for sellers are an important determinant of spread changes.[10] A \$1 billion capital loss to net sellers on positions from outside of r's industry results in an increase of 3.1 percent in the level of r's CDS spread. To put this in perspective, the standard deviation of spread movements across all firms in my sample is about 6 percent and the standard deviation of OCF^s_{rt} is \$730 million. In many respects, this is a lower bound on the effect of seller capital losses on prices because I exclude losses coming from positions on firms in r's industry for the purpose of identification.

Column (2) also indicates that capital fluctuations help explain an additional 9.5 percent of spread variation, which is large given that observable macroeconomic and firm fundamentals explain only 15.7 percent on their own. Another way to view the incremental R^2 in column (3) versus column (1) is that capital fluctuations for sellers of CDS protection can explain about one-tenth of the variation in CDS spreads. Because CDS spreads are anchored to corporate credit risk, these results are also in line with previous studies that link financial-intermediary activity to corporate bond pricing (e.g. Green, Hollifield, and Schurhoff (2007), Newman and Rierson (2004)).[11]

In most asset pricing theories with capital market frictions, risk premiums interact non-linearly with capital. To test this in the data, the specification in Column (3) of Table 1 interacts outside capital fluctuations for sellers with three different dummy variables that indicate whether OCF^s_{rt} is in a given tercile. This specification allows for capital to differentially impact spreads based on loss size.

[10] I double cluster all standard errors by time and by top two sellers (alphabetized). Clustering by top sellers accounts for the fact that sellers may be common across reference entities, thereby generating correlation in the OCF^s_{rt} measures. My results are robust to the choice of clustering by two sellers (e.g. versus top five sellers). The conclusions are also unchanged when clustering by reference entity, which makes sense given spread changes are not that autocorrelated.

[11] In the Online Appendix, I provide some evidence that movements in CDS spreads translate to movements in actual bond yields, as opposed to just changes in the CDS-bond basis (though this is likely to depend on my post-crisis sample). This evidence is consistent with the idea that a large number of bond market investors are long-term holders of the bond, whereas CDS market participants are more active and therefore serve as the marginal pricer of credit risk.

As is clear from the point estimates, there is a highly non-linear relationship between risk premiums and capital. Unsurprisingly, large losses for protection sellers have the biggest impact on spread changes. It is also important to recognize markets have been relatively calm during the sample period of 2010 to 2014. One can imagine that during a period of extreme market stress, like those seen during the financial crisis, the impact of seller losses on CDS pricing would be further amplified.

Columns (4) and (5) represent my most fundamental evidence that seller capital losses are an key determinant of CDS spread movements. Column (4) removes macroeconomic controls from the regression and replaces them with an industry-by-time fixed effect to absorb any unobservable characteristics common to the cross section of firms within each industry at each point in time. Importantly, the point estimate on OCF^s_{rt} remains relatively unchanged in magnitude and statistically significant. Because one might still be concerned that I am omitting important firm-level characteristics, column (5) adds both option-implied CDS spreads and at-the-money volatilities to the regression. The sample size in this specification is cut in half because of an imperfect match between CDS data and options prices. Still, the key message is that, even after incorporating information on the firm implied by options markets, the effect of seller losses on spread movements is statistically significant and about 2.3 to 2.5 percent in magnitude. In contrast, the impact of buyer capital on CDS spread dynamics appears to be negligible.

The fixed-effects specification in columns (4) and (5) also speak to the degree of segmentation in the CDS market. To see why, suppose that CDS markets were segmented across industries. This might occur because market participants specialize in certain industries, making it difficult for capital to move across industries on short notice. In this case, the relevant state variable for pricing the CDS of a given firm would be the capital of all market participants in that firm's industry. In turn, including an industry-by-time fixed effect in the regression would account for any industry-wide movements in capital, and should therefore drive out the explanatory power of the OCF variables for pricing. Thus, the fact that capital changes for sellers of a specific reference entity explain spread movements in the presence of industry-by-time fixed effects suggests that CDS markets are at least partially segmented at the reference entity level.

Based on these results, what would one have to believe to invalidate identification in this setup? Consider the example of Ford. The exclusion restriction for the regression would be violated if OCF^s_{rt} captures some factor that drives Ford's CDS spread, but in a way that is: (i) not common to the auto-industry, as ruled out by the industry-by-time fixed effect; (ii) not better captured by Ford's own equity price (an input to the EDF) or; (iii) Ford's own option prices. In my view, this alternative seems implausible.[12]

The final thing to note from this exercise is that, relative to seller capital, buyer capital appears to play much less of a role in explaining weekly CDS spread movements. In almost all specifications, the coefficient on buyer capital is statistically indistinguishable from zero and generally very small in economic magnitude. The sign of the coefficient on buyer capital is positive, indicating that gains at the CDS desks of buyers lead to price increases. This finding is consistent with the notion that buyers of protection largely use CDS to hedge an underlying corporate bond position, but also do not fully hedge. Hence, gains at the CDS desks of large buyers are more than offset by losses on their bond portfolio, and in turn, lead to an increased demand for CDS protection. However, like in the market for catastrophe insurance (Froot and O'Connell (1999)), these demand shocks are less important than supply shocks for understanding CDS price movements over this time period and at a weekly frequency.

4.2 Internal Capital Markets

As previously discussed, my use of CDS trader capital inherently assumes that the trading organizations in the CDS market (e.g. hedge funds or dealers) face *internal* capital constraints, at least in the short run. In the absence of internal capital market frictions, losses at the CDS desk of a large dealer or hedge fund could be replenished by instantaneously transferring capital from other segments of the firm that are not capital constrained. In turn, the overall capital position of the trading firm would be the relevant state variable for pricing, as opposed to the capital of CDS traders in particular.

To determine the importance of CDS desk capital versus overall firm capital,

[12] Implicit in this argument is that changes in Ford's CDS spread do not drive OCF^s_{rt}, which is reasonable given: (i) that OCF^s_{rt} is constructed using positions outside of the auto-industry and (ii) that the total position in Ford is extremely small relative to the entire CDS portfolio.

I include the change in the average market leverage ratio (book debt to market equity) of both sellers and buyers in the baseline regression (5). Increased market leverage is meant to proxy for increased capital constraints for the entire trading organization.[13] Columns (6) and (7) of Table 1 demonstrate the importance of internal capital market frictions within trading organizations. In column (6), I exclude my measure of CDS-trader capital and replace it with changes in buyer and seller leverage. At a weekly frequency, changes in market leverage do have statistically significant explanatory power for changes in CDS spreads, but the magnitude of the effect is fairly small. The standard deviation of leverage changes for net sellers is around 3, indicating that a one standard deviation increase in seller leverage only raises CDS spreads by about 0.27 percent. In column (7), I include both my OCF measure and changes in leverage. In this case, the point estimates on seller leverage shrinks by a factor of two. More importantly, the inclusion of leverage does not alter the response of CDS spreads to changes in seller capital, as measured by OCF^s_{rt}. The last thing to note from column (7) is the interaction term between changes in seller leverage and the CDS-desk specific capital. Though imprecisely measured, the negative sign of the interaction term accords with intuition — losses at the CDS desks of sellers have a larger impact on spreads when they coincide with firm-wide capital constraints also tightening.

To summarize, there are two main takeaways from Table 1. First, CDS markets are partially segmented at the reference entity level, so in the short run, one can view the CDS market for a particular reference entity as its own standalone market. In turn, each reference entity's CDS spreads are driven by the capital of its net sellers. Second, the pertinent measure of capital for net sellers is the capital of the CDS desk, as opposed to just the overall capital of the firm. This finding indicates the presence of nontrivial capital market frictions inside of large financial institutions.

[13] Obviously I can only compute this measure for dealers. As shown in Section 3.3, dealers are the primary net sellers of protection in the market from 2010 to 2012. When I run the analysis using this window, the results are qualitatively very similar.

4.3 Horizon Dependence

I have argued that the primary economic mechanism driving these results derives from asset pricing theories with limited investment capital. A signature prediction of these theories is that, given time, capital is able to flow into the market and thus any pricing effects should disappear over longer horizons. To evaluate this prediction, I run the following regression for various horizons h:

$$\log\left(\frac{CDS_{r,t+h-1}}{CDS_{r,t-1}}\right) = a_r + \boldsymbol{\beta}_1' \Delta \mathbf{Z}_{rt}(h) + \boldsymbol{\beta}_2' \Delta \mathbf{X}_t(h) - \zeta_s(h) \times OCF_{rt}^s + \zeta_b(h) \times OCF_{rt}^b$$

where $\Delta \mathbf{Z}_{rt}(h) = \mathbf{Z}_{r,t+h-1} - \mathbf{Z}_{r,t-1}$ and $\Delta \mathbf{X}_t(h) = \mathbf{X}_{t+h-1} - \mathbf{X}_{t-1}$. Note that the sign of $\zeta_s(h)$ in the above regression means that it represents the effect of seller capital losses on CDS spreads from $t-1$ to $t+h-1$. The regressions from the previous subsection were run for $h = 1$. If the effect of seller losses on CDS pricing decays with time, then $\zeta_s(h)$ should tend towards zero as h increases. Furthermore, the regression is run without industry-by-time fixed effects because I want to identify the time-series impact of the OCF variables.

Akin to an impulse response function, Figure 4 plots the point estimate of $\zeta_s(h)$ along with 95 percent confidence bands for various horizons h. Consistent with idea that the price impact of seller losses reverses as capital flows into the market, $\zeta_s(h)$ declines as h increases, with a half life of about two weeks. As illustrated by the 95 percent confidence bands in Figure 4, the pricing effects of seller capital losses are fully undone after about 8 weeks.

It is not surprising that the pricing effects die out rather quickly because segmentation at the reference entity level is a fairly extreme form of capital market segmentation; hence, one would not expect it to persist for long periods. Moreover, it seems reasonable that internal capital market frictions at large financial institutions are resolved over short horizons as well. This interpretation is also consistent with He, Kelly, and Manela (2015), who show that at a quarterly frequency, the leverage of primary dealers can explain CDS returns. Importantly, they measure capital at the holding company level, thereby implying that at a quarterly frequency external capital market frictions are an important component to pricing. Their findings, combined with the pattern of decay that I've documented, suggest that short-run price dynamics are influenced by internal capital market frictions and segmentation at the reference entity level, whereas external capital market frictions are more relevant for long-run dynamics.

5 Market Structure and Pricing

Given the detailed nature of my data, I am also able to investigate two additional channels through which the structure of the CDS market may interact with pricing. The first feature of the CDS market that I explore is whether heterogeneity in the types of active financial institutions impacts pricing. This channel is important because, in the wake of the 2008 financial crisis, hedge funds and asset managers (HFAMs) have steadily replaced dealers as the primary providers of default insurance (Section 3.3). The second aspect that I examine is concentration, which is particularly relevant given that CDS markets are dominated by a handful of key players (Section 3.2).

5.1 Heterogenous Institutions and Pricing

To test whether the type of institution that is actively selling protection in the CDS market affects pricing, I estimate the following variant of regression (5):

$$\begin{aligned}\Delta \log\left(CDS_{rt}\right) \quad = \quad & a_r + a_{it} + \boldsymbol{\beta}_1^{'}\Delta\mathbf{Z}_{rt} + \zeta_s OCF_{rt}^s + \zeta_{s,HF} OCF_{rt}^s \times HFS_{r,t-1}^s \\ & + \zeta_{HF} HFS_{r,t-1}^s + \zeta_b OCF_{rt}^b\end{aligned}$$

where $HFS_{r,t-1}^s$ is the share of net selling in r by hedge funds and asset managers at time $t-1$. a_r and a_{it} are reference entity and industry-by-time fixed effects, respectively. $\mathbf{Z}_{rt}$ is the base set of reference entity specific controls used in Section 4. The interesting feature of this regression is the interaction term between OCF_{rt}^s and $HFS_{r,t-1}^s$, which measures whether capital fluctuations have a differential impact on pricing when hedge funds are responsible for more selling.

Column (1) of Table 2 indicates that seller capital losses have a larger impact on spreads if sellers are HFAMs. The interaction term between HFAM share and seller capital is significantly negative and fairly large in magnitude. For example, from 2010 to 2014, the share of selling by HFAMs moved from roughly 15 percent to 55 percent. In turn, the interaction term implies that the impact of a \$1bn seller capital loss on spreads has changed from 2.31 to 3.47 percent over this same time period. One way to rationalize this finding is that HFAMs have a higher shadow cost of capital than dealers, who were the primary provider of credit insurance at the beginning of the sample. This interpretation seems reasonable given that, relative to dealers, hedge funds are more specialized and generally have a smaller capital base; thus, they are more likely to face external capital market frictions.

5.2 Why Does Concentration Matter?

My findings thus far suggest that the total capital of *all* sellers in a reference entity is important for explaining price dynamics. However, consider the following thought experiment: hold the total level of risk bearing capital fixed, but vary the distribution of capital within natural sellers of protection. In this case, it is not obvious from theory whether the *level* of risk premiums should change, even if all of the capital was allocated to a small set of traders. My next task is to argue why concentration, or the distribution of risk bearing capital, is also important for pricing.

At least one reason to care about concentration is fragility. If CDS markets are dominated by a handful of important sellers, then a capital shock to one of these key players will have a sizable effect on the total amount of risk bearing capital, and presumably, prices. A similar concept for macroeconomic growth has been studied recently by Gabaix (2011) and Kelly, Lustig, and Van Nieuwerburgh (2014).

To operationalize this idea in the data, I first compute the following share-weighted version of my buyer and seller capital variables:

$$\text{SWA-}OCF^s_{rt} := \sum_{c \in S_{r,t-1}} MS(c,r,t-1) \times \Delta V_{c,-r,t}$$

where again, $MS(c,r,t)$ is the market share of net selling by counterparty c in reference entity r at time t. $\Delta V_{c,-r,t}$ is the change in capital for counterparty c, and is the same variable that I used in Section 4. SWA-OCF^b_{rt} is the analogous measure, constructed for net buyers of reference entity r. Finally, I also construct equal-weighted counterparts for all of these measures, and these are denoted by EWA-OCF.

Table 2 presents the results of running variants of the following regression specification:

$$\Delta \log \left(CDS_{rt}\right) = a_r + a_{it} + \boldsymbol{\beta}_1^{'} \Delta \mathbf{Z}_{rt} + \zeta_s \text{SW-}OCF_{rt}^{s} + \zeta_b \text{SW-}OCF_{rt}^{b}$$

In column (2), I only include the share-weighted capital measures in the regression. Consistent with my findings in Section 4, when the share-weighted capital of sellers decreases, spreads rise. In column (3), I replace the share-weighted capital measures with their equal-weighted counterparts. In this case, both the buyer and seller estimates are measured precisely, as indicated by their standard errors. However, the magnitude of both capital measures is quite small, and the sign for seller capital flips. For instance, a one-standard deviation decrease in equal-weighted average seller capital lowers spreads by only 0.33 percent. The more salient takeaway from Table 2 is found in column (4), which includes both the share-weighted and equal-weighted capital measures. When both measures are included in the regression, the share-weighted capital of sellers is the dominant explanatory force. In terms of absolute magnitude, the impact of a one-standard deviation move in share-weighted capital has nearly three times the impact of a one-standard deviation move in equal-weighted seller capital. In addition, the point estimate on share-weighted seller capital is relatively unaffected by the inclusion of equal-weighed seller capital. If the distribution of capital was irrelevant for pricing, one would not expect the distinction between share-weighted and equal-weighted capital to be as important.

A natural way to interpret these results is that not all seller losses are equal — a capital loss at a large player has a bigger impact on spreads than a capital loss at a smaller player. Though this is perhaps not so surprising, it also highlights the fragility of the CDS market. Because there are only a handful of large sellers of default insurance, their capital position has an outsized impact on the pricing of corporate credit risk.

6 Robustness: The 2011 Japanese Tsunami as a Natural Experiment

I now turn to a natural experiment that will further establish a causal link between capital losses and CDS pricing. The event I focus on is the Japanese tsunami of March 2011, which was the result of a magnitude 9.0 earthquake off the coast of Tohoku. The tsunami occurred on a Friday, and had a significant impact on the risk of the country as a whole. For example, Japan's sovereign CDS spread increased by nearly 50 percent from 80 to 115 basis points on the following Monday. The Online Appendix 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, capital losses accrue to Hedge Fund A because the Japanese firms they have written protection on are now fundamentally more risky; however, the same does not hold true for Hedge Fund B. My hypothesized mechanism then suggests the U.S. firms for whom Hedge Fund A is large seller will experience increases in their CDS risk premiums. On the other hand, U.S. firms where Hedge Fund B is a large seller will not see their spreads rise.

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

6.1.1 Measurement

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

$$\begin{aligned} \Gamma_{S,r} &:= \sum_{c \in \mathcal{S}(r)} \left[\frac{NS(c,r)}{NO(r)} \right] \times NS(c, Japan) \\ \Gamma_{B,r} &:= \sum_{c \in \mathcal{B}(r)} \left[-\frac{NS(c,r)}{NO(r)} \right] \times NS(c, Japan) \end{aligned} \tag{6}$$

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. The negative sign in the definition of $\Gamma_{B,r}$ is just to make sure the weights are positive and sum to 1. 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 on U.S. reference entities.[14]

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

To tease out my main hypothesis, I estimate variants of the following cross-sectional regression:

$$\Delta \log(CDS_{r,1}) = a + \phi_1 \Gamma_{S,r} + \phi_2 \Gamma_{B,r} + \boldsymbol{\beta}' X_r + \varepsilon_r \tag{7}$$

[14] 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.

where X_r is a vector of observable reference entity characteristics that I will discuss shortly. $\Delta \log(CDS_{r,1})$ is the log-change in r's CDS spread in the week following the tsunami. To reiterate, I consider only U.S. reference entities. There are certainly identification issues with attributing changes in CDS spreads after the tsunami with high levels of Γ, as the regression (7) would suggest. One obvious example is that sellers with large Japanese exposures also specialize in U.S. reference entities that are fundamentally linked to the Japanese economy. In Section 5 of Online Appendix, I fully frame the identification issues and rule out this "specialization" hypothesis for both buyers and sellers of U.S. reference entities.

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

Because certain industries may have been more exposed to Japanese firms, X_r also contains a fixed effect corresponding to each reference entity's NAICS code. I also include level of CDS spreads for each reference entity on 3/11/2011 to allow for the possibility that Γ captures sellers/buyers who specialize in riskier credits. Finally, I include the 90-day running volatility of each reference entity's CDS spread (in log-changes); this allows for the possibility that reference entities who experienced large spread movements post-tsunami are those that have larger volatility. Table 3 summarizes the results of running variations of regression (7).

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)-(5) sequentially add other control variables and industry fixed effects. As expected, the coefficient on $\Gamma_{S,r}$ remains stable throughout. Interestingly, including an industry fixed effect in the regression has a negligible effect on the point estimate of $\Gamma_{S,r}$. This is because the regression (with full controls) also includes each firm's own equity return; thus, any information contained in the industry fixed effect is subsumed by the more granular information contained in individual equity returns.

To get a sense of magnitude, consider a U.S. reference entity whose sellers were in the 90th percentile in terms of their exposure to Japanese firms. Similarly, consider a U.S. reference entity whose sellers were in the 10th percentile. Firms in the 90th percentile saw their spread levels increase 2.5 percent, relative to the 10th percentile, in the week following the tsunami.[15]

As a placebo test, in columns (8) and (9) of Table 3, I replace CDS spread changes as the dependent variable in the regression with each firm's equity return in the week following the tsunami. The logic behind this placebo test is twofold. First, the equity holders for reference r are probably different than r's CDS sellers. Second, the capital market frictions that I have documented in CDS markets are much less likely to appear in equity markets because, presumably, capital can flow much faster to investment opportunities in equities. Thus, one would not expect losses in the CDS market to necessarily impact equity market pricing. Consistent with this intuition, columns (8) and (9) indicate that whether or not a reference entity's CDS sellers were exposed to Japan has no explanatory power for equity returns after the tsunami — the transmission mechanism appears to be specific to the CDS market.

6.3 Isolating Concentration

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

[15] i.e. $\Delta \log(CDS_{r\in 90,1}) - \Delta \log(CDS_{r\in 10,1}) = 0.025$ where, for instance, $\Delta \log(CDS_{r\in 90,1})$ is the log-CDS spread change for firms in the 90th percentile.

Japanese firms and S_2 had exposure of 10. In this example, the total exposure of r_A and r_B's sellers is the same since they have the same two sellers. Still, we might expect that r_A will be more sensitive to the shock of the tsunami because its primary seller had large exposure to Japanese firms.

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

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

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

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

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

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

where $\omega_{J,r}$ is the share of J in the net selling of r and X_r is the same set of reference entity controls used throughout this section. Fixing the seller and only varying J's share across reference entities allows me to focus on how concentration interacts with pricing. In addition, in Section 5 of the Online Appendix

I verify that $\omega_{J,r}$ is once again not just a proxy for reference entities with high fundamental exposure to the Japanese economy.

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

7 Conclusion

This paper uses detailed data on CDS transactions to study the behavior of asset prices when capital cannot flow frictionlessly to investment opportunities. My evidence strongly suggests that the capital of a small set of concentrated CDS protection sellers plays a significant role in determining CDS price dynamics. More specifically, I show that a firm's short run CDS spread fluctuations are partially driven by the capital of the CDS desks at financial institutions who provide default insurance on the firm. These findings imply that the CDS market for a given name is segmented in the short run. More broadly, my findings are consistent with a cascading model of capital markets where the depth of segmentation increases at shorter horizons.

In addition, my results suggest that internal capital market frictions at financial institutions — whether due to agency issues, optimal risk management, or simple lack of capital — can act as an additional layer of segmentation in markets where outside capital is slow to enter.[16] In this sense, one can view the trading desks at large financial institutions as individual silos whose capital base is not instantaneously integrated with the larger firm. These types of segmentation issues are most likely to impact pricing in asset classes where investment activity requires stable funding, specific trading technology, and expertise.

I also document additional ways in which the structure of the CDS market

[16] See the Online Appendix for a deeper discussion of why capital may be slow moving.

impact CDS pricing. A defining feature of the CDS market is that net sellers of protection are highly concentrated. This concentration leads to fragility, as evidenced by the fact that capital losses at large sellers have a bigger impact on the market than losses at smaller sellers. Another way to frame this finding is that, in markets with segmentation, the distribution of risk bearing capital is an important consideration for pricing.

Heterogeneity in the type of financial institution that acts as a net seller of protection is a key determinant of spread dynamics as well. A striking trend in the data is that dealers have been replaced as the primary sellers of protections by hedge funds and asset managers. A likely explanation for this pattern is that new regulation has made it less profitable (or even possible) for dealers to ultimately bear credit risk via CDS. Still, the evidence in this paper indicates that capital losses at hedge funds and asset managers have a stronger impact on pricing than losses at dealers. Put differently, a potential unintended consequence of post-crisis regulation is that CDS prices are now even more influenced by the capital positions of a few players in the market, as opposed to fundamental risk exposures.

Figure 1: Notional Size of the U.S. CDS Market

Panel A (Gross Notional):

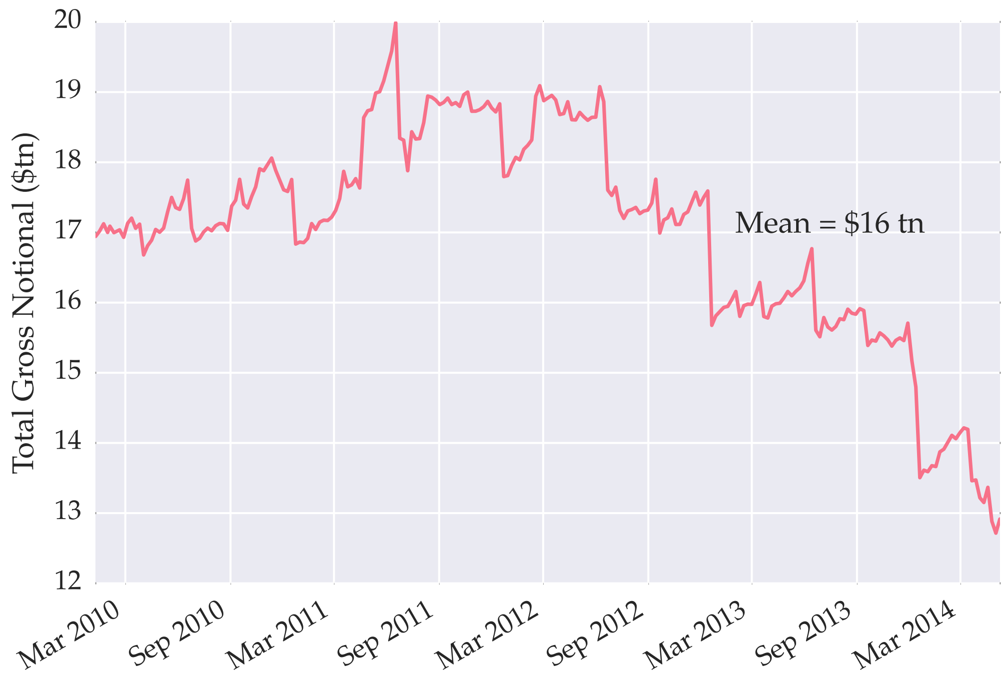

Panel B (Net Notional):

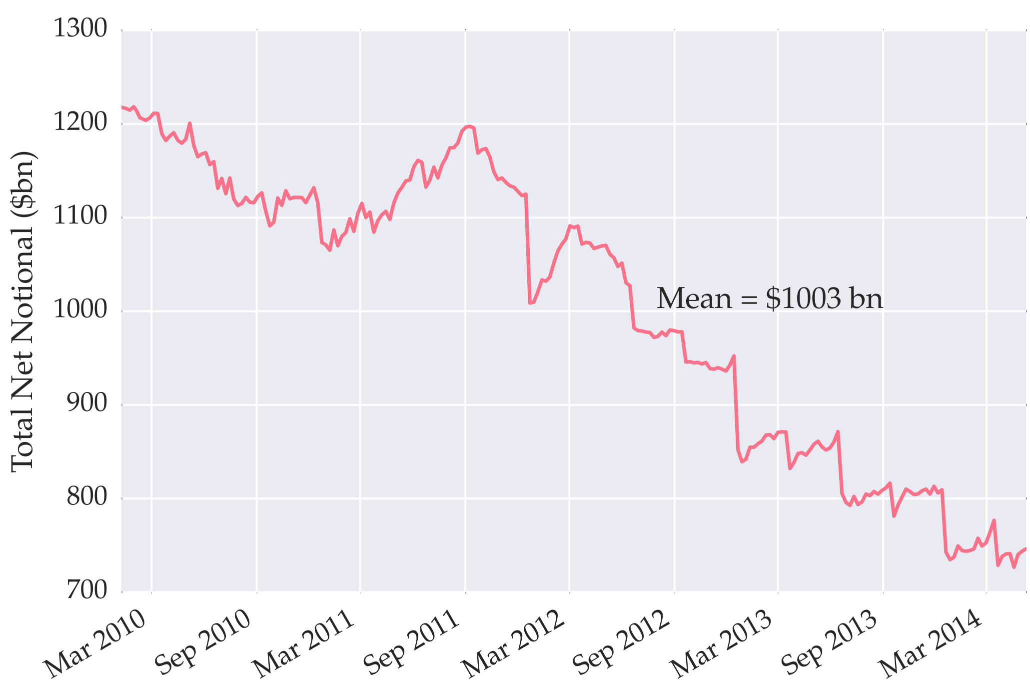

Notes: This figure plots the gross notional (Panel A) and net notional (Panel B) size of the U.S. CDS market. At each point in time, gross notional size is computed by adding the notional size of all positions written on U.S. reference entities. To determine total net notional size, I compute the net notional size for each U.S. reference entity, then sum over all reference entities. All computations account for positions that derive from index swaps. The Data Appendix contains details for the process through which I classify U.S. reference entities. Data is weekly and spans January 2010 to May 2014.

Figure 2: Aggregate Share and Persistence of Top Five Net Sellers and Buyers

Panel A (Aggregate Market Share):

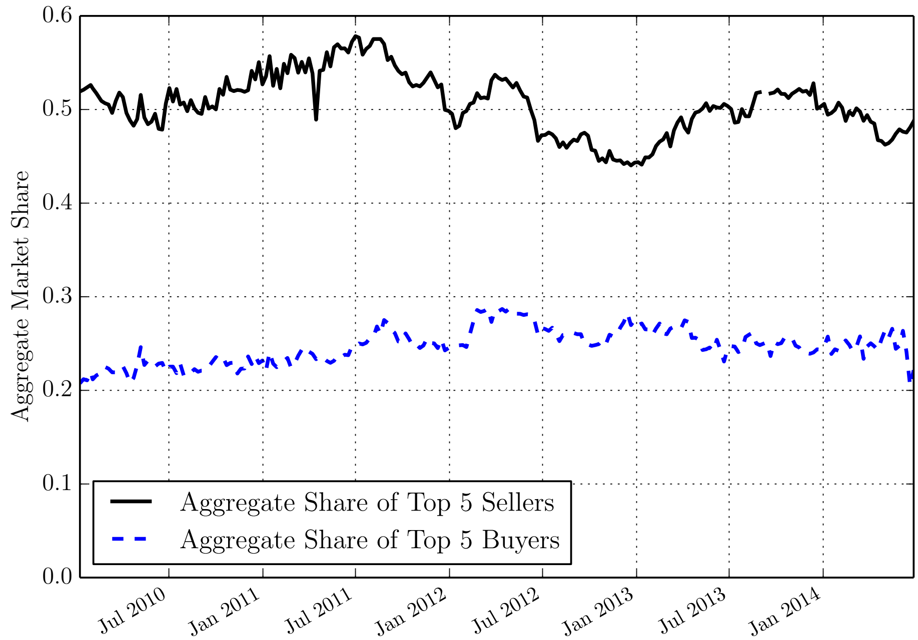

Panel B (Persistence):

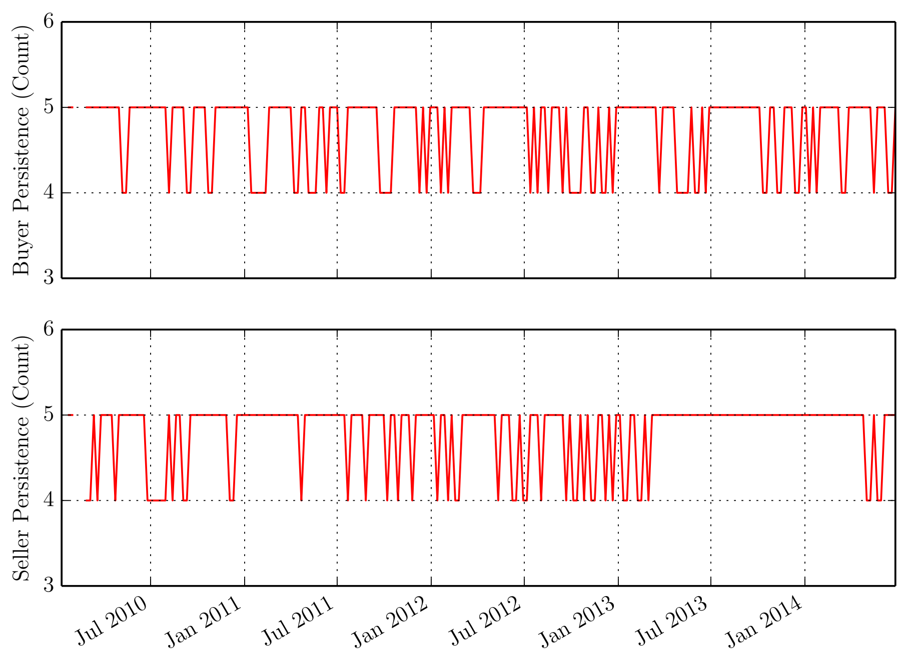

Notes: Panel A plots the aggregate share of the top five sellers and buyers of CDS protection through time. The share of a single counterparty c in a given reference entity is the proportion of net selling by c in that reference entity. The aggregate share of net selling by c is the size-weighted average 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. Panel B 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. Data is weekly and spans January 2010 to May 2014.

Figure 3: Aggregate Share of Net Buying and Net Selling by Counterparty Type

Panel A (Share of Buying):

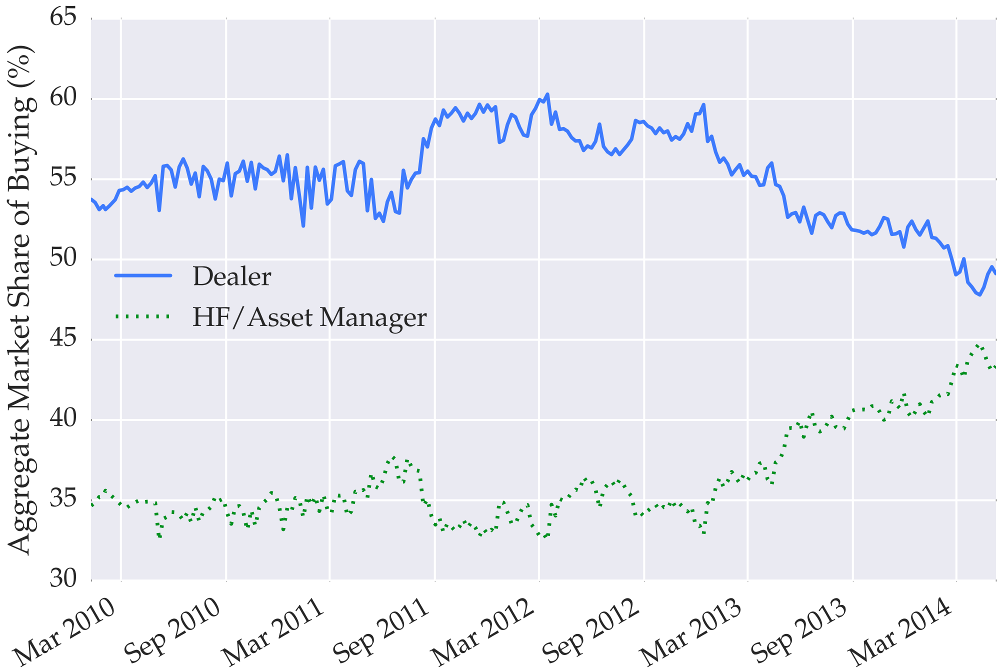

Panel B (Share of Selling):

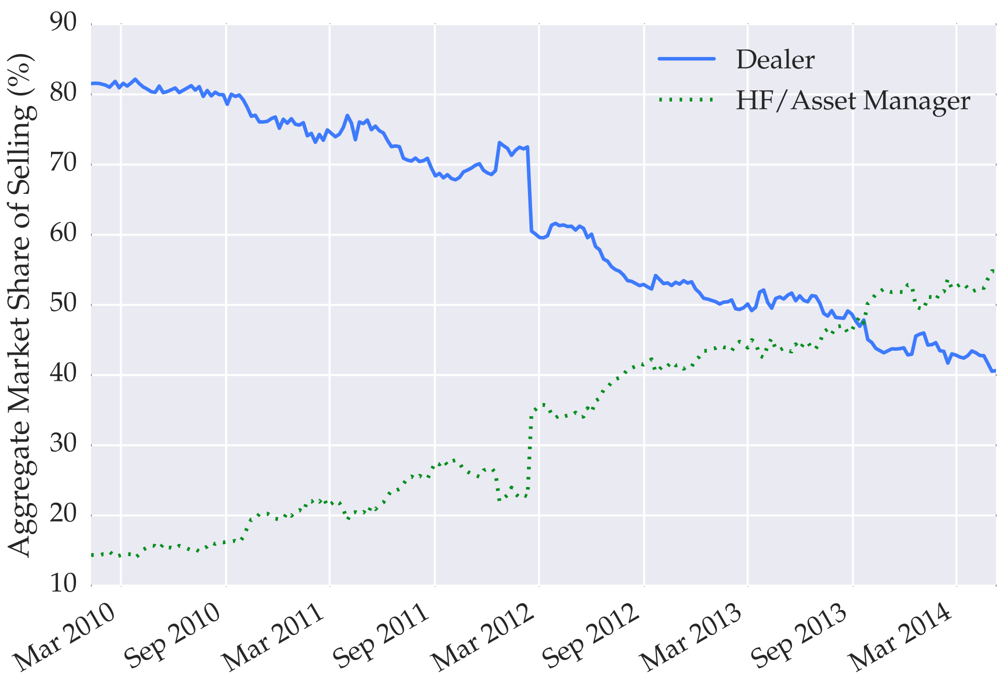

Notes: This figure plots the aggregate proportion of net buying (Panel A) and net selling (Panel B) by dealers and hedge funds/asset managers. For each reference entity, I compute the proportion of net buying and selling by each counterparty type. To aggregate, I compute the size-weighted average of the proportion bought by each type across reference entities. Data is weekly and spans January 2010 to May 2014.

Figure 4: The Impact of Capital on CDS Spreads Over Different Horizons

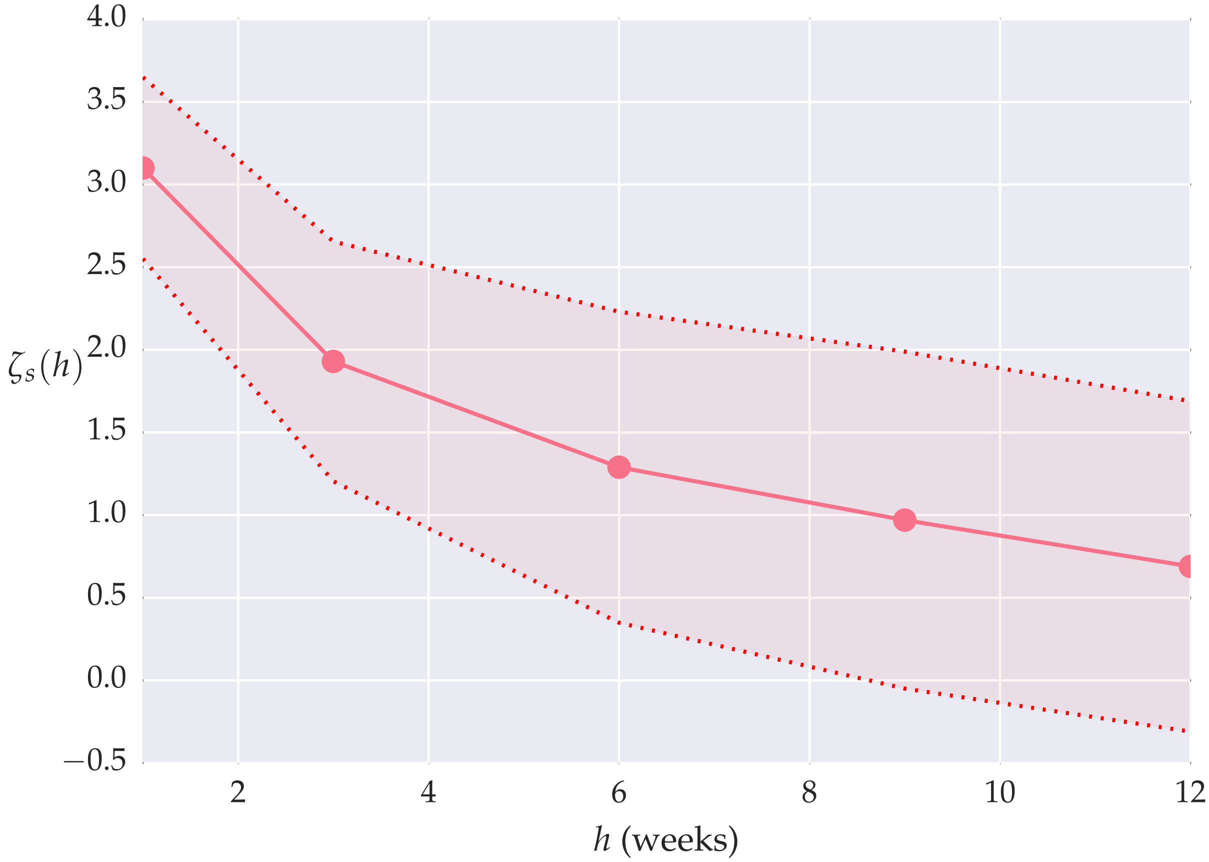

Notes: This figure plots $\zeta_b(h)$ from the following regression: $\log\left(CDS_{r,t+h-1}\right) - \log\left(CDS_{r,t-1}\right) = a_r + \boldsymbol{\beta}_1^{'}\Delta\mathbf{Z}_{rt}(h) + \boldsymbol{\beta}_2^{'}\Delta\mathbf{X}_t(h) - \zeta_s(h) \times OCF_{rt}^{s} + \zeta_b(h) \times OCF_{rt}^{b}$. CDS_{rt} is the 5-year CDS spread of reference entity r at time t, where the document clause of the contract is detailed in Data Appendix. OCF_{rt}^{s} measures the change from $t-1$ to t in capital ($bn) of r's net sellers, where capital is measured as the mark-to-market value of CDS positions coming from positions outside of r's industry. OCF_{rt}^{b} is the same variable, except for r's net buyers. The industry for each reference entity is defined by Markit. a_r is a reference entity fixed effect. $\mathbf{Z_{rt}}$ contains Moody's 5-year expected default frequency and Markit's estimate of loss-given-default, respectively. $\mathbf{X}_t$ includes the log equity-to-price ratio for the S&P 500, VIX, TED, CFNAI, 10 year Treasury yield, 10-year-minus-2-year Treasury yield, and the CBOE Option Skew index. Along the x-axis, I vary the horizon from $h = 1, 3, 6, 9, 12$. Note that because of the sign of $\zeta_b(h)$ in the regression, $\zeta_b(h)$ represents the response of CDS spreads to a billion dollar loss for large sellers. Standard errors are double-clustered by the top two sellers (alphabetized) and time, and are used to plot the confidence bands in the plot. Data is weekly and regressions are run over the period March 2010 to May 2014.

Table 1: The Impact of Capital Fluctuations on CDS Spreads

Dep. Variable	$\Delta \log(CDS_{rt}) \times 100$						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
OCF^s_{rt} ($bn)		**-3.10**		**-2.54**	**-2.34**		**-2.44**
		(0.28)		(0.25)	(0.27)		(0.26)
OCF^b_{rt}($bn)		0.13	0.13	0.04	0.05		0.13
		(0.33)	(0.28)	(0.34)	(0.35)		(0.34)
$\text{OCF}^s_{rt} \times 1_{\text{OCF}^s_{rt} \in \text{btm tercile}}$			**-3.57**				
			(0.31)				
$\text{OCF}^s_{rt} \times 1_{\text{OCF}^s_{rt} \in \text{mid tercile}}$			**-2.73**				
			(0.57)				
$\text{OCF}^s_{rt} \times 1_{\text{OCF}^s_{rt} \in \text{top tercile}}$			**-2.58**				
			(0.27)				
$\Delta\left(Leverage^s_{rt}\right)$						**0.09**	**0.05**
						(0.04)	(0.02)
$\Delta\left(Leverage^b_{rt}\right)$						**0.04**	0.13
						(0.02)	(0.34)
$\Delta\left(Leverage^s_{rt}\right) \times \text{OCF}^s_{rt}$							-0.02
							(0.01)
$\Delta \log\left(\text{EDF and LGD}\right)_{rt}$	✓	✓	✓	✓	✓	✓	✓
Δ Macro Variables$_t$	✓	✓	✓				
Ref. Entity FE	✓	✓	✓	✓	✓	✓	✓
Industry×Time FE				✓	✓	✓	✓
ΔOption-Based Measures$_{rt}$					✓		
Adj. R^2(%)	15.7	25.2	25.4	35.3	46.1	32.2	35.4
N	62,476	62,476	62,476	64,222	24,420	64,006	64,006

Notes: This table presents coefficients from regressions relating weekly (log) changes in CDS spreads to contemporaneous changes in capital for net sellers and net buyers in the CDS market. CDS_{rt} is the 5-year CDS spread of reference entity r at time t, where the document clause of the contract is detailed in Data Appendix. Log-changes in CDS spreads are expressed in percentage terms. OCF^s_{rt} measures the change from $t-1$ to t in capital ($bn) of r's net sellers, where capital is measured as the mark-to-market value of CDS positions coming from positions outside of r's industry. OCF^b_{rt} is the same variable, except for r's net buyers. EDF and LGD are the (log) change in Moody's 5-year expected default frequency and Markit's estimate of loss-given-default, respectively. Macroeconomic controls include the log equity-to-price ratio for the S&P 500, VIX, TED, CFNAI, 10 year Treasury yield, 10-year-minus-2-year Treasury yield, and the CBOE Option Skew index. The option based measures for each reference entity are: (i) the log of the option-implied CDS spread computed from option prices according to Carr and Wu (2011), and (ii) the implied volatility of at-the-money options. $Leverage^s_t$ is the share-weighted average market value of leverage for r's net sellers (if available), and $Leverage^b_t$ is the same for buyers. Industries for reference entities are defined according to Markit. Data spans March 2010 to May 2014. All standard errors are listed below point estimates and are double-clustered by the top two sellers (alphabetized) and time. Point estimates that are different from zero with 5% statistical significance are indicated in bold. The Adj. R^2 is computed within each reference entity group.

Table 2: How Seller Type and Concentration Impact CDS Spread Movements

Dep. Variable	$\Delta \log(CDS_{rt}) \times 100$			
	(1)	(2)	(3)	(4)
OCF^s_{rt} ($bn)	**-1.88**			
	(0.33)			
OCF^b_{rt} ($bn)	-0.19			
	(0.35)			
OCF$^s_{rt} \times HFS^s_{r,t-1}$	**-2.89**			
	(0.59)			
SWA-OCF$^s_{rt}$		**-1.48**		**-1.66**
		(0.12)		(0.11)
SWA-OCF$^b_{rt}$		-0.01		-0.09
		(0.05)		(0.06)
EWA-OCF$^s_{rt}$			**0.33**	**0.61**
			(0.11)	(0.07)
EWA-OCF$^b_{rt}$			**0.33**	**0.19**
			(0.10)	(0.07)
Controls	✓	✓	✓	✓
Adj. R^2(%)	35.7	34.4	32.3	35.1
N	64,222	64,222	64,222	64,222

Notes: This table presents coefficients from regressions relating weekly (log) changes in CDS spreads to contemporaneous changes in capital for net sellers and net buyers in the CDS market. CDS_{rt} is the 5-year CDS spread of reference entity r at time t, where the document clause of the contract is detailed in Data Appendix. Log-changes in CDS spreads are expressed in percentage terms. OCF^s_{rt} measures the change from $t-1$ to t in capital ($bn) of r's net sellers, where capital is measured as the mark-to-market value of CDS positions coming from positions outside of r's industry. OCF^b_{rt} is the same variable, except for r's net buyers. $HFS^s_{r,t-1}$ is the share of r's net selling by hedge funds and asset managers at time $t-1$. SWA-OCF^s_{rt} is the share-weighted average of the change in capital (excluding r's industry) for r's sellers, and SWA-OCF^b_{rt} is the same variable, defined for r's buyers. EWA-OCF^s_{rt} and EWA-OCF^b_{rt} are computed analogously, but instead use equal-weighted averages. All capital measures are standardized to have a mean zero and variance of one. Controls include: (i) the (log) change in Moody's 5-year expected default frequency and Markit's estimate of loss-given-default; (ii) industry-by-time fixed effects; and (iii) reference entity fixed effects. Industries for each reference entities are defined according to Markit. Data spans March 2010 to May 2014. All standard errors are listed below point estimates and are double-clustered by the top two sellers (alphabetized) and time. Point estimates that are different from zero with 5% statistical significance are indicated in bold. The Adj. R^2 is computed within each reference entity group.

Table 3: Transmission of Japanese Tsunami to U.S. CDS Markets

Dependent Variable	$\Delta\log(CDS_{r,1})\times 100$							$Ret^{equity}_{r,1}\times 100$	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
$\Gamma_{S,r}$	**3.23**		**3.35**	**3.44**	**3.00**		**3.07**	-0.57	0.92
	(1.00)		(1.00)	(0.94)	(1.37)		(1.27)	(0.83)	(0.84)
$\Gamma_{B,r}$		0.47	0.94	0.62	0.37		1.11	0.23	0.37
		(1.04)	(1.06)	(0.88)	(1.10)		(0.89)	(0.67)	(0.59)
$\Gamma^{avg}_{S,r}$						0.81	-0.77		
						(3.43)	(3.46)		
$\Gamma^{avg}_{B,r}$						-1.33	-1.47		
						(1.38)	(1.27)		
Control Variables				✓	✓	✓	✓		✓
Industry FE					✓	✓	✓		✓
Total N	288	288	288	288	288	288	288	288	288
Adj. R^2 (%)	2.3	0.0	2.4	23.2	20.4	17.5	21.0	0.0	13.7

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 from March 11, 2011 to March 18, 2011. $\Gamma_{S,r}$ 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 ($1 bn), meaning the units of $\Gamma_{S,r}$ and $\Gamma_{B,r}$ are in billions of dollar notional. The control variables are (for each reference entity r): the change in the 5-year Moody's expected default frequency (EDF), the change in Markit's loss-given-default, the weekly equity return, the 90-day trailing correlation of (changes in) r's CDS spread with the country of Japan's CDS spread, the 90-day trailing volatility of r's CDS spread, a fixed effect based on the NAICS code of each reference entity, and the level of the CDS spread for r on the day of the tsunami. In column (9), I exclude the change in 5-year Moody's EDF as a control because it effectively reflects equity returns. Bolded variables indicate statistical significance at the 5 percent level. Standard errors are clustered within each industry group and reported below point estimates. When industry fixed effects are included with the controls, the reported R^2 is within each industry group.

Table 4: Concentrated Exposures and Japanese Tsunami Transmission

Dependent Variable	$\Delta \log(CDS_{r,1})$
	(1)
$\omega_{J,r}$	**8.55**
	(3.00)
Control Variables	Yes
Total N	175
Adj. R^2	24.4%

Notes: The table presents results from the regression: $\Delta \log(CDS_{r,1}) = a + \eta_J \omega_{J,r} + \boldsymbol{\beta}' X_r + \varepsilon_r$. The dependent variable is the change in CDS spread for U.S. reference entities from March 11, 2011 to March 18, 2011. $\omega_{J,r}$ is the share of counterparty J in the net selling of reference entity r. J is the counterparty who had the largest exposure to Japanese firms prior to the tsunami. The regression includes only reference entities for which $\omega_{J,r} \neq 0$. The control variables are (for each reference entity r): the change in the 5-year Moody's expected default frequency, the change in Markit's loss-given-default, the weekly equity return, the 90-day trailing correlation of (changes in) r's CDS spread with the country of Japan's CDS spread, the 90-day trailing volatility of r's CDS spread, a fixed effect based on NAICS industry code, and the level of the CDS spread for r on the day of the tsunami. Bolded point estimates represent statistical significance at a 5 percent level. Standard errors are clustered within each industry group and reported below the point estimate. The reported R^2 is within each industry group.

References

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

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

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

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

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

Navneet Arora, Priyank Gandhi, and Francis A. Longstaff. Counterparty credit risk and the credit default swap market. *Journal of Financial Economics*, 103(2):280 – 293, 2012.

Andrew G. Atkeson, Andrea L. Eisfeldt, and Pierre-Olivier Weill. Entry and exit in otc derivatives markets. *Econometrica*, 83(6):2231–2292, 2015.

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

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

Jack Bao and Jun Pan. Bond illiquidity and excess volatility. *Review of Financial Studies*, 2013.

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

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

Peter Carr and Liuren Wu. A simple robust link between american puts and credit protection. *Review of Financial Studies*, 24(2):473–505, 2011.

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

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

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

Long Chen, Pierre Collin-Dufresne, and Robert S. Goldstein. On the relation between the credit spread puzzle and the equity premium puzzle. *Review of Financial Studies*, 22(9): 3367–3409, 2009.

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

Pierre Collin-Dufresne, Robert S. Goldstein, and J. Spencer Martin. The determinants of credit spread changes. *The Journal of Finance*, 56(6):2177–2207, 2001.

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

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

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

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

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

Darrell Duffie and Ken Singleton. Modeling term structures of defaultable bonds. *Review of Financial Studies*, 12(4):687–720, 1999.

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

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

Edwin J. Elton, Martin J. Gruber, Deepak Agrawal, and Christopher Mann. Explaining the rate spread on corporate bonds. *The Journal of Finance*, 56(1):247–277, 2001.

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

Erkko Etula. Broker-dealer risk appetite and commodity returns. *Journal of Financial Econometrics*, 2013.

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

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

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

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

Kenneth A. Froot, David S. Scharfstein, and Jeremy C. Stein. Risk management: Coordinating corporate investment and financing policies. *The Journal of Finance*, 48(5): 1629–1658, 1993.

Xavier Gabaix. The granular origins of aggregate fluctuations. *Econometrica*, 79(3):733–772, 2011.

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

Simon Gilchrist and Egon Zakrajsek. Credit spreads and business cycle fluctuations. *American Economic Review*, 102(4):1692–1720, 2012.

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

Robin Greenwood and David Thesmar. Stock price fragility. *Journal of Financial Economics*, 102(3):471 – 490, 2011.

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

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

Zhiguo He and Arvind Krishnamurthy. Intermediary asset pricing. *American Economic Review*, 103(2):732–70, 2013.

Zhiguo He, Bryan Kelly, and Asaf Manela. Intermediary Asset Pricing: New Evidence from Many Asset Classes. *SSRN*, 2015. URL http://ssrn.com/abstract=2662182.

Robert A. Jarrow and Stuart M. Turnbull. Pricing derivatives on financial securities subject to credit risk. *The Journal of Finance*, 50(1):53–85, 1995.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

A Appendix: Additional Computations

A.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.[17] 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:

$$CDS_{rt}(\tau) \quad = \quad \frac{\lambda_{rt}^{\mathbb{Q}}(\tau)}{\lambda_{rt}^{\mathbb{P}}(\tau)} \times LGD_{rt}^{\mathbb{Q}} \times \lambda_{rt}^{\mathbb{P}}(\tau) \tag{8}$$

where $LGD_{rt}^{\mathbb{Q}}$ is the loss given default under the risk-neutral measure. The ratio, $\Pi_{rt} := \lambda_{rt}^{\mathbb{Q}}(\tau)/\lambda_{rt}^{\mathbb{P}}(\tau)$, can be interpreted as the default premium for reference entity r. It quantifies the risk-reward tradeoff for bearing r's default risk.[18]

Taking the log of both sides of Equation (8) gives:

$$\log\left(CDS_{rt}\right) = \log\left(LGD_{rt}^{\mathbb{Q}}\right) + \log\left(\lambda_{rt}^{\mathbb{P}}\right) + \log\left(\Pi_{rt}\right) \tag{9}$$

where I have omitted the functional dependence of variables on the time to maturity. To make Equation (9) 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.[19]

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:

$$\log\left(LGD_{rt}^{\mathbb{Q}}\right) = \log(\eta_r) + \log\left(\widetilde{LGD}_{rt}^{\mathbb{Q}}\right) \tag{10}$$

When might this assumption be reasonable? For instance, the Moody's estimate of LGD

[17] 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).

[18] 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.

[19] 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.

is at the sectoral level. Assuming $LGD^{\mathbb{Q}}_{rt} = \eta_r \widetilde{LGD}^{\mathbb{Q}}_{rt}$ then means a firm's LGD is a time-invariant scalar transformation of the sectoral LGD. In other words, time-variation in reference entity LGD is common within a sector, which seems plausible.[20]

Substituting Equation (9) into Equation (10) yields:

$$\log\left(CDS_{rt}\right) = \log(\eta_r) + \log\left(\widetilde{LGD}^{\mathbb{Q}}_{rt}\right) + \log\left(\lambda^{\mathbb{P}}_{rt}\right) + \log\left(\Pi_{rt}\right) \tag{11}$$

Equation (11) 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 the firm-specific component of LGD, η_r. After controlling for firm specific variables, Equation (11) suggests the additional control variables that enter the regression capture the default risk premium for reference entity r, Π_{rt}.

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.[21] In the Online Appendix, I use actual bond yields to confirm my results do indeed pertain to the default risk premiums, as opposed to the CDS-bond basis.

A.2 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

[20] When $\widetilde{LGD}^{\mathbb{Q}}_{rt}$ comes from Markit, $\widetilde{LGD}^{\mathbb{Q}}_{rt}$ 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.

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

firm. The unit recovery claim is defined as follows:

$$U^O(t,T) = \frac{P_t(K_2,T) - P(K_1,T)}{K_2 - K_1} \tag{12}$$

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:

$$U^O(t,T) = \xi k \times \frac{1 - \exp\left(-(r + \xi k)(T - t)\right)}{r(t,T) + \xi k} \tag{13}$$

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 (13) 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.