

14-10 | December 22, 2014
Revised February 12, 2015
Revised June 4, 2015

Concentrated Capital Losses and the Pricing of Corporate Credit Risk: Data Appendix

Emil Siriwardane

Office of Financial Research

Emil.Siriwardane@treasury.gov

The Office of Financial Research (OFR) Working Paper Series allows members of the OFR staff and their coauthors to disseminate preliminary research findings in a format intended to generate discussion and critical comments. Papers in the OFR Working Paper Series are works in progress and subject to revision.

Views and opinions expressed are those of the authors and do not necessarily represent official positions or policy of the OFR or Treasury. Comments and suggestions for improvements are welcome and should be directed to the authors. OFR working papers may be quoted without additional permission.

Concentrated Capital Losses and the Pricing of Corporate Credit Risk*

Data Appendix

Emil Siriwardane[†]

June 4, 2015

Abstract

This is the data appendix to “Concentrated Capital Losses and the Pricing of Corporate Credit Risk”. In this appendix, I provide thorough documentation of how I use the underlying CDS data from the main paper.

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

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

1 List of Data Sources Used

This paper uses many different sources of data, which I list below. In addition, I describe the main uses in the paper of each data source:

- DTCC CDS Positional Data
 - Used to compute net exposures
 - Used to compute P&L for each trader
- DTCC CDS Transactional Data
 - Used as one identification strategy that looks at a single buyer purchasing protection from multiple sellers
- Markit CDS Pricing Data
 - As main dependent variable in tests of how capital impacts CDS pricing
 - Marking positions to market
 - Information on estimate loss-given-default
- Markit RED Information Data
 - Used to determine the constituents of each CDX index swap.
- Moody's EDF
 - Used as one measure of fundamental default risk
- OptionsMetrics
 - Used to compute option-implied CDS measures
- CRSP
 - Used in various instances when equity price information is needed

2 Building Basic Database of Positions

2.1 Background and Data Description

This section primarily describes how I build a database of positions that can be used to, among other things, compute next exposures of each counterparty, concentration measures (Section 2 of the main text), etc. The primary data source for this task originates from the Depository Trust & Clearing Corporation (DTCC). The DTCC provides trade processing and registration services for all major dealers in CDS markets. After a trade is registered with the DTCC, it is recorded into a Trade Information Warehouse (TIW). The data sample is a subset of the TIW, and includes trades that 1) involve a U.S. registered reference entity, or 2) involves at least one U.S. registered counterparty.¹ A simple example of the types of trades I am *not* privy to illustrates the scope of my CDS data: if a bank in Paris enters into a credit default swap with a bank in Germany on Portuguese sovereign debt, then I am not able to see this transaction. At a minimum, I can therefore make definitive statements about the nature of the U.S. Corporate CDS markets. I describe the process of filtering the database down to only U.S. based swaps in Section 2.3.

Additionally, the DTCC provides information on both transactions and accumulated positions. The primary data used in the paper is positions, but there is also a bit of analysis that uses the transactional data.

Defining a Reference Entity Depending on the context in the paper, I define a reference entity in one of two ways. The first way — and the predominant way — defines a reference entity simply by its Markit RedID[®], which I refer to henceforth as its RedID. Defining a reference entity by its RedID is useful for a broad understanding of credit risk. The second way refines the definition, and uses the RedID and the maturity of the swap. Specifically, I assign each swap into a “maturity bucket” based on its remaining term. The maturity buckets I consider are: 0-2 years, 2-4 years, 4-6 years, 6-8 years, 8-10 years, and 10+ years. A reference entity is then defined by its RedID and maturity bucket. This second way of defining a reference entity is useful for a more granular understanding of credit risk. When determining the overall size of the CDS market, I use the more granular definition since it is analogous to different maturity bonds on the same underlying firm. For the other computations in the main text, I collapse the term in order to make the computations more tractable.

Positions

The primary data source for this paper is the positional database. Positions represent the accumulation of all past transactions that contribute to a swap between counterparties. For instance, if Counterparty A and Counterparty B initiate a trade on date t , this also creates an existing position

¹In addition, for credit default swaps written on loans (typically called LCDX), we see the entire global market.

between the two counterparties. Suppose at some later date, $t + \tau$, both counterparties want to add more notional to their outstanding exposures. This is accomplished via a second transaction and can include termination of an offsetting position, an assignment from another counterparty, or a new trade, . Thus, at time $t + \tau$, their existing position gets modified as well. Another way to view transactions versus positions is as follows: transactions create a flow of exposures between counterparties, and positions keep track of the stock of exposure between two counterparties. The process of aggregating transactions to bilateral positions is handled by the DTCC according to their own internal algorithms. Positions naturally contain much of the same information as transactions (e.g. counterparty identity, reference entity, etc.).²

2.2 Building the Database

Step 1: Download Positions of Each Dealer and Central Clearing Parties The CDS markets are nearly 100% intermediated via dealers.³ Since all swaps are bilateral, it is therefore enough for me to examine swaps where a dealer is at least one of the counterparties in order to get a view of the *entire* market. To be safe, I also include any swap that is centrally cleared, though this is somewhat redundant given dealers are primary clearing members. On the Friday (close of business) of each week, the DTCC combines all of the existing positions with transactions from that week. The first day in the positional database begins on 1/1/2010, which represents the outstanding positions as of that date. The next date is 1/8/2010, and so forth. I refer to these weekly load dates as “as-of-dates”, since they represent the stock of positions as of that particular date. For each as-of-date in the database, I download the positions where dealers or the central-clearing party (CCP) are the “party” in the swap. Each swap has both a party and a counterparty; again, since swaps are bilateral, an identical swap appears in the database with the party and the counterparty reversed. By focusing on the swaps where a dealer or CCP is a party, I can also construct the aggregate positions for non-dealers since they will be listed as the counterparty.

The DTCC provides a classification of whether a counterparty is a dealer or not in its database. This is also the definition that I use as well. I cannot disclose the dealers who are classified as such by DTCC; however, a subset of this list can be found on DTCC’s website.⁴

²Because swaps are bilateral, each position is repeated twice. The party in a position either buys or sells protection, and this is mirrored by a symmetrical position for the other party in the swap. As a quality check of the data, I have verified that each position does indeed appear twice, once for each counterparty.

³As of 2/20/2014, 99.88% of all outstanding positions had at least one dealer as a counterparty in the swap.

⁴Note that each counterparty in a swap is associated with a PartyID that is defined by DTCC. A given dealer may have multiple PartyIDs. When aggregating positions, I consider the dealer level, which means I may sum/aggregate across positions spanning multiple PartyIDs.

Step 2: Disaggregation of Index Exposures

One of the main contributions of this paper is the so-called “credit risk topology” in CDS markets. In order to ascertain who ultimately bears corporate credit risk, it is vital to consider index exposures, as this constitutes a major way through which counterparties take on credit risk. The end goal is to take a position on a credit index and disaggregate it to the individual single name positions that are implied by the index exposure. Fortunately, this is quite easy given that each index is an equal weighted basket of its constituents.⁵ For example, due to how defaults on index exposures are settled, if I purchase \$100 of protection on a CDS index that contains 100 single names, then I am really purchasing protection of \$1 on each of the single names. Thus, to disaggregate a given index position I simply must ascertain each of the single name constituents for that particular index.

To obtain a list of the constituents (and weights) of each single name for each index, I use the Markit RED database. The Markit RED database contains index information such as constituents, weights, annex dates, inception dates, etc. for each of its published indices. I download all of the following index families: CDX North America, ITRAXX Asia, ITRAXX Europe, and ITRAXX SOVX. In addition, there are some swaps that do not have a RedID in the DTCC database. To address this case, I manually go through the set of swaps with missing RedIDs and use the field “Reference_Entity_Name_Unscrub” to try to infer what the correct RedID should be. I use the Markit RED database to get a list of indices, their RedIDs, their constituents, etc. I add the correct RedID to swaps for which I am able to confidently match a “Reference_Entity_Name_Unscrub” to a RedID. For a given date t and an index i , I then proceed as follows:

1. In order to ensure I use the correct weights for each constituent in the index, I first have to check if there has been a credit event on index i in between the beginning of the transaction and date t . If there has not been, then I use the original weights provided by Markit to create the single name exposures generated by i
2. If there has been a credit event for i , then I find the last version of the index whose annex date is before date t . The annex date for an index indicates which date its characteristics were made public on Markit’s website. Using the latest annex ensures that all credit events that have happened up to date t are accounted for. Call the index matching this criteria i^* . If a reference entity has defaulted, then its current weight in i^* is zero, so this is how I infer defaulted constituents. Based on this inference from i^* , I set the weights of any defaulted names in i equal to zero, and use the current weights in i for the remaining single name reference entities.⁶ I use the current weights in index i because the DTCC does not update

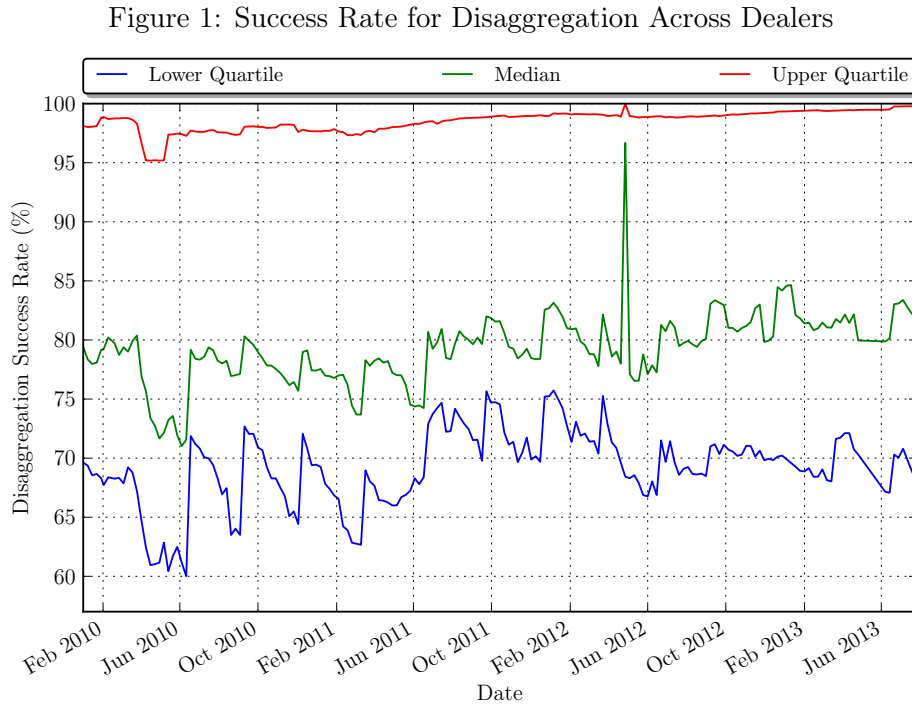
⁵To be precise, it is an equal weighted basket of each of the remaining non-defaulted constituents in it.

⁶The “current weights” provided by Markit correspond to the weights of the remaining non-defaulted constituents at the time of the index’s inception. For example, in an index of one hundred names, if one name defaults, a new version of the index is rolled out. The “current weights” for this new version become 1/99 for each member.

the outstanding notional on index contracts after default events.⁷ Disaggregation from this point then proceeds in the obvious fashion.

3. Finally, for indices that do not have a RedID or do not have a match in the Markit List, I simply do not disaggregate.

In order to verify that I am effectively capturing a large amount single name credit risk through this process, I plot the quartiles (across all dealers) of indices for which disaggregation succeeds.⁸ Note that this includes index exposures where no RedID was reported and I was not manually able to match a RedID.



Notes: This figure plots the quartiles of success rates of index disaggregation across twenty three dealers in the sample. The sample begins in January 2010 and ends in August 2013.

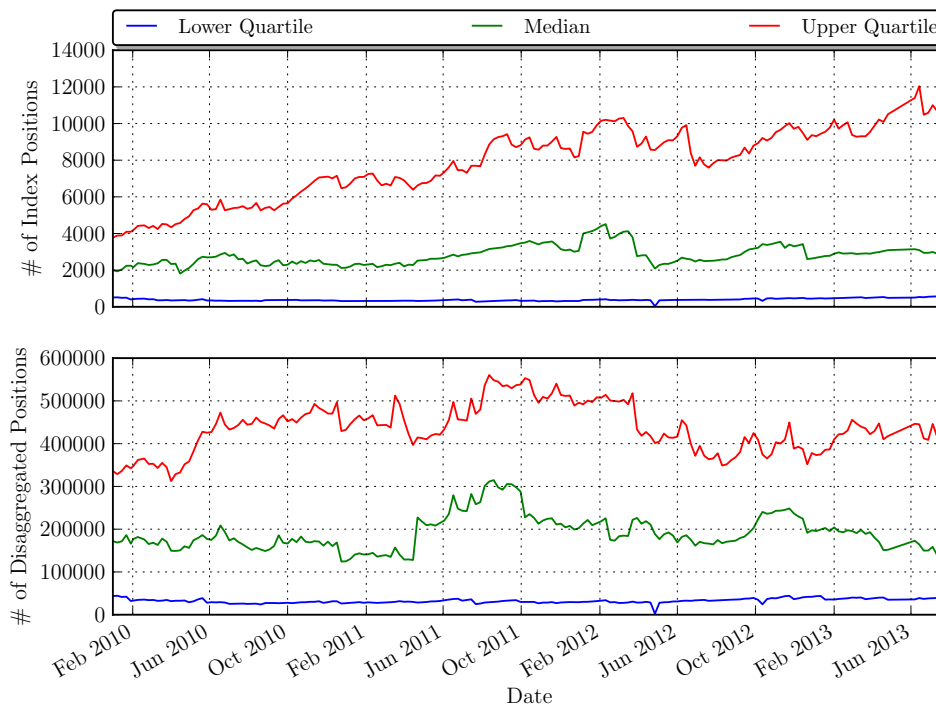
As is evident from Table 1, my methodology of disaggregating indices into single exposures is quite successful. Indeed, roughly 80% of all index exposures can be disaggregated further into single name exposures; thus, I conclude that the results of this process yield a reasonably comprehensive view of single name credit exposures in the CDS market. The importance of this procedure cannot be overstated, as a large portion of exposures in CDS market come through indices. To visually

⁷This is according to formal conversations with the DTCC.

⁸For each party ID in the DTCC database that is designated by DTCC as a broker-dealer, I aggregate up to the “bank” level. For example, if Bank A has three different party IDs that correspond to it, then I compute exposures for Bank A across all three party IDs. Within a bank, I take the average percentage of successful disaggregation. Finally, if there is a position extended within a bank (i.e. Bank A trades with Bank A), I ignore this.

emphasize this point, I plot the quartiles of simple counts of index exposure positions, as well as the resulting count of disaggregated positions in Figure 2

Figure 2: Index and Disaggregation Counts Across Dealers



Notes: This figure plots the quartiles of counts of 1) CDS index exposures (panel A) and 2) disaggregated index positions (panel B) across twenty three dealers in the sample. The sample begins in January 2010 and ends in August 2013.

Comparing the scales of the panels in the plot above, it is clear that there are roughly one hundred constituents in each index. In unreported results, the number of single name exposures that results from disaggregation is roughly equal to the number of direct single name exposures for most broker dealers. Disaggregating the indices is therefore crucial to understanding ultimate credit risk exposures.

2.3 U.S. Reference Entities

My analysis in Section 2 of the main text reports basic statistics on the size and concentration of the CDS market. In order to precisely define these statistics, I must be certain that I see the entire CDS market for a particular reference entity, which is why I restrict my analysis to only U.S. reference entities. For single name positions, the jurisdiction of the underlying reference entity is reported by DTCC. For swaps that arise from disaggregating index positions, I must determine the underlying jurisdiction myself. This subsection describes how I go about that process.

purposes of CDS pricing will typically include the term structure of riskless interest rates and the term structure of default likelihoods. Analogously, the value of the “default leg” is the risk-adjusted present value of the expected payouts to be made by the seller of protection. I denote the default leg by $DL_t(r, t_0, t_M, f; \mathcal{F}_t)$.

Let me begin with the value of the position at the contract’s inception. From the *perspective of the seller*, the net present value of the swap is therefore:

$$NPV_{t_0}^S(r, t_0, t_M, f) := PL_{t_0}(r, t_0, t_M, f; \mathcal{F}_{t_0}) - DL_{t_0}(r, t_0, t_M, f; \mathcal{F}_{t_0}) \quad (1)$$

The “fair value CDS spread” is the fixed coupon, f , that makes $NPV_{t_0}^S = 0$. In other words, it makes the net present value of the contract zero while ensuring that no money is exchanged upfront. This is the standard definition of a swap. In reality, the trading of CDS is done at fixed coupons, f , so that some money must be exchanged upfront in order for the contract to have an initial NPV of zero.¹² For example, if $NPV_{t_0}^S > 0$ then it means the value of the buyer leg is larger than the value of the seller leg; thus, in order to make the contract NPV zero, the seller must pay the buyer $NPV_{t_0}^S$ at time t_0 . This exchange is typically called the “upfront amount”.

Next, consider the value of the contract at time t :

$$NPV_t^S(r, t_0, t_M, f) := PL_t(r, t_0, t_M, f; \mathcal{F}_t) - DL_t(r, t_0, t_M, f; \mathcal{F}_t) \quad (2)$$

At this juncture, it is important to note that NPV_t^S is not necessarily the loss or gain to the seller on the position. The reason is that the net present value of the position may not have been zero at $t = t_0$. I will return to the issue of gains and losses shortly. When it is clear, I will drop the functional dependencies of key variables; for instance, I will refer to $PL_t(r, t_0, t_M, f; \mathcal{F}_t)$ simply as PL_t . I now document how I explicitly compute PL_t , DL_t , and hence NPV_t^S .

3.2.1 Default Intensity and Risk-Free Interest Rates

All credit derivative pricing requires some assumptions on the arrival rate of defaults. The seminal works of Duffie (1996) in pricing credit derivatives models defaults as Poisson processes that arrive at a specified rate, which is commonly called the default intensity. For reference entity r , the default intensity in the so-called physical measure is typically denoted by $\lambda_r^{\mathbb{P}}$. Similarly, the default intensity in the risk-neutral measure is denoted by $\lambda_r^{\mathbb{Q}}$.

Much of the research in credit derivatives focuses, in some way or another, on these default intensities. For instance, Berndt, Douglas, Duffie, Ferguson, and Schranz (2008) document how to use realized defaults to estimate $\lambda_r^{\mathbb{P}}$. On the other hand, a lot of research has been dedicated to estimating $\lambda_r^{\mathbb{Q}}$ from observed CDS spreads. Of course, the estimation procedure depends on the

¹²The fixed coupon, f , is a reported field in the DTCC positional data.

dynamics assumed for $\lambda_r^{\mathbb{P}}$ and $\lambda_r^{\mathbb{Q}}$.

For the purposes of this paper, I will be exclusively be focused on estimating $\lambda_r^{\mathbb{Q}}$ from observed CDS spreads. When the \mathbb{Q} superscript is absent, the reader can assume I am referring to the risk-neutral measure. I will assume the simplest dynamics for $\lambda_r^{\mathbb{Q}}$, as well as interest rates. To be clear, these assumptions are undoubtedly unrealistic, but my purposes do not require as much precision, and thus benefit from ease of computation. Specifically, I will assume that the default intensity is given to us exogenously. Moreover, at each point in time, the default intensity is constant over the remaining life of the position. The same assumption is made for the risk-free rate process, rf_t .

On any date t , we observe the CDS spread term structure for a variety of maturities. To incorporate maturity into my analysis, let $\lambda_{rt}(x)$ denote the default intensity up to time $t + x$. Given a position with remaining maturity τ_t , my goal is to estimate $\lambda_{rt}(\tau_t)$ from the observed CDS spreads for reference entity r on date t . First, I use a cubic spline to fit the observed term structure of CDS spreads. This allows me to compute the CDS spread for maturity τ_t , $CDS_{rt}(\tau)$.¹³ Under the assumption of constant default intensity and a non-stochastic loss-given-default, L_r , we obtain:

$$\lambda_{rt}^{\mathbb{Q}}(\tau_t) = \frac{CDS_{rt}(\tau_t)}{L_r} \quad (3)$$

In practice, there are a number of other alternatives to inferring the risk-neutral hazard rate curve.¹⁴ The underlying assumption I make is that, for a given maturity τ , there is a constant hazard rate that can be inferred from the CDS spread for τ (which itself may be interpolated from observed spreads).¹⁵

Similarly, I use the U.S. dollar swap rate curve to interpolate the default-free rate for a given maturity x , $rf_t(x)$. As per usual, this is assumed to be a continuously compounded rate. A more popular and accurate alternative is to “bootstrap” the risk-free rate curve recursively, but I avoid this route to make my computations simpler. With these two quantities in hand, I can now compute NPV_t^S .

¹³For more details, see the Data Appendix, Section 4.

¹⁴The ISDA standard uses a recursive method and a piecewise hazard rate function. I drop the risk-neutral designator here for convenience. For instance, suppose the first observed CDS is at date t_1 . Assume that $\lambda_s = \lambda_{[t, t_1]}$ for $s \in [t, t_1)$. Using the standard formula for a CDS spread, one can solve for $\lambda_{[t, t_1]}$. Then assume $\lambda_s = \lambda_{[t_1, t_2]}$ for $s \in [t_1, t_2)$. Using the observed spread for t_2 and $\lambda_{[t, t_1]}$, one can then solve for $\lambda_{[t_1, t_2]}$, and so forth.

¹⁵Let $\bar{\lambda}_\tau$ be the hazard rate inferred from using the CDS spread for τ . Suppose that $t + \tau \in (t_1, t_2)$ from the previous footnote. Compared to the piecewise approximation, I am essentially assuming the following approximation holds:

$$e^{-\lambda_{rt}(\tau) \times \tau} \approx e^{-\lambda_{[t, t_1]} \times (t_1 - t) - \lambda_{[t_1, t_2]} \times (t + \tau - t_1)}$$

3.2.2 Valuation

Let me begin by valuing the premium leg paid by the buyer to the seller on date t . Under my simplistic assumptions, the value of the premium leg is well-known and given by:

$$\begin{aligned} PL_t &= N \times f \times \int_0^{\tau_t} \exp\left(-\int_0^s (rf_t(\tau_t) + \lambda_{rt}^{\mathbb{Q}}(\tau_t))du\right) ds \\ &= \frac{N \times f}{rf_t(\tau_t) + \lambda_{rt}^{\mathbb{Q}}(\tau_t)} \left[1 - \exp\left(-(rf_t(\tau_t) + \lambda_{rt}^{\mathbb{Q}}(\tau_t)) \times \tau_t\right)\right] \end{aligned} \quad (4)$$

Analogously, the default leg is given by:

$$\begin{aligned} DL_t &= N \times L_r \times \int_0^{\tau_t} \lambda_{rt}^{\mathbb{Q}}(\tau_t) \times \exp\left(-\int_0^s (rf_t(\tau_t) + \lambda_{rt}^{\mathbb{Q}}(\tau_t))du\right) ds \\ &= \frac{N \times L_r \times \lambda_{rt}^{\mathbb{Q}}(\tau_t)}{rf_t(\tau_t) + \lambda_{rt}^{\mathbb{Q}}(\tau_t)} \left[1 - \exp\left(-(rf_t(\tau_t) + \lambda_{rt}^{\mathbb{Q}}(\tau_t)) \times \tau_t\right)\right] \end{aligned} \quad (5)$$

As previously defined, the net expected present value to the seller at time t is then $NPV_t^S = PL_t - DL_t$

3.3 Upfront Exchange in a CDS

At t_0 , the net present value of the CDS position will, in general, not be zero. This is because the fixed payment, f , is rarely equal to the “fair value spread”. As a reminder, the fair value spread would be the fixed coupon that would make $NPV_{t_0}^S = 0$.

Because $NPV_{t_0}^S \neq 0$, there must be some money exchanged upfront to make the transaction fair. When $NPV_{t_0}^S > 0$ it means that the value of the premium payments made by the buyer exceeds the expected payout of the seller to the buyer. This scenario indicates that the fixed coupon f is less than the fair value spread. In this case, the seller must pay the buyer this difference in order for the buyer to be willing to partake in the swap. The reverse holds when $NPV_{t_0}^S < 0$. Regardless of the sign of $NPV_{t_0}^S$, the upfront cash gained (or paid) by the seller must be equal $-NPV_{t_0}^S$.

3.4 Variation Margin, aka Mark to Market Gains, from a CDS

In standard trading agreements of OTC derivatives, both counterparties to a swap are responsible for maintaining what is called a variation margin.¹⁶ The variation margin reflects changes in the market value of the underlying derivative. A simple way to view the variation margin, or equivalently the mark to market gains/losses, is as follows: the balance in the variation margin account is meant

¹⁶In fact, this is the mechanism that caused AIG so much stress in the financial crisis. A position on which they had sold money lost a lot of value due to rising CDS spreads, and AIG was unable to make their variation margin payment.

to cover part of the costs of unwinding the trade at any given point in time. Based on how I have defined NPV_t^S , it represents the net profit to the seller if the position is terminated at that instant in time. Formally, the variation margin, or the mark to market gain, is defined from the perspective of the seller as:

$$MM_t^S = NPV_t^S - NPV_{t_0}^S \quad (6)$$

The situation depicted in Figure 3 is simple example that illustrates the mechanics and accounting of mark to market gains and losses. For now, it is enough to know that the fictitious swap has a maturity of one year. All NPV quantities correspond to the net present value from the perspective of the seller. A negative gain corresponds to a loss for the seller. For simplicity, I normalize $t_0 = 0$ and $t_M = 1$. Further suppose that the seller's name is Sally, and the buyers name is Barry. For now, ignore any accrued premiums so that we can focus solely on capital gains.


Let's start at $t = 0$. Since $NPV_0^S = -100$, it must mean that the fair value CDS spread is higher than the fixed coupon in the swap. Thus, the Barry must pay Sally \$100 in order for the total expected economic value of the contract to be zero. Typically, this money is placed into an escrow account, and whether it can be re-hypothecated depends on the terms of the swap. Initially, there is no money placed in the variation margin account. If Sally wanted to immediately close the position out, she could buy protection from Barry with the same terms, and would have to pay \$100 to do so. Sally would therefore take the proceeds from selling protection, and use them to buy protection; unsurprisingly, Sally's net profit if the position is unwound is zero.

At time $t = 0.5$, fair value CDS spreads have risen such that $NPV_{0.5}^S = -150$. Sally must pay to Barry \$50 in a variation margin. This payment can be interpreted as contributing to the cost of unwinding the position. To see why, note that in order to offset the existing position, Sally would need to buy protection from Barry, which would cost Sally \$150. \$50 of that \$150 comes from the variation margin; the remaining \$100 would be expected to come from the initial payment. Thus, the total out of pocket loss for Sally if she unwinds the position is \$50.

Finally, at expiration, the contract is, in some sense, unwound. An instant before expiration, Sally would have received \$100 in variation margin from Barry, which is meant to contribute to Barry's cost if he unwinds the position. However, since $NPV_1^S = 0$, Sally has received too much from Barry, so she must pay him the \$100 from the variation margin account back in order to close out the position fairly. This exchange makes the balance in the variation margin account zero, which makes sense. Sally still keeps the initial payment of \$100, which also makes sense since this was the amount that made the contract have an expected value of zero at its inception.

As the example illustrates, prior to expiration the balance in the variation margin corresponds to the net profit the seller makes from unwinding the position. In accounting terms, this is what is booked as profits and losses on a trading book. Importantly, it also represents the profits and losses that traders use as a mental accounting tool. It is for this reason that I treat the mark to market

Figure 3: Illustration of Mark to Market Losses



Account	Seller's Balance
Initial Payment	100
Variation Margin	0
Additional Gain to Seller if Position Closed	-100
Net Profit if Closed	0

Account	Seller's Balance
Initial Payment	100
Variation Margin	-50
Additional Gain to Seller if Position Closed	-100
Net Profit if Closed	-50

Account	Seller's Balance
Initial Payment	100
Variation Margin	100
Additional Gain to Seller if Position Closed	-100
Net Profit if Closed	100

Additional Gain to Seller if Position Closed = Gain from Offsetting Position – Variation Margin Contribution

gain as my measure of risk aversion.¹⁷

To this point, I’ve ignored any accrued premiums from selling the CDS in the first place. Based on my preceding analysis, I therefore define the total profit of the trade as the mark to market gain plus any accrued premiums:

$$V_t^S := NPV_t^S - NPV_{t_0}^S + N \times f \times (t - t_0) \quad (7)$$

By means of convention, I let a default event be the limiting case where $\lambda_{rt}^{\mathbb{Q}}(\tau_t) \rightarrow \infty$, which naturally means that $\lim_{\lambda_{rt}^{\mathbb{Q}}(\tau_t) \rightarrow \infty} NPV_t^S = N \times L_r$. Again, the way to view V_t^S — and one which is often used by traders — is the net profit to the seller if the CDS position is reversed at time t . Naturally, if a counterparty is the buyer in the position, then their total gain is $-V_t^S$.

Lastly, I have defined V_t^S for a single position and from the perspective of the seller; thus, the profit of all positions is simply the sum of the individual profits (and profits to sellers are interchangeable with losses to buyers), taking into account which side of the transaction the counterparty is on (e.g. buyer or seller).

3.5 Actual Implementation

To improve the computational efficiency of marking each position in the database to market, I first determine how many unique positions are in the database. Importantly, this occurs after I have disaggregated each index exposure. A position is uniquely defined by the transaction tag that spawns the position. All together, there are 132,774,580 unique positions. Of those, I am able to compute the initial NPV for 128,410,438, or 96.7% of them.

To further ease the process, I cycle through all positions and get the unique set of initial positions. Since positions roll through time, open positions appear each week in the data set, but to compute the initial value, I only need the information on the position once. Once I generate the set of unique initial positions, I compute their market value as with all other positions.

4 Data Used for Section 4 in Main Text

CDS Spreads and Loss-Given-Default

For the analysis in Section 4 of the main text, I obtain a weekly time series of 5-year CDS spreads and expected loss-given-default (LGD) from Markit for as many reference entities that are in my positional data set from January 2010 to May 2014. I focus on CDS written on senior unsecured debt in U.S. dollars, and use swaps that have a “modified restructuring” clause in their definition of what

¹⁷At expiration, the position is closed out by zeroing the variation account balance and the seller retains (or the buyer retains) the initial upfront amount. Mathematically, it is easy to see that, excluding premiums, this sequence of transactions ultimately results in a profit of $NPV_{t_M}^S - NPV_{t_0}^S = -NPV_{t_0}^S$ for the seller of protection.

triggers default in the CDS contract. In the Markit database, this corresponds to a restructuring clause of “MR”. I choose this restructuring clause for two reasons. The first is that it is the second most abundant contract over the course of my sample period.¹⁸ The second is that I will eventually compare CDS spreads to Moody’s EDF metric. Based on the definition of default in Moody’s computations, the modified restructuring clause in CDS contracts is comparable.¹⁹ Because my sample spans over four years at a weekly frequency, each reference entity must have at least 200 observations over this time period in order to be included in the subsequent analysis.

Moody’s EDF

When possible, I also obtain the 5-year annualized EDF series from Moody’s. To match RedIDs from Markit to Moody’s, I use the 6 digit CUSIP.

5 Details on Building Panel Data in Section 6 of Main Text

5.1 Converting Observed Upfront Amounts to Default Intensities

I assume a constant default intensity. From here, it is straightforward to convert an observed upfront amount to an implied default intensity. At the initial point of a transaction, the upfront paid by the buyer to the seller (can be negative) is given by:

$$UP_{t_0} = DL_{t_0} - PL_{t_0} \quad (8)$$

where, DL_t and PL_t are defined as:

$$\begin{aligned} PL_t &= N \times f \times \int_0^{\tau_t} \exp\left(-\int_0^s (rf_t(\tau_t) + \lambda_{rt}^{\mathbb{Q}}(\tau_t)) du\right) ds \\ &= \frac{N \times f}{rf_t(\tau_t) + \lambda_{rt}^{\mathbb{Q}}(\tau_t)} \left[1 - \exp\left(-(rf_t(\tau_t) + \lambda_{rt}^{\mathbb{Q}}(\tau_t)) \times \tau_t\right)\right] \end{aligned} \quad (9)$$

Analogously, the default leg is given by:

$$\begin{aligned} DL_t &= N \times L_r \times \int_0^{\tau_t} \lambda_{rt}^{\mathbb{Q}}(\tau_t) \times \exp\left(-\int_0^s (rf_t(\tau_t) + \lambda_{rt}^{\mathbb{Q}}(\tau_t)) du\right) ds \\ &= \frac{N \times L_r \times \lambda_{rt}^{\mathbb{Q}}(\tau_t)}{rf_t(\tau_t) + \lambda_{rt}^{\mathbb{Q}}(\tau_t)} \left[1 - \exp\left(-(rf_t(\tau_t) + \lambda_{rt}^{\mathbb{Q}}(\tau_t)) \times \tau_t\right)\right] \end{aligned} \quad (10)$$

¹⁸The count of available 5-year USD contracts on senior unsecured debt with a United States jurisdiction are: CR (495,075); MM(220,305); MR (1,036,081); XR (1,083,786).

¹⁹See: <http://www.moodyanalytics.com/~media/Insight/Quantitative-Research/Default-and-Recovery/2012/2012-28-06-Public-EDF-Methodology.ashx>

where $t = t_0$ to complete the computation. Notice here that we are trying to infer $\lambda_{rt_0}^{\mathbb{Q}}(\tau_{t_0})$ by observing UP_{t_0} , the maturity τ_{t_0} , the fixed coupon f , and the notional amount N . For all computations, I assume a loss-given-default of 60%. Because my panel regression includes a fixed effect that accounts for (r, t) , this is an innocuous assumption. To solve for the implied hazard rate, I use standard numerical methods to invert the right hand side of Equation (8). The risk-free rate is, as with the rest of the paper, derived from the daily US swap rate curve.

Finally, note that because I run my regression in logs, converting the implied hazard rate to a CDS spread simply changes the constant in the regression (under the constant hazard rate assumption). As a result, the results presented in the main text are run using the implied hazard rate as the dependent variable.

5.2 Data Filters

I include transactions in my final panel so long as they meet all of the following criteria:

- There are exactly two records of the transaction in the data (both sides report the trade)
- The buyer and seller in the transaction are the two counterparties listed in the transaction
- There is a non-null initial payment in the transaction
- The termination date is non-null
- The implied hazard rate is strictly positive and less than 30%
- The maturity is strictly positive

I further exclude transactions where $|RB_t^s| < 0.5$. This discards transactions by counterparties with very low gross numbers, so that their annualized portfolio return (with no leverage) is disproportionately large. The reason I exclude these firms is directly related to my discussion of leverage in the main text. Assuming that each counterparty can take the same leverage is not likely to be true in practice, and the outliers of these returns are very likely to be driven by the fact that I am treating the mega-players in the market as if they can take the same leverage as the small players. The results are qualitatively the same without this restriction, and are available upon request.