The Transaction Costs of Trading Corporate Credit *

Gopa Biswas† Stanislava Nikolova‡ Christof W. Stahel§

August 2014

ABSTRACT

We estimate individual roundtrip transaction costs for 851 single-name CDS contracts traded between August 2009 and May 2014. Effective half-spreads are 14 bps of the notional amount for dealer-to-enduser trades and 12 bps for dealer-to-dealer trades for the most common notional amount traded, \$2.5-7.5M. In the cross-section, effective spreads are smaller than and only weakly correlated with indicative quoted spreads. Effective spreads tend to increase with market activity consistent with the notion that the CDS market attracts informed trading. When we compare CDS transaction costs to those of the underlying bonds over a range of trade sizes where trading in the two markets overlaps, we find that CDSs are typically cheaper. At smaller trade sizes bonds are more than three times more expensive to trade than the CDS contracts written on them, but at larger trade sizes this pattern reverses.

Keywords: Credit Default Swaps, Bonds, Transaction Costs, Liquidity
JEL Classification: G11, G12, G14

---

*We thank Scott Bauguess, John Cooney, Sanjiv Das, Donna Dudney, Amy Edwards, Mark Flannery, Hans Heidle, Pankaj Jain, Amar Kuchinad, Scott Murray, Nimal Nimalendran, Hari Phatak, Dragon Tang, Laura Tuttle and seminar participants at the SEC, Texas Tech University and University of Nebraska-Lincoln for useful comments that have substantially benefited the paper. Bulgan Batsaikhan and Viet Nguyen provided excellent research assistance. All remaining errors are our own.

†Office of Financial Research, U.S. Department of the Treasury, Washington, DC 2005. Phone: (202) 927-8455, Email: gopa.biswas@treasury.gov. Views and opinions expressed are those of the authors and do not necessarily represent official OFR or Treasury positions or policy.

‡Department of Finance, University of Nebraska-Lincoln, Lincoln, NE 68588. Phone: (402) 472-6049, Email: snikolova2@unl.edu.

§Division of Economic and Risk Analysis, U.S. Securities and Exchange Commission, Washington, DC 20549. Phone: (202) 551-6623, Email: stahelc@sec.gov. The Securities and Exchange Commission disclaims responsibility for any private publication or statement of any SEC employee or Commissioner. This paper expresses the author's views and does not necessarily reflect those of the Commission, Commissioners, or other members of the staff.

I. Introduction

Facilitating price discovery is arguably one of the most important functions of financial markets, though market frictions can impede it. For closely-linked securities trading in different markets, price discovery will likely take place in the market characterized by fewer frictions. A number of recent studies document that price discovery in corporate credit has shifted from the cash to the single-name credit default swap (CDS) market and some authors have concluded that the CDS market must be more liquid and cheaper to trade in than the corporate bond market.[1] However, CDS data limitations have made testing the validity of this conclusion difficult. Despite the high level of trading activity,[2] the market for CDSs remains an opaque over-the-counter market with little pre-trade and post-trade transparency. As a result, the cost of trading CDSs is not well known to those outside the circle of dealers and major market participants who actively trade them.

Quoted bid-ask spreads can provide some indication of CDS transaction costs and more generally CDS liquidity. However, in many dealer markets quoted prices are simply a starting point for negotiations between dealers and customers, and trades often occur at prices other than the quoted bids and asks (e.g. Bessembinder and Venkataraman (2010)). Moreover, available CDS quotes are typically indicative and do not represent a binding commitment by a dealer to trade at these prices. As a result, an indicative quoted bid-ask spread may significantly differ from the transaction costs actually incurred by investors. We are aware of only one study that estimates effective bid-ask spreads, Fulop and Lescourret (2009), which reports average roundtrip transaction costs

[1] See Longstaff, Mithal, and Neis (2005), Blanco, Brennan, and Marsh (2005), Ashcraft and Santos (2009), and Das, Kalimipalli, and Nayak (2014) among others.

[2] DTCC Trade Information Warehouse Data, Table 17: Summary of Weekly Transaction Activity, available at http://www.dtcc.com/products/derivserv/data_table_iii.php.

of 64 bps from 2006 to 2008 for six actively traded contracts on U.S. reference entities. In comparison, studies of the U.S. corporate-bond market report estimates of significantly smaller roundtrip transaction costs for institutional-sized trades.[3] In this paper we attempt to resolve this seeming inconsistency between the perception of some academics that trading corporate credit is cheaper in the CDS than in the cash market, and the effective bid-ask spread estimates found in empirical studies of the two markets.

We estimate transaction costs for a sample of 851 single-name CDSs traded from 08/01/2009 to 05/31/2014. Our model of transaction costs is similar in spirit to models used to estimate bond transaction costs in Schultz (2001), Bessembinder, Maxwell, and Venkataraman (2006), Goldstein, Hotchkiss, and Sirri (2007), Harris and Piwowar (2006), and Edwards, Harris, and Piwowar (2007). In general, in these models the difference between an observed trade price and a contemporaneous fundamental value estimate is regressed on a buy-sell indicator of the trade. The coefficient on the buy-sell indicator is the estimate of transaction costs. We modify these models to take into account the availability of CDS quotes and use the midpoint of the indicative quoted bid-ask spread to proxy for CDS fundamental value. As in Harris and Piwowar (2006) and Edwards, Harris, and Piwowar (2007) we allow transaction costs to vary with trade size. However, unlike them we specify a model that permits for the estimation of transaction costs between dealers and customers (referred to as endusers in the CDS market) and between dealers and dealers. This is important because dealer-to-dealer trades comprise the majority (73%) of trades in our sample. Specifically, we first estimate dealer-to-enduser and dealer-to-dealer transaction-cost functions for each CDS contract using traded and quoted CDS prices. We then evaluate the estimated functions at each

[3] See Hong and Warga (2000), Schultz (2001), Chakravarty and Sarkar (2003), Bessembinder, Maxwell, and Venkataraman (2006), Goldstein, Hotchkiss, and Sirri (2007), and Edwards, Harris, and Piwowar (2007).

CDS's actual trade sizes.[4] We document that effective half-spreads are on average 12 bps of the notional amount for dealer-to-dealer trades and 14 bps for dealer-to-enduser trades for the most common notional amount traded, between \$2.5 and \$7.5 million. These estimates are significantly smaller than those reported in Fulop and Lescourret (2009) and consistent with the perception that single-name CDSs are indeed very liquid and relatively cheap to trade.

Next, we examine the relationship between our estimates of CDS transaction costs and other measures of CDS liquidity. We recognize that the concept of liquidity is multifaceted and there is no single metric that fully describes the liquidity of an asset. We focus on measures commonly used in the literature: quoted bid-ask spread, number of trades, number of quotes, quote matching intensity (ratio of trades to quotes), and quote depth (number of quote-submitting dealers). We document several interesting patterns in the relationship between these measures and CDS transaction costs. First, the average effective spread is less than half of the indicative quoted spread, but CDS trades often take place outside the indicative quoted spread. Second, quoted spreads generally explain little of the cross-sectional variation in effective spreads, except for contracts written on non-investment grade bonds. Third, we find that effective spreads tend to be larger for contracts characterized by higher level of market activity. This is consistent with Tang and Yan (2007), who show that adverse selection in an important concern in the pricing of actively traded contracts.

Finally, we show that CDS transaction costs are typically lower than the transaction

[4]Harris and Piwowar (2006) and Edwards, Harris, and Piwowar (2007) evaluate the transaction-cost functions they estimate at hypothetical trade sizes, while Jostova, Nikolova, Philipov, and Stahel (2013) use actual trade sizes. We investigate the robustness of our findings to following either approach in Section V.A.

costs of the bonds on which they are written.[5] We obtain trading data for 739 bonds underlying the CDSs in our sample and compare trading activity in the bonds to that in the CDSs. We find that small trades tend to be executed in the cash market while large trades tend to be executed in the CDS market. This is consistent with the view that the latter is an institutional market, and that CDS trading is not constrained by bond amount outstanding. Substantial overlap in CDS and underlying-bond trading occurs in our sample only over the $100K-$5M trade-size range. Over this range we estimate the bonds' dealer-to-customer transaction-cost functions as in Edwards, Harris, and Piwowar (2007) and evaluate them as in Jostova, Nikolova, Philipov, and Stahel (2013). We then compare the bond transaction costs to the dealer-to-enduser transaction costs of the corresponding CDSs. We find that at smaller trade sizes bonds are three times more expensive to trade than the CDSs written on them. However, this difference decreases as trade size increases. Surprisingly, for the largest trade sizes for which we can estimate both bond and CDS transaction costs, trades are cheaper to execute in the cash than the CDS market. These results are similar regardless of whether we compare CDS transaction costs to the reference bond's transaction costs, the average costs across the reference entity's traded bonds, or to the costs of the reference entity's cheapest-to-trade bond.

We contribute to the literature in three important ways. First, our study is the first to estimate single-name CDS transaction costs using newly available transaction data from Depository Trust & Clearing Corporation (DTCC). DTCC is the only central trade registry for CDSs and covers 90% of the notional amount traded. Hence, our sample

[5]In theory, an investor can make the same economic trade in the cash or derivative market by either adding a risk-free floating-rate note to the CDS, or subtracting it from the bond (Duffie (1999)). In either case, an investor will incur the costs of trading in Treasuries, which Fleming, Mizrach, and Nguyen (2014) estimate to be 0.9 bps for a 5-year note roundtrip. Since these costs are negligible relative to the CDS and corporate-bond effective spreads we estimate, we neglect them when we compare CDS to corporate-bond transaction costs.

reflects the vast majority of CDS market activity. Second, we show that indicative quoted bid-ask spreads are significantly different from effective spreads, and that the cross-sectional variation in indicative quoted bid-ask spreads does not fully capture the variation in actual transaction costs. This finding has important implications for studies that use indicative quoted bid-ask spreads to proxy for CDS transaction costs or more generally CDS liquidity. Finally, we are the first to compare single-name CDS transaction costs to the transaction costs of the bonds they reference. Our finding that CDSs are typically much cheaper to trade provides one explanation for why price discovery and trading corporate credit have shifted from the cash to the CDS market as documented in Blanco, Brennan, and Marsh (2005), Norden and Wagner (2008), Norden and Weber (2009), Forte and Pena (2009), and Das, Kalimipalli, and Nayak (2014).

Our paper also has relevance to CDS and bond investors. Trading strategies are costly to implement and when evaluating a strategy's profitability it is important to account for the anticipated costs of the required trades. In the CDS market, these costs have been largely unknown. In addition, if investors want to take a view on corporate credit, they can do so in the cash or CDS market. Thus, it is important for them to know the relative costs of executing trades in these two markets and the factors affecting the costs. Our CDS-bond transaction cost comparisons might be of particular interest to arbitrageurs of the CDS-bond basis, the difference between the CDS spread of a specific firm and the credit spread paid on a bond of the same firm.[6]

Finally, our findings have important implications for regulators, who have expressed continued interest in improving CDS market quality. Transaction costs are one aspect of market quality and our finding that trading corporate credit is generally more expensive

[6]For more details on the magnitude and time-series pattern of profitability of CDS-bond basis arbitrage see Bai and Collin-Dufresne (2013), Blanco, Brennan, and Marsh (2005) and Nashikkar, Subrahmanyam, and Mahanti (2011) among others.

in the cash than in the CDS market suggests that the marginal benefits of further transparency in the largely institutional CDS market with already low transaction costs may be limited. Regulatory efforts to further improve the microstructure of the corporate-bond market might be of more value.

The rest of the paper is organized as follows. Section II provides background information on the CDS market. Section III explains our transaction cost estimation methodology. Section IV describes the data we use and presents summary statistics for our sample. Section V summarizes our findings on the magnitude of CDS transaction costs and on their correlation with other measures of CDS liquidity. Section VI compares our CDS transaction costs to those of their underlying bonds. Finally, Section VII concludes.

II. CDS Market Background

Single-name CDSs are privately negotiated financial contracts that allow for the transfer of credit risk of a reference obligation from one party to another. If the reference obligation experiences a credit event (e.g. default), the protection seller either takes delivery of the defaulted obligation and pays its par value to the buyer (physical settlement), or pays the buyer the difference between the par value and recovery value of the obligation (cash settlement). Note that while single-name CDS contracts reference both an entity and an obligation, it is an oversimplification to assume that they are priced off the credit risk of the entity alone. Indeed, our sample includes several CDS contracts referencing different bonds issued by the same firm. Furthermore, several obligations are deliverable on a CDS contract in addition to the reference obligation, which is why CDSs are likely priced based on the credit risk of the cheapest deliverable obligation. This makes CDS to bond comparisons difficult as identifying a single underlying security in a CDS contract is not straightforward.

CDS contracts trade in a large and active market, which in recent years has undergone several structural improvements. Most notably, in April 2009 ISDA implemented a number of CDS contract and convention changes, popularly known at the "CDS Big Bang." One of these changes relates to the quoting convention for CDS contracts. Prior to April 2009 contracts were quoted in par spreads. The par spread is the coupon a buyer is willing to pay on a quarterly basis over the life of the contract in exchange for protection against the default of the reference obligation, with no initial cash exchange.[7] CDS par spreads are an intuitively appealing quoting convention in that they are directly related to the yield spreads on the CDS reference obligations. The CDS Big Bang standardized CDS coupon levels at either 100 or 500 bps to create fungibility across contracts, which forced a change in the way market participants quote and trade the contracts. Generally, CDS are now quoted in upfront payments (or "upfronts"), which are the initial cash payments that compensate for the difference between the fixed coupon and the actual par spread. Upfronts are typically expressed in terms of basis points as a percentage of the notional amount.[8] That is, an upfront of $125,000 for $10M notional is quoted as 125 bps.

Despite ISDA's efforts to improve the infrastructure of CDS trading, CDS contracts continue to trade in an over-the-counter market characterized by little pre-trade transparency. Information on CDS quotes is not publicly disseminated, though it can be purchased through third-party providers such as Markit and Bloomberg. More importantly,

[7]Cash may still be exchanged as a collateral on the trade.

[8]CDSs can also be quoted in terms of a conventional, or quoted, spreads, which are obtained using ISDA's standard CDS converter, `http://www.cdsmodel.com/cdsmodel`. Some of the assumptions underlying the conversion include a fixed recovery rate, a market convention for risk-free discount rates, and a constant instantaneous default probability across tenors.

CDS quotes are typically indicative, providing only imprecise information to market participants interested in trading. Nonetheless, indicative quotes have been widely used in the academic literature to proxy for CDS transaction prices and indicative bid-ask quote spreads to proxy for CDS liquidity.[9] Post-trade transparency in the CDS market is also limited. Currently, DTCC collects contract-by-contract pricing, trading and counterparty information, but publicly disseminates only weekly trading statistics.[10] We use the transaction data underlying the DTCC's weekly statistics, rather than indicative quotes, to estimate transaction costs for single-name CDS contracts.

III. Transaction Cost Estimation Methodology

Our estimation approach builds on the concept introduced by Roll (1984) that an observed transaction price can be decomposed into a fundamental value, and a transitory component that is related to the trade itself. That is,

$$P_t = V_t + c\, Q_t, \tag{1}$$

where P_t is the transaction price, V_t is the fundamental value, Q_t is an indicator variable that equals 1 if the trade is initiated by the buyer and -1 if it is initiated by the seller, and c is the effective bid-ask half-spread. This approach has been expanded by Schultz (2001), Bessembinder, Maxwell, and Venkataraman (2006), Harris and Piwowar (2006), Edwards, Harris, and Piwowar (2007), and Goldstein, Hotchkiss, and Sirri (2007) among others to estimate transaction costs in the municipal and corporate bond markets.

[9] See Blanco, Brennan, and Marsh (2005), Forte and Pena (2009), Das, Hanouna, and Sarin (2009), Zhang, Zhou, and Zhu (2009), and Das, Kalimipalli, and Nayak (2014) among others.

[10] DTCC releases weekly warehouse inventory statistics every Tuesday after 5:00 p.m. ET (2200 GMT).

Following the above studies, we assume that the observed transaction price, P_t, is equal to an unobserved fundamental value, V_t, plus a transitory component that depends on whether the liquidity demander is the buyer or the seller of the contract. In the context of our model, P_t is the CDS upfront as a fraction of the notional (in basis points). The midpoint of the quoted indicative bid-ask spread, M_t, serves as an imprecise proxy for the unobserved fundamental value, V_t.

To determine whether a trade is buyer- or seller-initiated, we follow the extant market microstructure literature and sign trades using a derivative algorithm based on Lee and Ready (1991). Specifically, we assume that a trade is buyer-initiated if the transaction price is above the contemporaneous quote midpoint, M_t, and seller-initiated if the price is below it.[11] That is, $Q_t = 1$ if $P_t > M_t$, and $Q_t = -1$ if $P_t < M_t$.[12]

A. Time-Series Estimation Model

We allow the transitory price component c to depend on the notional amount of the contract traded (S_t) and to differ for dealer-to-dealer (DD) and dealer-to-enduser (DE) trades.[13] To be more specific, let $D_t = 1$ for DD trades and $D_t = 0$ for DE trades. If $d(S_t)$ and $c(S_t)$ are the respective transaction cost functions for DD and DE trades that depend on the notional dollar amount of the trade, S_t, then the observed price of a CDS

[11] We are limited in our implementation of the Lee and Ready (1991) algorithm by the lack of time stamps on CDS trade prices. Thus, rather than time-matching trades to quotes, we date-match them. We confirm the robustness of our findings to alternative ways of signing trades in Section V.A.

[12] In our sample, no trades take place at the midpoint of the bid-ask spread.

[13] This is consistent with the bond literature that allows for transaction costs to differ across trade sizes, and for dealer-to-dealer versus dealer-to-customer trades (e.g. Harris and Piwowar (2006) and Edwards, Harris, and Piwowar (2007)). In the CDS market, customers are commonly referred to as endusers, since they are using the contract to hedge or mitigate credit risk. However, in the CDS market, all endusers are institutional traders and retail investors are not present.

can be modelled as

$$P_t = V_t + c(S_t)(1 - D_t)Q_t + d(S_t)D_tQ_t + \epsilon_t \tag{2}$$

where ϵ_t is an asset-specific valuation factor with mean zero and variance σ_ϵ^2. We assume that the fundamental value V_t can only be observed as M_t with an error κ_t as in

$$V_t = M_t + \kappa_t \tag{3}$$

where κ_t has mean zero and variance σ_κ^2. Adopting the approach in Edwards, Harris, and Piwowar (2007), we model the transaction cost functions using the following specifications:

$$d(S_t) = d_0 + d_1\frac{1}{S_t} + d_2 \log S_t + d_3S_t + d_4S_t^2 + \varphi_t \tag{4}$$

and

$$c(S_t) = c_0 + c_1\frac{1}{S_t} + c_2 \log S_t + c_3S_t + c_4S_t^2 + \vartheta_t \tag{5}$$

where the errors represent variations in actual transaction costs unexplained by the average cost functions.[14] They could be random deviations for a given trade from the mean cost function or an error for a given trade size that an average cost function cannot represent. We assume that the errors φ_t and ϑ_t have mean zero and variances σ_φ^2 and $(1 + \alpha)\sigma_\varphi^2$ respectively.

The above assumptions sufficiently identify the variance of the reduced form pricing

[14] We confirm robustness of our findings to alternative functional forms in Section V.A.

model. Combining equations (2), (3), (4) and (5), we obtain

$$P_t - M_t = \Bigg[\left(c_0 + c_1\frac{1}{S_t} + c_2\log S_t + c_3 S_t + c_4 S_t^2\right)(1-D_t) + \left(d_0 + d_1\frac{1}{S_t} + d_2\log S_t + d_3 S_t + d_4 S_t^2\right)D_t\Bigg]Q_t + \varepsilon_t \tag{6}$$

where the bracketed term represents c in equation 1. The combined error term ε_t has mean zero and variance given by

$$\sigma_{\varepsilon t}^2 = \sigma_\epsilon^2 + \sigma_\kappa^2 + \sigma_\varphi^2 + (\alpha\sigma_\varphi^2)(1-D_t) = \sigma_0^2 + \sigma_1^2(1-D_t) \tag{7}$$

We estimate (6) separately for each CDS using feasible generalized least squares (FGLS) with the error variance structure given by (7). There are two points worth noting. First, because our econometric specification requires that we observe at least eleven trades, our transaction cost estimates are potentially biased towards more actively traded contracts. Second, if the decision to trade is endogenous and depends on transaction costs, an average transaction cost estimate might be further downward biased.

B. Cross-Sectional Transaction Costs

Using the estimated DE and DD transaction-cost functions for each CDS in our sample, we evaluate these functions at each trade of each contract to obtain an estimate of the effective half-spread. Let $\hat{c} = [\hat{c}_0 \quad \hat{c}_1 \quad \hat{c}_2 \quad \hat{c}_3]'$ be the vector of estimated coefficients for a given contract, $\hat{\Sigma}_c$ the corresponding covariance matrix of the estimates, and

$$\boldsymbol{S}_t = \begin{bmatrix} 1 & \frac{1}{S_t} & \log S_t & S_t & S_t^2 \end{bmatrix}. \tag{8}$$

Then for a DE trade with a notional amount S_t the estimated effective half-spread is

$$\hat{c}(S_t) = \boldsymbol{S}_t \hat{c} \tag{9}$$

and its estimated error variance is

$$\mathrm{Var}(\hat{c}(S_t)) = \boldsymbol{S}_t \hat{\Sigma}_c \boldsymbol{S}_t'. \tag{10}$$

Similarly, we define and calculate for each DD trade $\hat{d}(S_t)$ and $\mathrm{Var}(\hat{d}(S_t))$. We use the inverse of the estimated error variances as precision measures for $\hat{c}(S_t)$ and $\hat{d}(S_t)$.[15]

For each CDS in our sample we first calculate trade-by-trade transaction costs by evaluating either $\hat{c}(S_t)$ or $\hat{d}(S_t)$ at each trade's actual notional amount, S_t, depending on whether the trade is a dealer-to-enduser or dealer-to-dealer trade. This approach is similar to the one used by Jostova, Nikolova, Philipov, and Stahel (2013) to estimate corporate-bond transaction costs.[16] Then based on S_t and D_t, we assign the transaction cost estimate to one of the notional size bins (≤\$2.5M, >\$2.5-7.5M, >\$7.5-12.5M, >\$12.5-25M, >\$25-50M, >\$50-100M, and >\$100M), and calculate for each bin the precision-weighted cross-sectional average DE and DD transaction costs.

IV. Data and Summary Statistics

Our source of CDS trade data is the Trade Information Warehouse (TIW) of DTCC. DTCC provides clearing, settlement and other services for OTC derivatives. In November 2006, DTCC established its automated TIW as the electronic central registry for

[15]This follows Harris and Piwowar (2006) and Edwards, Harris, and Piwowar (2007).

[16]Harris and Piwowar (2006) and Edwards, Harris, and Piwowar (2007) report estimated transaction costs for municipal and corporate bonds respectively, using hypothetical instead of actual trade sizes. We investigate the robustness of our findings to following their approach in Section V.A.

CDS contracts. Since that time, the vast majority of CDS contracts traded have been registered with the TIW. In addition, all of the major global CDS dealers have registered with the TIW many contracts executed prior to November 2006 to which they were a counterparty. It is important to note that although DTCC provides information on when a trade was reported to the TIW, there are no reporting requirements for CDS trades similar to reporting requirements for corporate bonds. Therefore, reported time stamps do not indicate the time when a trade actually occurred.[17]

From DTCC's TIW we collect information on new confirmed trades in single-name corporate CDSs for the period of August 1, 2009 to May 31, 2014. We remove from the sample non-price-forming or redundant records, which arise either out of multilateral netting (compression), intra-family trades, or duplicate reporting by prime brokers.[18] We also exclude centrally cleared trades because we cannot unambiguously match buys with their corresponding sells in order to eliminate the redundancy introduced by the interjection of a central counterparty between the original buyer and seller.[19] We retain only USD-denominated standard ISDA contracts with no restructuring clause.[20] Since CDS trading is concentrated in the 5-year tenor of contracts with 100 or 500-bps coupon, we limit our sample to only these contracts.[21] We further exclude CDS trades that take place within six months of the reference bond migrating from investment grade (IG) to

[17] In general, FINRA rules require broker/dealers to report to TRACE the execution of secondary-market transactions in TRACE-eligible securities within 15 minutes of the execution time. Thus, a TRACE time stamp can reasonably proxy for trade execution time.

[18] Intra-family trades are transfers of CDS contracts from one to another division within the same firm. Since it is unclear whether the reported price is the current market price, we exclude these trades from our sample. Duplicate reporting by prime brokers occurs when a prime broker is used and the prime broker is different from the executing dealer, since then both entities may submit the trade to DTCC.

[19] Loon and Zhong (2014) show that central clearing improves CDS liquidity and trading activity.

[20] The CDS Big Bang made "No Restructuring" the standard for North American Corporate transactions. However, parties may still elect for Modified Restructuring to apply.

[21] The coupon-based filter excludes 120 CDS contracts with non-standard coupons that comprise 0.02% of trades. Almost all contracts trade with a 5-year tenor, so the tenor-based filter excludes only 160 CDS contracts, which are very thinly traded (511 trades during our sample period).

non-investment grade (NIG) (or vice versa), in order to ensure that our findings are not driven by the increased trading volume and/or quote volatility around the time bonds cross the investment-grade threshold.[22] We also filter out trades likely to be reporting errors by using standard filters based on intra-day deviations from the daily median. We are left with a sample of 376,685 trades on 1,217 CDS contracts.

We obtain intra-day indicative bid and ask CDS quotes from Markit. Markit uses a real-time quote parsing algorithm to extract indicative over-the-counter quotes from dealer messages. Messages are from multiple sources and are parsed and scrubbed in near real-time to produce consistent quote data on CDSs.[23] For the period prior to June 28, 2010, Markit provides indicative quote midpoints based on trader/broker reports and artificially generates bid/ask spreads using a fixed band around the midpoints.[24] After June 28, 2010, Markit provides the actual bid/ask quotes as transmitted by traders/brokers. For each CDS in the Markit database we calculate a daily indicative quote midpoint as the time-weighted average of midpoints provided from 8am to 4pm EST, where the weight is the number of seconds a quote remains outstanding.[25] We match these daily quotes to CDS trades by reference entity, date, seniority, tenor and coupon, which reduces our sample to 332,369 trades on 1,024 CDS contracts.

We have sufficient number of trades to estimate transaction costs for 851 of these CDS contracts, which comprise our final sample. These contracts account for 331,873 trades totaling \$1.9T of notional amount traded. Of these trades, 243,275 (73%) are DD trades

[22]Including these CDS trades in the sample leaves our transaction cost estimates largely unchanged. On average, effective half-spreads are never more than 0.5 bps larger than half-spreads estimated when trades within six months of a major rating change are excluded.

[23]For more details see `http://www.dadd.co.uk/en/products/data/quotes/quotes.page?`.

[24]Since our methodology uses the quoted spread midpoint, and not the quoted spread itself, this limitation of the data should not affect the transaction-cost estimates. When we compare quoted to effective spreads, we do so only for the period after June 28, 2010.

[25]In Section V.A we confirm that out transaction-cost estimates are robust to using alternative quote midpoints to proxy for M_t and sign CDS trades. Robustness is likely the result of quote midpoints being relatively stable from 8am to 4 pm EST.

and 88,598 (27%) are DE trades. Panels A and B of Table I present additional descriptive statistics for our CDS sample and make three noteworthy points. First, interdealer transactions account for most of the daily notional amount traded ($1,200M) and number of trades (202), which is why we are reluctant to ignore them and intentionally choose a methodology that allows us to estimate DD transaction costs. This is in contrast to Harris and Piwowar (2006) and Edwards, Harris, and Piwowar (2007) who focus their analysis on bond transaction costs of customer trades.[26] Second, DD and DE trades do not significantly differ in terms of trade size. The average notional traded is $5.9M for DD trades and $4.8M for DE trades. This is another departure from the corporate bond market where the average trade size of institutional-size trades is more than twice that of interdealer trades.[27] Finally, trades in our sample are close to evenly split into buyer- and seller-initiated, both in terms of number of trades and notional amount trade. Buyer- and seller-initiated trades are of a similar average size.

Data on the industry and credit-rating composition of our sample of 851 single-name CDS contracts come from DTCC and Mergent's Fixed Income Securities Database (FISD) respectively. Panel C of Table I shows that the financial services industry is the most represented accounting for roughly a quarter of CDS trading. The credit-quality breakdown of the sample indicates that more than half of the trades are of CDS written on IG bonds. These represent 65% of the notional amount traded.

Bond trade data for the reference obligations come from the Enhanced Trade Reporting and Compliance Engine (TRACE) of the Financial Industry Regulatory Authority (FINRA). Introduced in July 2002, TRACE provides information on secondary mar-

[26]They report that interdealer trades account for 27% of weekly and 25% of daily dollar volume in the municipal and corporate bond markets respectively.

[27]See Edwards, Harris, and Piwowar (2007). We compare CDS trades to institutional- rather than retail-size corporate-bond trades, because the CDS market is an institutional-investor market and CDSs almost always trade in sizes higher than $100K.

ket transactions in publicly traded TRACE-eligible securities, which comprise 99% of the over-the-counter corporate-bond market activity.[28] The enhanced version of the database distinguishes inter-dealer from dealer-to-customer trades, indicates customer trading direction, and reports untruncated trade sizes – information previously not disseminated to the public. Since Enhanced TRACE is only available with an 18-month lag, our corporate-bond trade sample ends on December 31, 2012. We filter out trade cancellations and corrections using the approach in Dick-Nielsen (2009). We are able to obtain bond trade information and characteristics matching 739 of the 851 CDSs in our sample during the period of common DTCC/Markit and TRACE coverage, August 1, 2009 – December 31, 2012.

V. CDS Transaction Costs

A. CDS Transaction Cost Estimates

We use the methodology described in detail in Section III to estimate transaction costs for the 851 CDS contracts in our sample.[29] Table II presents the cross-sectional precision-weighted average transaction costs for DE and DD trades for different trade-size bins. For each DE or DD trade we measure the precision of the transaction-cost estimate by the inverse of $\mathrm{Var}(\hat{c}(S_t))$ or $\mathrm{Var}(\hat{d}(S_t))$ respectively. Figure I plots these along with the precision-weighted average of the 95% confidence intervals associated with each CDS's cost estimate, and the proportion of DE and DD notional amount traded in each trade-size bin.

[28] See http://www.finra.org/web/groups/industry/@ip/@comp/@mt/documents/appsupportdocs/p014320.pdf

[29] 0.1% of our estimated effective half-spreads are negative. In comparison, Edwards, Harris, and Piwowar (2007) report that, depending on trade size, between 7.9% and 38.4% of the estimated effective half spreads for corporate bonds are negative.

Average effective half-spreads range from 13.8 to 22.9 bps of the notional amount for DE trades and from 10.9 to 23.8 bps for DD trades. For trade sizes in the $2.5M to $25M range, where most notional is traded, DD transaction costs are lower than DE transaction costs. For smaller trade sizes ($\leq$ $2.5M), we find the opposite. For larger trade sizes (>$50M), DD and DE transaction costs are not significantly different. However, fewer trades at the largest trade sizes as well as wide average confidence intervals suggest that the latter estimates are less accurate and make it difficult to draw definitive conclusions. Interestingly, CDS transaction costs do not decrease as significantly with trade size as do corporate-bond transaction costs. Edwards, Harris, and Piwowar (2007) find that half-spreads in the corporate-bond market decline close to tenfold (from 62 to 9 bps) when trade size increases from $20,000 to $1,000,000. In contrast, CDS transaction costs first decrease and then slightly increase, but remain relatively flat in a narrow range of 10 bps.

We assess the robustness of our transaction cost estimates along several dimensions: functional form of the cost function, algorithm for signing of trades, and evaluation of the estimated cost functions. First, we examine whether the functional form of the cost function we use affects the cost estimates' magnitude. We specify, estimate, and evaluate three alternative functions similar to Edwards, Harris, and Piwowar (2007), including

$$d(S_t) = d_0 + d_1 \frac{1}{S_t} + d_2 \log S_t + d_3 S_t + d_4 S_t^2 + d_5 S_t^3 + \varphi_{4t} \tag{11}$$

$$c(S_t) = c_0 + c_1 \tfrac{1}{S_t} + c_2 \log S_t + c_3 S_t + c_4 S_t^2 + c_5 S_t^3 + \vartheta_{4t}$$

and various nested models (M1-M4) obtained by setting different combinations of coefficients in (11) to zero. Figure II presents the four estimated function curves. DE transaction costs are essentially identical regardless of the cost function's form. DD transaction costs are also very similar for all but the largest trade sizes, where trading is sparse and the estimates are less accurate. Overall, the proximity of the curves to each other suggests that average CDS transaction costs are not significantly affected by the choice of functional form. The average adjusted R^2 statistics of the estimated models are 38.5%, 38.8%, 39.1% and 39.2% for M1 through M4 respectively. Since there is no substantial improvement in model fit when moving from M3 to M4, in the remainder of the paper we report results based on the five-parameter model, M3 (equations 4 and 5).

Next, we examine whether the CDS transaction-cost estimates are importantly affected by the algorithm we use for signing trades. As a first alternative, we estimate daily quote midpoints using the following subsets of quotes: (1) only opening quotes (5 minutes before 8am EST), (2) only morning quotes (from 8am to 12pm), (3) only afternoon quotes (12pm to 4pm), or (4) quotes from 7am to 4pm EST. Then we re-categorize trades into buyer- and seller-initiated using daily quote midpoints based on (1)-(4). In unreported results, we find that using these alternative daily quote midpoints to sign trades produces effective half-spreads that are within 2 bps of those estimated using quotes from 8am to 4pm EST. Thus, for the remainder of the paper we report results using the latter.

As a second alternative, we follow the approach in Schultz (2001), Harris and Piwowar (2006) and Edwards, Harris, and Piwowar (2007) and assume that customers always demand liquidity. That is, if an enduser is the buyer/seller, we assume this to be a buyer-/seller-initiated trade. We recognize that this assumption, though appropriate for the corporate-bond market, might not accurately represent the functioning of the CDS market where dealers account for most of the trading activity and customers are large institutional investors (e.g. hedge funds) who likely demand but also supply liquidity. However, misclassifying buyer-initiated trades as seller-initiated will have the effect of underestimating transaction costs. Indeed, when we assume that customers always demand liquidity, 31% of DE trades carry a different sign from the one assigned under our base signing algorithm and not surprisingly, the resultant transaction cost estimates are now half of their original magnitude. Furthermore, now 10.6% of the DE transaction-cost estimates are negative compared to 0.2% when using our base signing methodology. This suggests that misclassification might be an issue if we assume that customers in the CDS market always demand liquidity, which is why in the remainder of the paper we use CDS transaction costs estimated without requiring this assumption.

Finally, the average CDS transaction costs reported thus far are constructed by evaluating the estimated cost functions at actual trade sizes as in Jostova, Nikolova, Philipov, and Stahel (2013).[30] In contrast, Harris and Piwowar (2006) and Edwards, Harris, and Piwowar (2007) use hypothetical trade sizes. We investigate whether using the approach in the latter two papers affects our conclusions. Figure III plots the cross-sectional precision-weighted average DE and DD transaction costs for different hypothetical trade sizes. The figure shows that for trade sizes up to $30M the transaction cost estimates are similar to the ones shown in Figure I, but for trade sizes larger than $30M DE trades are cheaper to execute than DD trades. This suggests that by evaluating the cost functions at trade sizes at which some CDSs do not trade, one may underestimate the trading costs incurred by CDS endusers trading large notional amounts. Thus, in the remainder of the paper we use half-spreads generated by evaluating the cost function of each CDS at the actual notional amounts over which the function is estimated.

[30]For a detailed description of the approach see the Appendix in Jostova, Nikolova, Philipov, and Stahel (2013)

B. CDS Transaction Costs and Quoted Bid-Ask Spreads

Given the limited availability of CDS trade data, existing studies have often used quoted bid-ask spreads as a proxy for CDS transaction costs or more generally CDS liquidity.[31] However, in many dealer markets quoted prices are often just a starting point for negotiations between dealers and customers, and trades often occur at prices other than the quoted bids and asks. Moreover, in the CDS market quotes are typically indicative. As a result, a quoted bid-ask spread may significantly differ from the transaction costs actually incurred by investors and the gross revenue earned by liquidity providers. We next investigate the relationship between the estimated CDS transaction costs and indicative quoted spreads, and present our findings in Table III.

Panel A provides summary statistics on the indicative quoted half-spreads observed in our sample and our estimated effective half-spreads. Average quoted spreads are 36.9 bps and effective spreads are only 13 bps. Since Markit quotes are for trade sizes of \$10M, we separately examine trades of sizes between \$7.5M and \$12.5M, and find that average effective spreads in this subsample are even lower (11.5 bps). Table III, Panel B summarizes our analysis of whether observed trade prices tend to fall within the quoted bid-ask spread. We find that a significant portion of trades, 24% overall and 21% in the \$7.5-12.5M trade-size range, are outside the quoted spread. We recognize that this finding might be partially due to our inability to time-match trades to quotes, so to assess robustness we also compare transaction prices to the intra-day highest ask and lowest bid quote. Even then, 11% of all trades and 9% of trades in the \$7.5-12.5M trade-size range are outside the quoted spread. Finally, to directly analyze to what extent effective spreads can be proxied by indicative quoted spreads in the cross-section, we estimate a weighted least-squares (WLS) regression, using as weights the inverses of

[31] See Acharya and Johnson (2007) among others.

the estimated error variances, $\text{Var}(\hat{c}(S_t))$ for DE trades and $\text{Var}(\hat{d}(S_t))$ for DD trades, and clustering standard errors at the CDS level. The results are presented in Table III, Panel C. We find that the relationship between quoted and effective bid-ask spreads is positive and significant, but also weak as suggested by the low R^2 (0.029). Furthermore, a 10-bps increase in quoted spreads corresponds to only a 0.65-bp increase in effective spreads. In columns (2)-(4) we investigate whether the relationship between quoted and effective spreads varies with the type of trade and credit risk of the reference obligation. We find that the model's fit significantly improves when we control for DE/DD trades and IG/NIG reference obligations. The magnitude of the quoted spread coefficients also increases suggesting a closer link between quoted and effective spreads for DE trades of NIG contracts. This is intuitively appealing, since Markit quotes are likely to be geared towards customers rather than other dealers, and since quotes are likely to be more responsive for CDS on high-credit-risk bonds, which are more information sensitive. In sum, the analysis presented in Table III highlights the limitations of using indicative quoted bid-ask spreads to proxy for effective spreads.

C. CDS Transaction Costs and Market Activity

We next examine the relationship between CDS transaction costs and CDS market activity. On one hand, higher transaction demand should lead to more profits and competition among dealers, and hence cheaper provision of liquidity services (Demsetz (1968)). On the other hand, classical models of asymmetric information suggest that it is not total trading activity that matters, but rather the proportion of informed to uninformed order flow (Kyle (1985)). Since banks are major players in the CDS market and since they possibly enjoy an informational advantage (Acharya and Johnson (2007)), a high level of CDS market activity may simply mean a high level of informed trading.

Consistent with this notion, Tang and Yan (2007) show that adverse selection is an important concern in the pricing of actively traded CDS contracts.

We estimate several WLS regressions of transaction costs on measures of CDS market activity. We construct these measures at the individual CDS level and then average them cross-sectionally. The measures include: daily number of trades, daily number of quotes, monthly quote matching intensity (ratio of trades to quotes as used by Tang and Yan (2007) to measure market search frictions), and daily quote depth (number of quote-submitting dealers as reported by Markit). Table IV presents the results from the WLS regressions, using as weights the inverses of the estimated error variances, $\mathrm{Var}(\hat{c}(S_t))$ or $\mathrm{Var}(\hat{d}(S_t))$ depending on whether the estimated effective spread is for a dealer-to-enduser or dealer-to-dealer trade, and clustering standard errors at the CDS level. We find that in the cross-section transaction costs are positively correlated with the daily number of CDS trades. In unreported results, we confirm robustness to measuring trading activity with the daily notional amount traded or proportion of in-sample days without a trade. CDS transaction costs are generally not related to daily number of quotes or quote depth, though for certain trade sizes quote depth carries a marginally significant positive coefficient. Finally, quote matching intensity is also positively related to transaction costs for all but the largest trade sizes. In sum, we find that measures of CDS market activity tend to be positively related to CDS transaction costs. This is consistent with the notion that ease of trading in the CDS market may facilitate informed trading, which increases the risk to dealers of adverse selection and widens effective bid-ask spreads.

VI. CDS and Corporate Bond Transaction Costs

Investors can trade corporate credit in either the cash or the derivatives market. Duffie (1999) establishes a theoretical relationship between the price of a bond and that of a CDS written on it by showing that the spread on a risky floating-rate note over a risk-free floating-rate note exactly equals the CDS premium.[32] This implies that when comparing the costs of trading corporate credit in the bond versus CDS market, one should also consider Treasury-bond transaction costs. Fleming, Mizrach, and Nguyen (2014), using data from the BrokerTec ECN platform, show that transaction costs in the Treasury market are very low with an average bid-ask spread for a 5-year note of 0.9 bps. Since we have no access to data that would allow us to estimate Treasury-bond transaction costs for our sample period and condition them on trade size, and since the magnitude of these transaction costs appears insignificantly small relative to the CDS effective spreads we estimate, we neglect them when we compare corporate-bond and CDS transaction costs below.

We showed in Section V.A that DE trades in the CDS market incur average roundtrip transaction costs in the range of 28-46 bps (half-spreads in the 14-23 bps range). When compared to existing transaction-cost estimates in the corporate-bond literature, our results suggest that large customer trades are not necessarily cheaper to execute in the CDS market. For instance, using matched buys and sells for the same bond on the same day, a methodology favoring more liquid bonds, Chakravarty and Sarkar (2003) estimate roundtrip corporate-bond transaction costs on institutional trades of 21 bps. Using a similar approach, Hong and Warga (2000) estimate 13 bps (19 bps) for institutional

[32]Though in theory the relationship is between a CDS on one hand and a portfolio of a risky and risk-free floating-rate notes on the other, empirically it tends to hold even when floating-rate notes are substituted with fixed-rate bonds though some have documented deviations from parity during the recent financial crisis. See Bai and Collin-Dufresne (2013), Blanco, Brennan, and Marsh (2005) and Nashikkar, Subrahmanyam, and Mahanti (2011) among others.

trades in investment-grade (non-investment-grade) bonds, and Goldstein, Hotchkiss, and Sirri (2007) analyzing only BBB-rated bonds traded by the same dealer estimate round-trip transaction costs of 27 bps for trade sizes larger than $1,000,000.[33] Using regression methodologies as opposed to matching buys and sells, Schultz (2001) documents transaction costs of 27 bps for institutional trades in investment-grade bonds, which is similar to the estimates in Bessembinder, Maxwell, and Venkataraman (2006). Finally, for institutional-sized trades Edwards, Harris, and Piwowar (2007) estimate roundtrip transaction costs of 28 bps for a trade size of $500K and 18 bps for a trade size of $1,000,000.

It is possible that our estimates of CDS transaction costs are higher than the corporate-bond transaction cost estimated in the literature, because of post-trade transparency differences across the two markets. While TRACE makes security-level trade information publicly available with a short time lag, no comparable contract-by-contract trade reporting and disseminating mechanism exists in the CDS market. Bessembinder, Maxwell, and Venkataraman (2006), Goldstein, Hotchkiss, and Sirri (2007) and Edwards, Harris, and Piwowar (2007) document that increased post-trade transparency in the corporate-bond market, brought about by the introduction of TRACE, lowered transaction costs. It is thus not surprising that the average single-name CDS may be more expensive to trade than the average corporate bond. On the other hand, the decision whether to trade credit in the CDS or the cash market may be endogenous and may depend on transaction costs. Thus, the CDS contracts for which we observe trading may be cheaper to trade than their underlying bonds, yet still more expensive to trade than the average bond. This is consistent with the evidence in Das, Kalimipalli, and Nayak (2014) that bonds of issuers with CDSs are somewhat less liquid than other bonds.

[33] These roundtrip spread estimates are for bonds with disseminated price information, which should be the relevant estimates since now most bonds' prices are disseminated via TRACE.

In this Section we compare the transaction costs of the CDS contracts in our sample to the transaction costs of their underlying bonds.[34] More specifically, we construct a sample of matched CDS-bond pairs, for which we are able to calculate transaction costs during the period of joint DTCC/Markit and Enhanced TRACE coverage, August 1, 2009–December 31, 2012. Corporate-bond transaction-cost functions are estimated as in Edwards, Harris, and Piwowar (2007) and evaluated as in Jostova, Nikolova, Philipov, and Stahel (2013). Since the methodology in these studies does not allow for the estimation of interdealer transaction costs, we limit our CDS sample to only DE trades.

Matching CDSs to their underlying bonds effectively controls for firm-specific and bond-specific characteristics that may impact both CDS and bond transaction costs, but matching can be difficult for at least two reasons. First, the identification of the underlying bond relevant for a comparison of cross-sectional transaction costs is not unambiguous, since a number of bonds could be delivered on any given CDS. Thus, we match each CDS in our sample in three different ways: (1) to its reference obligation ("reference bond"), (2) to all traded bonds of its reference entity ("issuer bond"), or (3) to the cheapest-to-trade bond of its reference entity ("cheapest bond"). To identify a reference entity's traded bonds, we use Mergent's FISD to find all 6-digit CUSIPs associated with the reference entity. We then use Enhanced TRACE to calculate transaction costs for all bonds matching any of the reference entity's 6-digit CUSIPs. Of these bonds, we flag the cheapest one to trade as the bond with the lowest average transaction cost within each actual trade-size bin.

Our second challenge in comparing CDS to bond transaction costs arises from the

[34]Note that this is a comparison of the relative execution costs for a roundtrip trade. That is, we do not compare the cash outflow needed *to establish a position*, but instead compare the proportional cost *to enter and exit a position* in the CDS versus bond market. Clearly, it is cheaper to establish a position in a derivative than the cash market because the former allows an investor to leverage their initial investment.

fact that CDSs generally trade in sizes larger than those for a typical corporate bond. We find that a significant overlap among CDSs, reference bonds, cheapest bonds, and issuer bonds occurs for trade sizes between \$100K and \$5M. The majority of corporate bonds trade in sizes lower than \$100K (e.g. 77% of our reference bonds), while few CDSs do (2% in our sample). In contrast, many CDSs trade in sizes of \$5M or larger (46% in our sample), while almost no bonds do (e.g. 2% of our reference bonds). Hence, we limit our comparison of CDS and bond transaction costs to the \$100K to \$5M trade-size range.

For all CDS and the corresponding reference/issuer/cheapest bond matches in our sample we calculate precision-weighted cross-sectional average transaction costs by actual trade-size bins. Figure IV plots these transaction-cost averages along with the proportion of trades in each bin. Table V presents t-tests of the hypothesis that precision-weighted average transaction costs are different for the CDSs compared to the reference/cheapest bonds.The Figure and Table show that CDSs are typically cheaper to trade than the corporate bonds they reference. For trade sizes in the \$100-500K range, reference bonds incur transaction costs that are more than three times those of the corresponding single-name CDSs. However, this difference in the cost of trading declines as trade size increases, and at trade sizes larger than \$2M it becomes statistically insignificant. Nonetheless, only 6% of the reference bonds trade at sizes that large. The results are similar when we compare CDS transaction costs to those of the reference entity's cheapest to trade bond. If one considers all bonds of an issuer as close substitutes, trading the reference entity's credit is no longer more expensive in the cash than the derivatives market at sizes larger than \$750K. Moreover, 19% of the cheapest-bond trades fall in trade-size bins where bond half-spreads are lower than CDS half-spreads.

Overall, the analysis in this section generally supports the assumption in the literature that trading corporate credit is cheaper in the derivatives than in the cash market. However, it also reveals that the relative costs of trading CDSs versus bonds importantly depend on trade size. Although at smaller trade sizes CDSs are cheaper, at larger trade sizes they no longer are. This finding is consistent with the notion that when asymmetric information concerns are high, a more likely scenario for large trades, market transparency tends to reduce the adverse-selection component of bid-ask spreads.

VII. Conclusions

Our study contributes to the liquidity literature in three important ways. First, we are the first to estimate single-name CDS transaction costs using newly available trade-by-trade data from DTCC. The TIW of DTCC is the only central trade registry for CDSs and it covers 90% of the total notional amount traded. Hence, our sample reflects the vast majority of CDS trading activity. We estimate individual transaction costs for 851 single-name CDS contracts traded between August 1, 2009 and May 31, 2014. Effective half-spreads are 14 bps of the notional amount for dealer-to-dealer trades and 12 bps for dealer-to-enduser trades for the most common trade sizes, between \$2.5M and \$7.5M.

Second, we show that indicative quoted spreads are twice as large as effective spreads, and explain little of their cross-sectional variation. Moreover, a large proportion of trades take place outside the quoted spreads. Taken together, these findings underscore the limitations of indicative quoted spreads as a proxy for actual transaction costs in the CDS market. We also examine the relationship between CDS transaction costs and market activity, and show that effective spreads are larger for more actively traded contracts. This supports the notion that the CDS market is characterized by a high degree of informed trading, which causes dealers to protect themselves from adverse selection by widening bid-ask spreads.

Finally, we compare single-name CDS transaction costs to the transaction costs of the bonds they reference. We find that trading corporate credit is typically cheaper in the CDS market, but the difference between CDS and bond effective spreads importantly depends on trade size. At smaller trade sizes bonds are more than three times more expensive to trade than the CDS contracts written on them, but at larger trade sizes bonds are no costlier and eventually cheaper. Since most bond trades are in sizes over which they are more expensive than CDSs, our findings provide one explanation for why price discovery and trading corporate credit has shifted from the cash to the derivatives market as documented in Blanco, Brennan, and Marsh (2005), Norden and Wagner (2008), Norden and Weber (2009), Forte and Pena (2009), and Das, Kalimipalli, and Nayak (2014). However, since CDS markets are much less transparent, the findings raise the question of why trading in a less transparent market might be cheaper. We believe that this question will be a fruitful venue for future research.

References

Acharya, Viral V., and Timothy C. Johnson, 2007, Insider trading in credit derivatives, *Journal of Financial Economics* 84, 110–141.

Ashcraft, Adam B., and Joao A.C. Santos, 2009, Has the CDS market lowered the cost of corporate debt?, *Journal of Monetary Economics* 56, 514 – 523.

Bai, Jennie, and Pierre Collin-Dufresne, 2013, The CDS-bond basis, Working paper, Georgetown University.

Bessembinder, Hendrik, William Maxwell, and Kumar Venkataraman, 2006, Market transparency, liquidity externalities, and institutional trading costs in corporate bonds, *Journal of Financial Economics* 82, 251–288.

Bessembinder, Hendrik, and Kumar Venkataraman, 2010, Bid-ask spreads: measuring trade execution costs in financial markets, in Rama Cont, eds.: *Encyclopedia of quantitative finance* (John Wiley & Sons, Ltd.,).

Blanco, Roberto, Simon Brennan, and Ian Marsh, 2005, An empirical analysis of the dynamic relation between investment-grade bonds and credit default swaps, *Journal of Finance* 60, 2255–2281.

Chakravarty, Sugato, and Asani Sarkar, 2003, Trading costs in three U.S. bond markets, *Journal of Fixed Income* 13, 39–48.

Das, Sanjiv, Paul Hanouna, and Atulya Sarin, 2009, Accounting-based versus market-based cross-sectional models of CDS spreads, *Journal of Banking and Finance* 33, 719–730.

Das, Sanjiv, Madhu Kalimipalli, and Subhankar Nayak, 2014, Did CDS trading improve the market for corporate bonds?, *Journal of Financial Economics* 111, 495–525.

Demsetz, Harold, 1968, The cost of transacting, *Quarterly Journal of Economics* 82, 33–53.

Dick-Nielsen, Jens, 2009, Liquidity biases in TRACE, *The Journal of Fixed Income* 19, 43–55.

Duffie, Darrell, 1999, Credit swap valuation, *Financial Analysts Journal* 55, 73–87.

Edwards, Amy, Lawrence E. Harris, and Michael S. Piwowar, 2007, Corporate bond market transaction costs and transparency, *Journal of Finance* 62, 1421–1451.

Fleming, Michael J., Bruce Mizrach, and Giang Nguyen, 2014, The microstructure of a U.S. Treasury ECN: the BrokerTec platform, FRBNY Staff Report No. 381.

Forte, Santiago, and Juan Ignacio Pena, 2009, Credit spreads: An empirical analysis on the informational content of stocks, bonds, and CDS, *Journal of Banking and Finance* 33, 2013–2025.

Fulop, Andras, and Laurence Lescourret, 2009, Intra-daily variations in volatility and transaction costs in the credit default swap market, Working paper available at SSRN: http://ssrn.com/abstract=1509323.

Goldstein, Michael A., Edith S. Hotchkiss, and Erik Sirri, 2007, Transparency and liquidity: a controlled experiment on corporate bonds, *Review of Financial Studies* 20, 235–273.

Harris, Lawrence E., and Michael S. Piwowar, 2006, Secondary trading costs in the municipal bond market, *Journal of Finance* 61, 1361–1397.

Hong, Gwangheon, and Arthur Warga, 2000, An empirical study of bond market transactions, *Financial Analysts Journal* 56, 32–46.

Jostova, Gergana, Stanislava Nikolova, Alexander Philipov, and Christof W. Stahel, 2013, Momentum in corporate bond returns, *Review of Financial Studies* 26, 1649–1693.

Kyle, Albert S., 1985, Continuous auctions and insider trading, *Econometrica* 53, 1315–1336.

Lee, Charles M. C., and Mark J. Ready, 1991, Inferring trade direction from intraday data, *Journal of Finance* 46, 733–746.

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

Loon, Yee Cheng, and Zhaodong Ken Zhong, 2014, The impact of central clearing on counterparty risk, liquidity, and trading: Evidence from the credit default swap market, *Journal of Financial Economics* 112, 91 – 115.

Nashikkar, Amrut, Marti G. Subrahmanyam, and Sriketan Mahanti, 2011, Liquidity and arbitrage in the market for credit risk, *Journal of Financial and Quantitative Analysis* 46, 627–656.

Norden, Lars, and Wolf Wagner, 2008, Credit derivatives and loan pricing, *Journal of Banking and Finance* 32, 2560 – 2569.

Norden, Lars, and Martin Weber, 2009, The co-movement of credit default swap, bond and stock markets: an empirical analysis, *European Financial Management* 15, 529–562.

Roll, Richard, 1984, A simple implicit measure of the effective bid-ask spread in an efficient market, *Journal of Finance* 39, 1127–1139.

Schultz, Paul, 2001, Corporate bond trading costs: a peak behind the curtain., *Journal of Finance* 56, 677–698.

Tang, Dragon, and Hong Yan, 2007, Liquidity and credit default swap spreads, Working paper.

Zhang, Benjamin Yibin, Hao Zhou, and Haibin Zhu, 2009, Explaining credit default swap spreads with the equity volatility and jump risks of individual firms, *Review of Financial Studies* 22, 5099–5131.

Table I: CDS Sample Characteristics

The table presents summary statistics for the sample of 331,873 trades for 851 single-name CDS contracts traded during 8/1/2009-5/31/2014. DD indicates interdealer trades and DE indicates dealer-to-enduser trades. The sample includes 243,275 DD and 88,598 DE trades. Buyer-/seller-initiated trades are those where the transaction price is above/below the same-day quote midpoint. Industry classification is based on DTCC indicators. Credit rating is the bond's S&P rating if available, supplemented with Moody's or Fitch otherwise.

Panel A. Dealer-to-Dealer (DD) Versus Dealer-to-Enduser (DE) CDS Trades

Characteristics	Mean	1%	Median	99%
Number of trades per day	276	12	264	581
DD	202	8	185	485
DE	74	2	71	198
Daily mean dollar volume (in $M)	1,556	69	1,454	3,540
DD	1,200	52	1,062	3,096
DE	358	13	330	1,047
Notional amount by trade (in $M)	5.6	0.1	5.0	26.0
DD	5.9	0.3	5.0	27.0
DE	4.8	0.1	3.4	25.0
Upfront (in bps)	337	-1,532	120	4,900
DD	353	-1,501	111	5,040
DE	291	-1,601	147	4,350

Panel B. Seller-Initiated Versus Buyer-Initiated CDS Trades

	Number of Trades		Total Notional Amount Traded		Average Trade Size
Types	N	%	$B	%	$M
Seller-initiated trades	169,270	51.0	928.5	49.6	5.49
DD	130,430	39.3	750.2	40.1	
DE	38,840	11.7	178.3	9.5	
Buyer-initiated trades	162,603	49.0	942.2	50.4	5.79
DD	112,845	34.0	691.8	37.0	
DE	49,758	15.0	250.4	13.4	

Panel C. CDS Trades by Reference Obligation Characteristics

	Number of Trades		Total Value Traded	
Characteristics	N	%	$B	%
Industry				
Basic materials	19,864	6	104.34	6
Consumer goods	56,702	17	287.36	15
Consumer services	72,454	22	378.56	20
Energy	21,281	6	109.97	6
Financial services	77,128	23	546.59	29
Healthcare	13,054	4	72.57	4
Industrial	22,919	7	128.31	7
Technology	18,929	6	92.16	5
Telecommunications	13,336	4	66.53	4
Utilities	11,760	4	59.13	3
Unknown	4,446	1	25.21	1
Credit rating				
Investment grade (IG)	177,996	54%	1209.7	65%
Non-investment grade (NIG)	143,942	43%	604.2	32%
Unknown	9,935	3%	56.8	3%

Table II: Estimated CDS Transaction Costs

The table presents precision-weighted cross-sectional average CDS transaction costs for various actual trade sizes. The sample includes 331,873 trades for 851 single-name CDS contracts traded during 8/1/2009-5/31/2014. Transaction costs are the effective half-spreads implied by the estimated coefficients of the observed-price model (6):

$$P_t - M_t = \left[\left(c_0 + c_1 \tfrac{1}{S_t} + c_2 \log S_t + c_3 S_t + c_4 S_t^2 \right)(1 - D_t) + \left(d_0 + d_1 \tfrac{1}{S_t} + d_2 \log S_t + d_3 S_t + d_4 S_t^2 \right) D_t \right] Q_t + \varepsilon_t$$

P_t is the CDS upfront as a fraction of the notional (in bps). M_t is the daily midpoint of the quoted indicative bid-ask spread. S_t is the CDS notional amount traded. D_t is an indicator variable that equals 1 for DD trades and 0 for DE trades. Q_t is an indicator variable that equals 1 for buyer-initiated trades (i.e. $P_t > M_t$), and -1 for seller-initiated trades (i.e. $P_t < M_t$). We estimate the model separately for each CDS using time-series regressions. We evaluate the estimated transaction-cost functions at actual trade sizes, and then assign the resultant DD/DE cost estimates to different bins based on the notional amount traded. The weights used to compute the cross-sectional averages are the inverses of the estimated variances of the respective cost estimates. DD indicates interdealer trades and DE indicates dealer-to-enduser trades. ***, **, and * indicate statistical significance at the 1%, 5% or 10% for a t-test that the precision-weighted cross-sectional DE and DD averages are equal.

Trade Size Bin ($M)	DE Transaction Costs (bps)	DE Total Notional Traded ($B)	DD Transaction Costs (bps)	DD Total Notional Traded ($B)	DE – DD Transaction Costs (bps)
0-2.5	14.9	41	23.8	92	-8.9 ***
2.5-7.5	14.5	159	12.3	579	2.2 ***
7.5-12.5	13.8	113	10.9	390	2.8 ***
12.5-25	13.8	84	11.8	272	2.0 ***
25-50	14.7	23	14.5	82	0.2 ***
50-100	17.7	7	17.1	22	0.7
>100	22.9	1	20.3	4	2.5

Table III: CDS Transaction Costs and Quoted Bid-Ask Spreads

Panel A presents upfront half-spread summary statistics and Panel B presents trade price statistics for the sample of 851 single-name CDS traded during 8/1/2009-5/31/2014. Panel C presents the results from a weighted least-squares regression on the sample of 742 single-name CDS traded during 6/28/2010-5/31/2014. The weights are the inverses of the estimated estimator variances of the respective cost estimates. The dependent variable is CDS transaction costs, which is the effective half-spreads estimated as described in Section III, for trade sizes in the 7.5-12.5M range. Quoted spread is the half spread between the daily time-weighted average bid and ask quotes. DD is an indicator variable equal to 1 for interdealer trades, 0 otherwise. IG is an indicator variable equal to 1 for investment-grade reference obligations, 0 otherwise. Standard errors are clustered at the CDS level. T-statistics are reported in parentheses. F-test(1) is the F-statistic for the hypothesis that the coefficients on Quoted spread and Quoted spread*DD sum to 0. F-test(2) is the F-statistic for the hypothesis that the coefficients on Quoted spread and Quoted spread*IG sum to 0. Statistical significance is denoted by ***, **, and * at the 1%, 5%, and 10% level respectively.

Panel A. Upfront Half Spreads

Upfront Half Spreads	Mean	1%	Median	99%	St Dev
Quoted	36.9	8.2	37.5	98.7	20.0
Effective (all trade sizes)	13.0	5.3	10.8	48.9	67.4
Effective (7.5-12.5M trade sizes)	11.5	5.2	10.4	30.1	58.7

Panel B. Trade Price Relative to Quoted Bid-Ask Spread

	Average Daily Quoted Spread				Maximum Daily Quoted Spread			
Trade Price	Count	% Trades	Average Trade Size	Average Transaction Costs	Count	% Trades	Average Trade Size	Average Transaction Costs
All trades								
Inside quoted spread	251,605	76%	5.68	12.52	294,424	89%	5.67	12.81
Below quoted spread	44,616	13%	5.13	15.71	23,315	7%	4.85	16.90
Above quoted spread	35,652	11%	5.94	14.94	14,134	4%	6.29	15.12
Trades in 7.5-12.5M size range								
Inside quoted spread	40,079	79%	9.90	11.22	46,167	91%	9.89	11.42
Below quoted spread	5,596	11%	9.83	12.80	2,688	5%	9.77	13.12
Above quoted spread	5,234	10%	9.85	13.05	2,054	4%	9.81	13.40

Panel C. Weighted Least Squares Regressions

	(1)	(2)	(3)	(4)
Quoted spread	0.065***	0.116***	0.403***	0.437***
	(2.63)	(4.50)	(3.32)	(3.57)
DD		-1.985***		-2.249***
		(-2.64)		(-3.07)
Quoted spread*DD		-0.050		-0.031
		(-1.59)		(-1.06)
IG			-4.911*	-4.739*
			(-1.96)	(-1.87)
Quoted spread*IG			-0.382***	-0.382***
			(-3.12)	(-3.04)
Intercept	10.289***	11.558***	15.711***	17.039***
	(14.82)	(17.18)	(6.49)	(6.79)
Observations	42,149	42,149	42,149	42,149
Adjusted R^2	0.029	0.082	0.308	0.357
F-test(1)		6.1***		10.6***
F-test(2)			1.2	4.8**

Table IV: CDS Transaction Costs and Liquidity

These are the results from a weighted least-squares regression on the sample of 851 single-name CDS traded during 8/1/2009-5/31/2014. The weights are the inverses of the estimated estimator variances of the respective cost estimates. CDS transaction cost estimates are the effective half spreads obtained from time-series regressions estimated separately for each CDS contract in the sample. For each CDS we evaluate the estimated model at actual trade sizes, and then assign the resultant DD/DE transaction cost estimates to different bins based on the notional amount traded. DD is an indicator variable equal to 1 for interdealer trades, 0 otherwise. Number of trades is the daily average number of trades over the sample period. Number of quotes is each CDS' daily average number of quotes over the sample period. Quote depth is Markit's "Composite five-year depth", which is largely based on the average number of quote submitting dealers for the reference entity. T2Q is quote matching intensity intended to measure search frictions and constructed as the monthly ratio of number of trades to number of quotes. Standard errors are clustered at the CDS level. T-statistics are reported in parentheses. Statistical significance is denoted by ***, **, and * at the 1%, 5%, and 10% level respectively.

	Actual Trade Size ($M)						
Variable	0-2.5	2.5-7.5	7.5-12.5	12.5-25	25-50	50-100	>100
DD	8.936***	-2.168***	-2.857***	-2.227***	-0.838	0.813	-1.852
	(7.36)	(-4.11)	(-5.41)	(-2.93)	(-0.48)	(0.34)	(-0.32)
Intercept	14.848***	14.492***	13.814***	13.627***	14.466***	15.760***	21.705***
	(18.26)	(26.69)	(29.84)	(34.24)	(17.95)	(9.56)	(3.39)
Observations	100,465	155,603	48,678	20,916	5,383	754	74
Adjusted R^2	0.090	0.012	0.046	0.031	0.003	-0.000	-0.010
DD	10.143***	-2.095***	-2.666***	-2.116***	-1.267	0.421	-6.767
	(9.11)	(-4.43)	(-5.95)	(-3.06)	(-0.88)	(0.19)	(-1.34)
Number of trades	4.002**	1.952**	1.735***	1.697***	1.761**	1.455	-5.082***
	(2.20)	(2.56)	(3.07)	(3.14)	(2.54)	(1.37)	(-2.73)
Intercept	9.838***	11.768***	10.907***	10.669***	11.292***	13.241***	31.123***
	(6.38)	(13.90)	(13.34)	(11.17)	(7.25)	(6.25)	(5.39)
Observations	100,465	155,603	48,678	20,916	5,383	754	74
Adjusted R^2	0.151	0.077	0.159	0.140	0.086	0.030	0.098
DD	8.839***	-2.214***	-2.819***	-2.149***	-0.811	0.320	-1.410
	(7.25)	(-4.10)	(-5.34)	(-2.70)	(-0.46)	(0.16)	(-0.27)
Number of quotes	-3.967	-5.012	-1.894	-4.725	-5.443	-16.982**	21.101
	(-0.63)	(-1.10)	(-0.37)	(-0.87)	(-0.80)	(-2.21)	(1.10)
Intercept	18.352***	18.913***	15.444***	17.632***	18.956***	30.256***	3.932
	(3.16)	(4.62)	(3.50)	(3.94)	(3.57)	(4.86)	(0.21)
Observations	100,465	155,603	48,678	20,916	5,383	754	74
Adjusted R^2	0.091	0.018	0.048	0.039	0.010	0.032	0.002
DD	9.053***	-2.137***	-2.825***	-2.218***	-0.864	0.966	-1.583
	(7.82)	(-4.04)	(-5.39)	(-2.89)	(-0.50)	(0.41)	(-0.27)
Quote depth	0.215	0.163	0.233*	0.242**	0.358**	0.306	0.165
	(0.67)	(1.25)	(1.88)	(2.17)	(2.04)	(0.97)	(0.23)
Intercept	13.385***	13.364***	12.180***	11.982***	12.089***	13.659***	20.485***
	(5.96)	(12.39)	(12.98)	(14.65)	(8.63)	(6.31)	(2.81)
Observations	99,933	154,996	48,507	20,857	5,372	752	74
Adjusted R^2	0.091	0.014	0.052	0.037	0.013	0.002	-0.024

DD	9.236***	-2.049***	-2.598***	-1.943**	-0.789	-0.932	-1.638
	(7.61)	(-4.03)	(-5.11)	(-2.12)	(-0.52)	(-0.64)	(-0.38)
T2Q	24.952***	18.852***	18.494***	21.046***	48.533***	44.614***	-51.733**
	(2.99)	(3.10)	(3.23)	(2.97)	(4.03)	(2.77)	(-2.75)
Intercept	12.999***	13.196***	12.362***	12.297***	11.777***	14.924***	25.384***
	(16.70)	(23.24)	(24.01)	(16.14)	(9.54)	(10.23)	(6.76)
Observations	100,465	157,326	50,909	19,718	3,005	412	38
Adjusted R-squared	0.103	0.038	0.086	0.075	0.101	0.055	0.040

Table V: Estimated CDS and Reference Bond Transaction Costs

These are weighted cross-sectional averages of DE CDS and bond transaction costs for various actual trade sizes. The sample includes 739 single-name CDS and either (1) the corporate bonds they reference or (2) the cheapest bond of the entity they reference. CDS and bonds are traded during the period August 1, 2009 – December 31, 2012. The transaction cost estimates, which are effective half-spreads, are obtained from time-series regressions estimated separately for each CDS/bond in the sample. For each CDS/bond we evaluate the estimated model at actual trade sizes, and then assign the resultant transaction cost estimates to different bins based on the notional amount traded. The weights used to compute the cross-sectional averages are the inverses of the estimated estimator variances of the respective cost estimates. T-tests of the hypothesis that average CDS and bond transaction costs are equal use Satterthwaite approximation of the standard errors (i.e. do not assume equal CDS and bond transaction-cost variances). Significance of the resultant t-statistics is denoted by ***, **, and * at the 1%, 5%, and 10% level respectively.

Trade Size Bin	CDS DE Transaction Costs (bps)	Ref Bond Transaction Costs (bps)	CDS-Ref Bond Transaction Costs (bps)	Cheapest Bond Transaction Costs (bps)	CDS-Cheapest Bond Transaction Costs (bps)
100K to <500K	14.5	47.1	-32.6***	22.4	-7.9***
500K to <750K	14.8	32.0	-17.3***	15.1	-0.4***
750K to <1.5M	15.0	23.8	-8.7***	14.1	0.9
1.5M to <2M	15.3	17.9	-2.6	8.7	6.6
2M to <4M	15.3	13.8	1.5	8.2	7.2*
4M to <5M	14.2	10.1	4.1**	3.7	10.5***

Figure I: Estimated CDS Transaction Costs by Actual Trade Sizes

These are weighted cross-sectional averages of CDS transaction costs and proportion of notional amount traded for various actual trade-size bins. The sample includes 331,873 trades of 851 single-name CDS contracts traded from August 1, 2009 to May 31, 2014. Transaction costs are implied by the estimated coefficients of the transaction-cost model (equation 6):

$$P_t - M_t = \left[\left(c_0 + c_1 \tfrac{1}{S_t} + c_2 \log S_t + c_3 S_t + c_4 S_t^2 \right) (1 - D_t) + \left(d_0 + d_1 \tfrac{1}{S_t} + d_2 \log S_t + d_3 S_t + d_4 S_t^2 \right) D_t \right] Q_t + \varepsilon_t$$

The cost estimates, which are effective half-spreads, are obtained from time-series regressions estimated separately for each CDS contract in the sample. For each CDS we evaluate the estimated model at actual trade sizes, and then assign the resultant DD/DE transaction cost estimates to different bins based on the notional amount traded. The weights used to compute the cross-sectional averages are the inverses of the estimated estimator variances of the respective cost estimates. The dashed lines on either side of the transaction cost functions represent the weighted means of the 95% confidence interval for the individual CDS cost estimates.

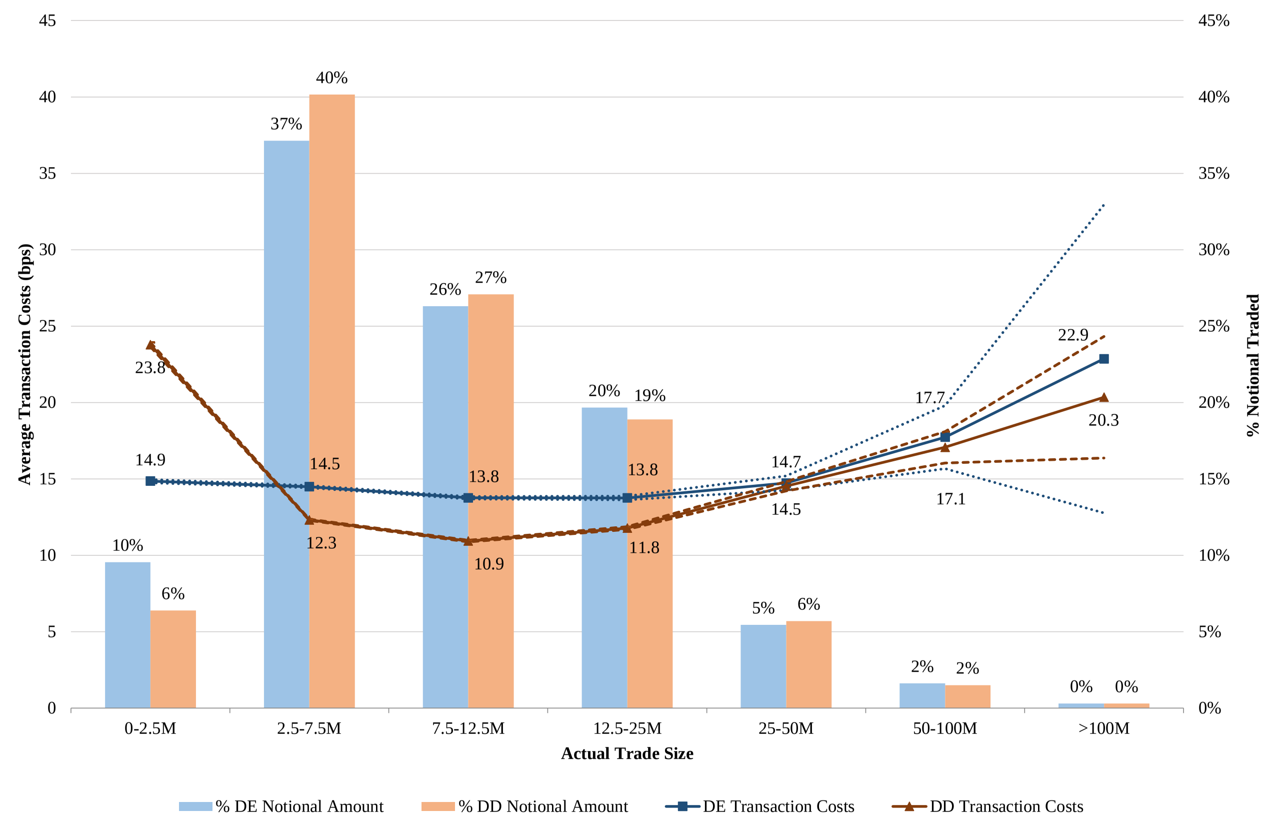

Figure II: Estimated CDS Transaction Costs Using Different Cost Functions

These are weighted cross-sectional averages of CDS transaction costs and proportion of notional amount traded for various actual trade-size bins using different cost functions. The functions are variants of:

$$d(S_t) = d_0 + d_1 \frac{1}{S_t} + d_2 \log S_t + d_3 S_t + d_4 S_t^2 + d_5 S_t^3 + \varphi_{4t}$$
$$c(S_t) = c_0 + c_1 \frac{1}{S_t} + c_2 \log S_t + c_3 S_t + c_4 S_t^2 + c_5 S_t^3 + \vartheta_{4t}$$

obtained by setting different combination of coefficients equal to zero. Specifically,

	M1	M2	M3	M4
d_1/c_1				
d_2/c_2				
d_3/c_3	0			
d_4/c_4	0	0		
d_5/c_5	0	0	0	

The cost estimates, which are effective half-spreads, are obtained from time-series regressions estimated separately for each CDS contract in the sample. For each CDS we evaluate the estimated model at actual trade sizes, and then assign the resultant DD/DE transaction cost estimates to different bins based on the notional amount traded. The weights used to compute the cross-sectional averages are the inverses of the estimated estimator variances of the respective cost estimates.

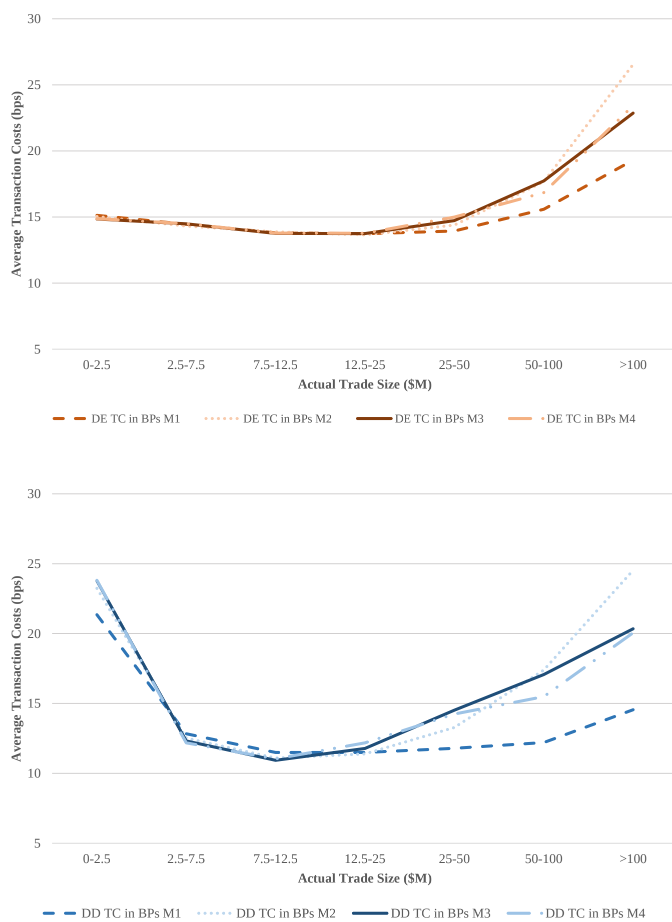

Figure III: Estimated CDS Transaction Costs by Hypothetical Trade Sizes

These are weighted cross-sectional averages of CDS transaction costs for various hypothetical trade sizes. The sample includes 331,873 trades of 851 single-name CDS contracts traded from August 1, 2009 to May 31, 2014. Transaction costs are implied by the estimated coefficients of the transaction-cost model (6):

$$P_t - M_t = \left[\left(c_0 + c_1 \tfrac{1}{S_t} + c_2 \log S_t + c_3 S_t + c_4 S_t^2 \right) (1 - D_t) + \right.$$
$$\left. \left(d_0 + d_1 \tfrac{1}{S_t} + d_2 \log S_t + d_3 S_t + d_4 S_t^2 \right) D_t \right] Q_t + \varepsilon_t$$

The cost estimates, which are effective half-spreads, are obtained from time-series regressions estimated separately for each CDS contract in the sample. For each CDS we evaluate the estimated transaction-cost functions at hypothetical trade sizes. The weights used to compute the cross-sectional averages are the inverses of the estimated estimator variances of the respective cost estimates. The dashed lines on either side of the transaction cost functions represent the weighted means of the 95% confidence interval for the individual CDS cost estimates.

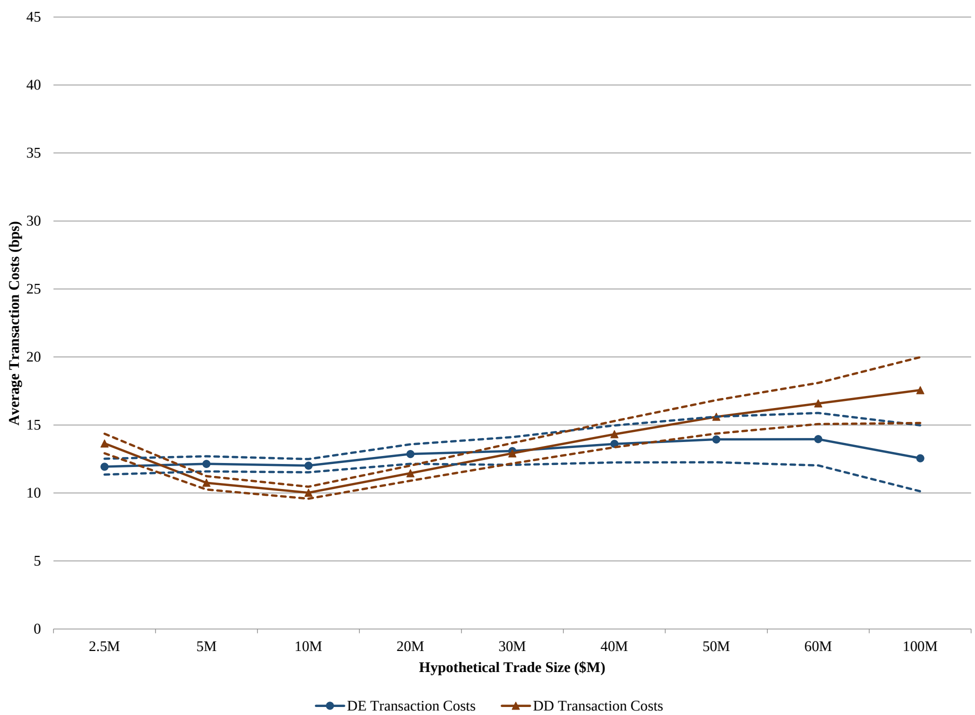

Figure IV: CDS and Corporate-bond Transaction Costs

These are precision-weighted cross-sectional averages of CDS DE transaction costs and corporate-bond transaction costs by actual trade-size bin. The sample includes 739 single-name CDS matched to either (1) the corporate bonds they reference, (2) the cheapest bond of the entity they reference, or (3) all traded bonds of the entity they reference. DE CDS and bond trades are for the period August 1, 2009 – December 31, 2012. The cost estimates, which are effective half-spreads, are obtained from time-series regressions estimated separately for each CDS/bond contract in the sample. For each CDS/bond we evaluate the estimated model at actual trade sizes, and then assign the resultant transaction cost estimates to different bins based on the notional amount traded. The weights used to compute the cross-sectional averages are the inverses of the estimated estimator variances of the respective cost estimates.