Market Structure and Transaction Costs of Index CDSs*

Pierre Collin-Dufresne† Benjamin Junge‡ Anders B. Trolle§

Abstract

Using transaction data, we study the two-tiered structure of the post-Dodd-Frank index CDS market. We identify dealer-to-client (D2C) trades and interdealer (D2D) trades. Average transaction costs are higher for D2C trades, reflecting higher average price impact. D2C trades Granger-cause D2D trades consistent with the interdealer market being used for managing inventory risk. Unique order-book data show that D2D transaction costs and price impacts vary across trading protocols, with mid-market matching and workup attracting liquidity-motivated trades. D2C prices are typically better than those available on the main interdealer limit order book, which may explain the endurance of the two-tiered market structure.

JEL Classification: G12, G13, G14, G28

Keywords: CDX, Dodd-Frank Act, Matching, Swap Execution Facility, Workup

This version: April 7, 2017

*We thank Damien Ackerer, Bruno Biais, Darrell Duffie, Thierry Foucault, Michael Johannes, René Kallestrup, Damien Klossner, Laurence Lescourret, Rémy Praz, Ioanid Rosu, Søren Willemann, Hongjun Yan, Alex Zhou, Haoxiang Zhu, and seminar participants at BI Norwegian Business School, Bank of England, Boston University, the Commodity Futures Trading Commission, Cornerstone Research, Deutsche Bundesbank, ESSEC, EPFL, the European Central Bank, the Federal Reserve Bank of New York, the Federal Reserve Board, Frankfurt School of Finance & Management, HEC Paris, McGill University, Rutgers University, University of New South Wales, the 12th Annual Central Bank Conference on Microstructure of Financial Markets, the 2016 Paris Finance Meeting, the 2016 SFI Research Days, the 2017 Chicago Financial Institutions Conference, and the 2017 Midwest Finance Association Annual Meeting for comments and suggestions. We gratefully acknowledge research support from the Swiss Finance Institute.

†EPFL and Swiss Finance Institute. E-mail: pierre.collin-dufresne@epfl.ch

‡EPFL and Swiss Finance Institute. E-mail: benjamin.junge@epfl.ch

§EPFL and Swiss Finance Institute. E-mail: anders.trolle@epfl.ch

1 Introduction

The index credit default swap (CDS) market constitutes an important component of the corporate credit market. Index CDSs allow banks, asset managers, and other institutional investors to efficiently hedge and trade aggregate credit risk in the economy. Unlike single-name CDSs, index CDSs have remained popular since the financial crisis with tens of billion dollars of notional amount traded on a daily basis. Nevertheless, little is known about the cost of trading in this important market.

The index CDS market is also interesting as a test case of how recent regulation introduced in the wake of the financial crisis affects the structure of swap markets. Since its inception in 2003, the index CDS market has operated as a classic two-tiered over-the-counter (OTC) market in which global derivatives dealers provide liquidity to clients in the dealer-to-client (D2C) segment of the market, and dealers trade among themselves in the interdealer (D2D) segment of the market. New swap market regulation following the Dodd-Frank Act had the potential to change this market structure by mandating trades in the most liquid index CDSs to be executed on so-called swap execution facilities (SEFs).[1] These regulated trading platforms are required to offer trading in order books, which opens up the market to all-to-all trading where clients can compete with dealers for liquidity provision. However, SEFs also offer trading via request for quote (RFQ), which mimics traditional trading in OTC markets. Interestingly, several years after the new regulation was fully implemented, all-to-all trading has yet to materialize. Instead, the two-tiered market structure persists, with D2C trades taking place on one group of SEFs (almost exclusively via name-disclosed RFQs) and D2D trades taking place on another group of SEFs (via a diverse set of trading protocols) run by interdealer brokers (IDBs).[2]

---

[1] Other key elements of the new swap market regulation are post-trade transparency via the immediate public dissemination of trades and central clearing of index CDSs with standardized contract terms.

[2] Referring to both the index CDS and the interest rate swap markets, a recent article summarize the current situation as "...dealer banks still trade together privately in one segment of the market and the buy side

From a regulatory and market design perspective it is important to understand the endurance of this bifurcated market structure. On the one hand, it could suggest that this is indeed the optimal structure of a market in which trades occur relatively infrequently and in very large sizes; see, e.g., Giancarlo (2015).[3] On the other hand, some market participants have accused dealers of resisting a transition to an all-to-all market structure in order to limit competition from non-dealer liquidity providers; see, e.g., Managed Funds Association (2015).

Motivated by these issues, the paper has three related objectives. First, we characterize the two-tiered structure of the post-Dodd-Frank index CDS market. Second, we analyze transaction costs and price impacts across market segments and trading protocols. Third, we estimate dealer profits from liquidity provision to clients.

We use transaction data from October 2, 2013 (the date on which the first SEFs started operating) to October 16, 2015, and we focus on the two most popular credit indices, CDX.IG and CDX.HY, which cover the investment-grade and high-yield components, respectively, of the North American corporate credit market. The transaction data include execution timestamps, transaction prices, and trade sizes up to certain notional caps. In addition, we develop algorithms that allow us to identify, for each trade, the SEF on which the trade took place and the type of trade (outright trade, index roll, curve trade, or delta hedge of an index swaption or tranche swap). The SEF on which the trade took place in turn reveals whether the trade is D2C or D2D.[4]

Trading volumes are large. The average daily notional amounts traded in the D2C segment are USD 9.843 billion for CDX.IG and USD 3.705 billion for CDX.HY. In the D2D

still executes via RFQ to the dealers in another. Proponents of this view say that nothing really changed in terms of how firms execute swaps except that the buy side has gone from RFQ-ing one dealer to RFQ-ing three. This appears to be in stark contrast to the all-to-all trading model envisioned for the swaps markets by regulators under Dodd-Frank." See "SEFs: A Market Divided," *Profit & Loss*, October 22, 2015.

[3]For example, for CDX.IG, one of the two most actively traded credit indices, there are 138 trades per day, on average, and the median trade size is USD 50 million.

[4]Because we identify D2C and D2D trades based on the SEF on which the trade took place, our sample is limited to the period during which SEFs were in operation and to trades that are executed on SEFs.

segment, the corresponding numbers are USD 1.354 billion and USD 0.402 billion. Outright trades account for the majority of trading volume. Index rolls constitute the second most important type of trade. Among outright trades, trading activity concentrates in five-year CDSs on the most recently issued (on-the-run) index. These trades are the main focus of the paper. Among index rolls, trading activity concentrates in rolls between five-year CDSs on the on-the-run index and five-year CDSs on the previous on-the-run (immediate off-the-run) index. These trades form an interesting comparison group as the trading motive for index rolls is different from that for outright trades.

We measure transaction cost by the effective half-spread, which is the difference between the transaction price and the contemporaneous value of the intraday mid-quote from Markit (both expressed in terms of par spreads). We measure price impact as the change in the mid-quote over a period of approximately 15 minutes following a trade. In case of outright trades, transaction costs of D2C trades are consistently higher than those of D2D trades. For CDX.IG, average transaction costs of D2C and D2D trades are 0.137 basis points (bps) and 0.088 bps, respectively, with the difference of 0.049 bps being statistically significant. The corresponding numbers for CDX.HY are 0.674 bps and 0.402 bps, respectively, with the difference of 0.273 bps again being statistically significant. These transaction cost differentials are largely due to D2C trades having higher price impact than D2D trades. For CDX.IG, average price impact is 0.043 bps higher for D2C trades, and for CDX.HY, average price impact is 0.262 bps higher for D2C trades. After taking price impact into account, there is no significant difference in transaction costs of D2C and D2D trades.

In contrast to outright trades, index rolls are motivated by investors seeking to maintain a liquid credit exposure with a relatively constant maturity profile. As such, they do not involve a significant risk transfer. Consistent with this, we find that transaction costs and price impacts of index rolls are both smaller than those of outright trades and similar across D2C and D2D index rolls.

We also investigate how trade characteristics and market conditions affect transaction costs and price impacts. Trade-by-trade regressions show that transaction costs and price impacts increase with trade size, Markit's intraday bid-ask spread, and volatility implied by index swaptions; i.e., options on index CDSs. However, our findings regarding differences in transaction costs and price impacts of D2C and D2D trades are robust to controlling for these determinants. Moreover, our findings also prevail in subsamples of pairs of D2C and D2D trades with matching trade characteristics that are executed at around the same time.

In addition, we analyze the dynamics of D2C trades, D2D trades, and quote revisions using a vector autoregressive (VAR) model in the spirit of Hasbrouck (1991a, 1991b). The VAR model allows to distinguish between permanent (information-driven) and transitory (inventory-driven) price impact, and to quantify the relative importance of trades for price discovery. In line with our findings based on the above-mentioned 15-minute price impact measure, the model-implied price impact is higher for D2C trades. In addition, price impact is largely permanent. The relatively high permanent price impact of D2C trades is suggestive of clients trading on information. This likely reflects the institutional nature of the index CDS market in which clients are professional investors who may have private information about the credit risk of certain index constituents (see, e.g., Acharya and Johnson (2007) and Ivashina and Sun (2011)) or may have an advantage over dealers in interpreting public information in relation to the aggregate credit risk in the economy.[5] The relatively low permanent price impact of D2D trades is suggestive of dealers mainly using the interdealer market to manage their inventory risk (as documented, for instance, by Reiss and Werner (1998) for the equity market). Consistent with this, we find one-way Granger causality from D2C trades to D2D trades. In terms of price discovery, D2C trades play a much more important role than D2D trades.

[5] In support of superior information processing by institutional investors, Hendershott, Livdan, and Schürhoff (2015) show that institutional order flow predicts the occurrence and sentiment of news as well as news-announcement-day equity market returns.

While virtually all D2C trades are executed via RFQs, a number of different trading protocols are used in the interdealer market. We refine the characterization of the interdealer market by investigating how transaction costs and price impacts vary across trading protocols. To this end, we exploit unique order-book data from the main IDB SEF, the GFI Swaps Exchange. In addition to a standard limit order book, this SEF offers two trading protocols—mid-market matching and workup—that facilitate trade by means of size discovery where orders are crossed at a fixed price (see, e.g., Duffie and Zhu (2017)).[6] In contrast to standard market orders, execution of orders for matching and workup is uncertain.

Mid-market matching is the dominant trading protocol; of the trading volume in five-year on-the-run CDX.IG and CDX.HY, it accounts for 52.2% and 58.6%, respectively. Workup is also frequently used and accounts for 19.9% and 15.6%, respectively. Mid-market matches have significantly lower average transaction cost and price impact than trades in the limit order book. This is consistent with Zhu's (2014) venue-selection model, in which liquidity traders prefer a mid-point dark pool (roughly equivalent to mid-market matching) that offers price improvement but does not guarantee execution, while informed traders prefer the certainty of executing against limit orders. By design, a workup is initiated by a trade in the limit order book and occurs at the same price. However, the average price impact of workups is close to zero implying that this trading protocol allows for increasing the size of a trade in the limit order book at little additional price impact. These results show that average transaction cost and price impact of D2D trades mask significant heterogeneity across trading protocols, with size-discovery protocols attracting liquidity-motivated trades.

Finally, we use the GFI data to estimate dealer profits from liquidity provision to clients in

[6] The two trading protocols differ in how the price is fixed, for how long orders can be crossed, and what information about unfilled interests is available to market participants. In case of mid-market matching, the price is fixed by a broker, orders can be crossed until the next time the broker resets the price, and market participants are informed when there is interest for matching. The direction and size of interests are not revealed. In case of workup, the price is fixed by an initiating trade in the limit order book, orders can be crossed for a short period of time following the initiating trade, and market participants are informed about the direction and size of interests.

five-year on-the-run index CDSs. Assuming that dealers immediately close D2C trades at the mid-market level, estimated daily profits are USD 0.433 million for CDX.IG and USD 0.808 million for CDX.HY. However, assuming instead that dealers immediately close positions at the best bid or offer on the limit order book, estimated profits are negative. Because of the execution risk associated with mid-market matching, this suggests that dealers only make profits through their willingness to bear inventory risk.

From the perspective of clients who value immediacy, our results show that the current market structure delivers very low transaction costs. The prices that clients obtain via name-disclosed RFQs are typically better than those available on the limit order book of the main IDB SEF. Indeed, 96.0% and 96.6% of the D2C trades in five-year on-the-run CDX.IG and CDX.HY, respectively, are executed at prices that are strictly more favorable than the contemporaneous best bid or offer on the GFI limit order book. This suggests that the two-tiered market structure—at least when combined with regulatory measures limiting dealer market power, such as post-trade transparency and a requirement to put a minimum number of dealers in competition for trades—constitutes a viable alternative to all-to-all trading in swap markets.[7]

1.1 Related Literature

The paper relates to a number of recent studies of how various provisions of the Dodd-Frank Act affect swap market liquidity. Loon and Zhong (2016) show that post-trade transparency and central clearing have a positive impact on liquidity in the index CDS market. Benos, Payne, and Vasios (2016) show that pre-trade transparency (the mandate to trade on SEFs)

[7] Proponents of bringing all market participants onto one limit order book typically argue that it would (i) increase quote competition among dealers and (ii) allow clients to occasionally supply liquidity via limit orders thereby lowering overall transaction costs (although at the cost of execution risk). However, a limit order book arguably works best when trading is continuous and it is not necessarily optimal when trading is more episodic as is the case for index CDSs. For instance, Barclay, Hendershott, and Kotz (2006) document a precipitous drop in electronic trading (via limit order books) when Treasuries go off-the-run and trading volumes decline.

has a positive impact on liquidity in the interest rate swap market. In contrast, we focus on the structure of the index CDS market after the full implementation of the new swap market regulation and analyze transaction costs across market segments and trading protocols.

The paper also relates to studies of transaction costs in the related markets for single-name CDSs and corporate bonds, both of which function as traditional OTC markets with relatively high search costs. For single-name CDSs, Biswas, Nikolova, and Stahel (2015) report average effective half-spreads in upfront terms of 14 bps for D2C trades in typical sizes of approximately USD 5 million. For a recent sample of corporate bonds, Harris (2015) reports average relative effective half-spreads in price terms of 39 bps for institutional-sized D2C trades.[8] For comparison, for D2C trades in CDX.IG and CDX.HY, average effective half-spreads in upfront terms are 0.66 bps and 3.03 bps, respectively, and average relative effective half-spreads in price terms are 0.65 bps and 2.86 bps. As such, transaction costs of index CDSs are about an order of magnitude lower than those of single-name CDSs and corporate bonds. Also, in contrast to virtually all corporate bond studies, but consistent with standard models of asymmetric information and inventory control, we find that transaction costs increase with trade size.

Finally, the paper relates to empirical studies of size-discovery trading protocols. Size discovery is widely used in the equity market (in the form of mid-point dark pools) and in the Treasury market (in the form of workup). In both markets, the trades that occur through size discovery tend to be less informed; see, e.g., Comerton-Forde and Putniņš (2015) and Fleming and Nguyen (2015). We provide the first analysis of size discovery in swap markets. In addition, we study a trading platform that offers two size-discovery protocols—mid-market matching and workup—providing insights into the relative importance and different impacts of the two trading protocols.

[8] We benchmark against Harris (2015) because his sample period overlaps with ours and his method of computing transaction costs is similar to ours. Earlier studies using different methodologies also report large transaction costs of corporate bonds; see, e.g., Edwards, Harris, and Piwowar (2007), Goldstein, Hotchkiss, and Sirri (2007), and Hendershott and Madhavan (2015).

Theoretically, a number of papers show how a two-tiered market structure can arise endogenously (see, e.g., Atkeson, Eisfeldt, and Weill (2013), Babus and Parlatore (2016), and Wang (2016)). Dunne, Hau, and Moore (2015) model price formation in a two-tiered market structure but take the structure as given.

The paper is organized as follows: Section 2 describes the structure of the index CDS market and the regulatory reforms set forth by the Dodd-Frank Act. Section 3 discusses the data and the identification algorithms. Section 4 compares D2C and D2D transaction costs and investigates how transaction costs vary with trade characteristics and market conditions. Section 5 analyzes the dynamics of trades and quotes using a VAR model. Section 6 uses GFI data to compare D2D transaction costs across different trading protocols and to estimate dealer profits from liquidity provision to clients. Section 7 concludes, and data-related details and robustness checks are contained in an Internet Appendix.

2 The Index CDS Market

This section describes index CDSs and the structure of the market in which these contracts trade. Furthermore, it discusses regulatory reforms set forth by the Dodd-Frank Act.

2.1 Index Credit Default Swaps

An index CDS is a standardized credit derivative contract on a diversified index of reference entities (typically, companies). Over the life of the contract, the credit protection seller provides default protection on each index constituent and, in return, receives periodic premium payments according to the fixed spread of the contract. At initiation, counterparties exchange an upfront amount equal to the present value of the contract. However, when quoting a contract, market participants either use the “par spread” or the “price.” The par spread is the fixed spread that makes the upfront amount equal to zero, and the price is one minus the upfront amount per dollar of notional.[9] We use par spreads throughout. Typically, contract tenors between one and ten years can be traded but the five-year contract tenor is the most liquid.

Twice a year, on the so-called index roll dates in March and September, a new index—or, more precisely, a new series of an index—is launched, with companies being revised according to credit rating and liquidity criteria.[10] Companies that fail to maintain a credit rating within a specified range, due to either up- or downgrades, and companies whose single-name CDSs have deteriorated significantly in terms of their trading activity are replaced by the most actively traded companies meeting the credit rating requirements. Typically, liquidity is concentrated in the most recently launched index, which is referred to as the on-the-run index. All previously launched indices are referred to as off-the-run indices.

The administrator of the most popular credit indices is Markit, and its benchmark credit indices of investment-grade and high-yield credit risk in North America are CDX.IG and CDX.HY, respectively. The former comprises 125 North American companies with investment-grade credit ratings, and the latter comprises 100 North American companies with non-investment-grade credit ratings. These indices are the focus of the paper.

2.2 Pre-Dodd-Frank Market Structure

Index CDSs used to be traded in a relatively opaque two-tiered OTC market. In the D2C segment of the market, dealers provided liquidity to their institutional clients. D2C trades were either negotiated over the phone or executed electronically on trading platforms, such as MarketAxess or Tradeweb.[11] Electronic trade execution was typically via name-disclosed

[9] The two ways of quoting contracts are equivalent under the market convention of using the ISDA CDS Standard Model to convert par spreads to prices, and vice versa.

[10] An index's series number uniquely determines the reference entities in the index.

[11] Electronic trading platforms for index CDSs emerged in 2005 (see "MarketAxess Launches CDS Index Trading Platform," *Risk Magazine*, September 12, 2005, and "TradeWeb Launches its Global Online Market for Credit Derivatives: TradeWeb CDS," *Press Release*, October 26, 2005), but their share of trading volume is unknown.

RFQs that enable querying multiple dealers simultaneously for an executable price for a given notional amount.

In the D2D segment of the market, dealers traded with each other typically involving IDB intermediation. D2D trades were either voice brokered or executed electronically on an IDB's order book. IDB intermediation guaranteed that trades were executed anonymously and that access to the interdealer market was restricted to dealers.

2.3 The Dodd-Frank Act and Current Market Structure

The Dodd-Frank Act tasked the Commodity Futures Trading Commission (CFTC) with regulating the index CDS market in order to promote financial stability as well as post- and pre-trade transparency. Pursuing these objectives, the CFTC enacted a clearing requirement for index CDSs with standardized contract terms as well as a reporting requirement and a trade execution requirement.[12]

The reporting requirement mandates real-time trade reporting of all index CDS trades to so-called swap data repositories (SDRs). SDRs publicly disseminate the received transaction data; dissemination is immediate unless the trade qualifies as a block in which case dissemination is delayed by at least 15 minutes.[13]

The trade execution requirement mandates that the most liquid index CDSs trade on SEFs and via one of two trading methods: an order book or an RFQ that is transmitted to at least three other market participants on the SEF.[14] Since the trade execution requirement took effect, trades in five-year on-the-run and immediate off-the-run index CDSs on CDX.IG

[12] See Part 50, Part 43, and Part 37 of Chapter I of Title 17 of the Code of Federal Regulations (17 CFR) and Section 2(h) of the Commodity Exchange Act (CEA).

[13] Block trades have notional amounts that exceed certain minimum block sizes and are exempt from immediate dissemination to protect liquidity providers in block-sized trades from front running. Minimum block sizes depend on the par spread and contract tenor (see Appendix F to Part 43 of Chapter I of 17 CFR for the mapping of spread-contract-tenor pairs to block sizes).

[14] For an interim one-year period, it was sufficient to transmit RFQs to at least two other participants.

and CDX.HY have been subject to the requirement.[15] Block trades are exempt from the trade execution requirement.

The implementation of Dodd-Frank Act provisions for index CDSs was rolled out in stages over a period of about one year. For dealers the reporting requirement took effect on December 31, 2012 and the clearing requirement took effect on March 11, 2013. By the time the first SEFs started operating on October 2, 2013, the trade reporting and clearing requirements were in effect for all market participants. Finally, the trade execution requirement took effect on February 26, 2014. Appendix A provides a timeline with additional details concerning the CFTC's implementation of Dodd-Frank Act provisions.

Through the introduction of SEFs and the requirement that they offer trading in order books, the new regulation had the potential to open up the index CDS market to all-to-all trading.[16] However, several years into the new regulatory regime, the index CDS market remains two-tiered and all-to-all trading has yet to materialize. The D2C segment of the market migrated onto SEFs run by incumbent operators of D2C trading platforms where the vast majority of trades are executed via name-disclosed RFQs. These are Bloomberg SEF, ICE Swap Trade, MarketAxess SEF, and TW SEF; collectively called D2C SEFs. The D2D segment of the market migrated onto SEFs run by IDBs where most trades are executed on order books. These are GFI Swaps Exchange, ICAP SEF, tpSEF, and Tradition SEF; collectively called IDB SEFs.

Several reasons have been given for the persistence of the two-tiered market structure. At one end of the spectrum, some observers argue that this is the optimal structure of a market in which trades occur relatively infrequently and in very large sizes (see, e.g., Giancarlo

[15] In addition, trades in five-year on-the-run and immediate off-the-run index CDSs on iTraxx Europe and iTraxx Europe Crossover are subject to the trade execution requirement. iTraxx Europe and iTraxx Europe Crossover are Markit's benchmark credit indices of investment-grade and high-yield credit risk in Europe.

[16] Implicitly, the CFTC had hoped that the introduction of SEFs would push the index CDS market, and other active OTC derivatives markets, towards all-to-all trading. For instance, when discussing the benefits of SEF rules, the CFTC stated that the "...rules provide for an anonymous but transparent order book that will facilitate trading among market participants directly without having to route all trades through dealers." See 78 Federal Register at 33565 (Jun. 4, 2013).

(2015)). At the other end of the spectrum, some market participants argue that dealers try to build barriers to entry to the interdealer market (see, e.g., Managed Funds Association (2015)). One such barrier is post-trade name give-up on IDB SEFs; i.e., the practice of informing anonymously matched traders about the identity of their counterparty after the trade is executed. This makes participation on IDB SEFs unattractive for many clients because of the risk of uncontrolled information leakage of proprietary trading strategies.[17]

3 Data and Identification Algorithms

This section describes the transaction and quote data and the algorithms that identify SEFs and package transactions.

3.1 Data

Our empirical analysis is based on trades and quotes over a two-year period from October 2, 2013 to October 16, 2015. All trades are executed on SEFs. The transaction data come from the three SDRs that disseminate trade reports of index CDS transactions: the Bloomberg Swap Data Repository (BSDR), the Depository Trust & Clearing Corporation Data Repository (DDR), and the Intercontinental Exchange Trade Vault (ICETV). Trade reports contain execution timestamps, transaction prices, and trade sizes up to a cap of at least USD 100 million,[18] and they indicate whether the trade is centrally cleared, whether it features non-standard (or bespoke) contract terms, and whether it is subject to an end-user

[17]Trading via RFQ also entails a certain amount of information leakage, but in this case the client has control over which dealers receive the information. Because the vast majority of index CDSs are centrally cleared, there is no reason for post-trade name give-up from a counterparty risk perspective. However, some dealers argue that name give-up is needed to prevent predatory trading by clients (see, e.g., "How to Game a SEF: Banks Fear Arrival of Arbitrageurs," *Risk Magazine*, March 19, 2014).

[18]The actual cap size is the larger of USD 100 million and the minimum block size (see §43.4(h) of Chapter I of 17 CFR).

exception that exempts the trade from the clearing and trade execution requirements.[19] The trade reports also indicate whether the trade is executed on a SEF, but they do not specify which one. They also do not specify whether the trade is part of a package; i.e., a transaction that involves more than one index CDS or an index CDS and a related instrument, such as an index swaption or tranche swap (both of which are conventionally traded with delta, see below).[20] Fortunately, SEFs and package transactions can be identified from trade reports; the details of the respective identification algorithms are discussed in subsequent sections.

Intraday composite bid and offer quotes for index CDSs come from Markit. These quotes constitute the main real-time reference in the index CDS market that is available to all market participants. The composites average over quotes of individual dealers that Markit parses from so-called dealer runs; i.e., electronic messages that dealers send to their institutional clients throughout the trading day to keep them up to date with indicative quotes of index CDSs and other credit derivatives. A composite is computed whenever a dealer sends out a run and only the quotes from each dealer's latest run are eligible for composite computation.[21] The average number of quotes per day for CDX.IG and CDX.HY are 443 and 391, respectively, and 97% of quotes take place between 7:00 a.m. and 5:30 p.m., New York time.

Figure 1 shows trades and Markit intraday mid-quotes on a representative trading day, May 6, 2015, for the five-year index CDS on the then on-the-run series of CDX.IG. There are 401 mid-quotes and 165 trades. Most striking are the trades at 64 bps and 66 bps that appear to be outliers in comparison to the other trades that tend to be relatively close to the mid-quote. After processing the data through our identification algorithms, these trades

[19] This would be the case if one counterparty is a non-financial entity that uses the trade to hedge commercial risks (see Sections 2(h)(7) and 2(h)(8) of the CEA).

[20] There are other important trade characteristics that are not specified in the trade reports. For instance, trade reports do not specify whether the trade is buyer- or seller-initiated, whether it is D2C or D2D, and whether it is executed on an order book or via a RFQ.

[21] Quotes from runs older than 15 minutes are discarded from the computation and a five-minute memory prevents repeated computations of the same composite.

turn out to be delta hedges of index swaption trades, see below. Data processing also shows that, out of the 165 trades, 139 are executed on D2C SEFs (identified as D2C trades) and 26 are executed on IDB SEFs (identified as D2D trades).

[Figure 1 about here.]

3.2 Identification of SEFs

In devising the SEF identification algorithm, we use SEF-reported trading volumes from Clarus FT.[22] Each of the on-SEF trade reports must have been submitted by one of the eight aforementioned SEFs. Bloomberg SEF submits trade reports to the BSDR and ICE Swap Trade submits trade reports to the ICETV. The remaining SEFs submit trade reports to the DDR and the trade-report-submitting SEF can be identified based on the format of the trade report. Specifically, we associate with each SEF the format of trade reports whose aggregate trade size corresponds to the SEF-reported trading volume (the Internet Appendix contains the details).

Because of the two-tiered market structure, the SEF on which the trade took place reveals whether the trade is D2C or D2D. It should be emphasized that limiting the sample to trades executed on SEFs is not restrictive because we focus on trades in the most actively traded on-the-run and immediate off-the-run index CDSs. After February 26, 2014 trades in these index CDSs are required to be executed on SEFs. In the initial period from October 2, 2013 to February 25, 2014, on-SEF trade execution was non-mandatory but the majority of trades were executed on SEFs (the Internet Appendix contains additional details).

[22] Clarus FT is the standard data source for SEF-reported daily trading volumes. In the Internet Appendix, we describe the Clarus FT data in detail.

3.3 Identification of Package Transactions

We identify four popular types of package transactions: index rolls, curve trades, delta-hedged index swaptions, and delta-hedged index tranche swaps (the Internet Appendix contains the details). A typical index roll involves an on-the-run and an off-the-run index CDS with the same contract tenor. Protection is sold on one index series and simultaneously bought on the other. Index rolls are popular because many institutional investors like to maintain liquid credit exposure with a relatively constant maturity profile. We identify index rolls as simultaneously executed index CDS trades on the same SEF that have the same contract tenor and reference two different series of the same index.

A typical curve trade involves two index CDSs with different contract tenors.[23] Protection is sold on one contract tenor and simultaneously bought on the other. Curve trades are popular because they are relatively directional (index CDS term structures tend to become flatter when spreads widen and steeper when spreads contract; see, e.g., Erlandsson, Ghosh, and Rennison (2008)) and require less capital outlay than outright index CDS trades. We identify curve trades as simultaneously executed index CDS trades on the same SEF that have different contract tenors and reference the same index (but not necessarily the same index series).

We also account for the fact that index swaptions and tranche swaps are conventionally traded "with delta;" i.e., together with a delta hedge in the corresponding index CDS. Quotes of index swaptions and tranche swaps incorporate both the delta and the so-called "reference level" at which the delta hedge will be traded. Usually, the reference level is set close to the par spread at which the index CDS trades at the beginning of the trading day (see, e.g., Hünseler (2013)), but it might be updated throughout the trading day because of spread movements. For CDX.IG, the reference level is usually set in spread multiples of 0.5 bps.[24]

[23] Typically, the two index CDSs reference the same index series, but there are also curve trades in which the two index CDSs reference different index series.

[24] Because CDX.HY is quoted in terms of a price, the reference level is usually set in price multiples of 0.125%.

We identify index swaption and tranche swap delta hedges as index CDS trades that have the same underlying index and contract tenor as an index swaption or tranche swap trade. Trade executions must be near simultaneous and notional amounts must be reconcilable with a delta that is quoted on the same trading day.

Index swaptions and tranche swaps can also be traded without delta, but usually at less favorable prices that incorporate the dealer's cost of establishing the hedge. Therefore, investors may find it beneficial to trade index swaptions and tranche swaps with delta and unwind the hedge themselves (see, e.g., Hünseler (2013)). We identify such delta unwinds as trades with the same transaction price and notional amount as a delta hedge of an index swaption or tranche swap trade that occurs on the same trading day and SEF.

Whether a trade is part of a package is important because package transactions are either quoted in relative terms (index rolls and curve trades) or along with a price-forming quote for another instrument (delta hedges of index swaption and tranche swap trades). Therefore, transaction prices on the individual index CDS legs of package transactions do not necessarily have to reflect the par spread at which outright trades in the respective index CDSs would be executed. This is clearly the case for most of the delta hedges in Figure 1.

3.4 SEF Order Flow

Table 1 displays descriptive statistics of the enriched transaction data that allows to distinguish between D2C and D2D trades and between outright trades and package transactions. Descriptive statistics are computed separately for D2C and D2D trades in CDX.IG (Panels A1 and A2, respectively) and CDX.HY (Panels B1 and B2, respectively) and, within these broad categories of trades, descriptive statistics are computed separately for trades that are executed on a given SEF.

[Table 1 about here.]

The index CDS market is characterized by relatively few trades in very large sizes. For CDX.IG, there are 114 D2C trades and 24 D2D trades per day, on average, and the median trade size is USD 50 million in both segments. For CDX.HY, there are somewhat more trades (164 D2C trades and 27 D2D trades per day, on average), but the median trade size is smaller (USD 10 million in both segments) because of the significantly higher volatility of high-yield contracts.

Trading volumes are large. The average daily D2C trading volumes are USD 9.843 billion for CDX.IG and USD 3.705 billion for CDX.HY. The corresponding D2D trading volumes are USD 1.354 billion and USD 0.402 billion.[25] These averages appear in parenthesis in the table because they are based on SEF-reported daily trading volumes from Clarus FT instead of transaction data. They cannot be reproduced with transaction data because trade reports contain capped trade sizes. Indeed, the table shows that the fraction of D2C (D2D) trades that are disseminated with capped trade sizes are 21.2% (6.8%) for CDX.IG and 2.3% (1.4%) for CDX.HY.[26] As a consequence, transaction-data-based average daily trading volumes are downward biased.[27]

[25] D2D trading accounts for 10% (for CDX.HY) to 12% (for CDX.IG) of total volume in the index CDS market. The International Swaps and Derivatives Association (2014, ISDA) estimates that, in case of interest rate swaps, D2D trading accounts for 35% of total volume. However, the ISDA (2014) argues that as much as two-thirds of D2D trading is due to non-price-forming trades, such as amendments, novations, and terminations, all of which are excluded from our sample. This brings the ISDA's (2014) estimate for interest rate swaps more in line with the one we find for index CDSs in our sample.

[26] In comparison to trades in CDX.IG, the percentage of trades that are disseminated with capped trade sizes is lower for trades in CDX.HY because the latter tend to be of smaller size (in absolute terms and relative to the cap). The median size of trades in CDX.IG is five times that of trades in CDX.HY but caps typically differ by USD 10 million only (for trades in CDX.IG the cap is typically USD 110 million and for trades in CDX.HY the cap is typically USD 100 million).

[27] The actual volumes allow to impute by how much the size of trades that are disseminated with capped trade sizes exceeds the cap on average. For instance, the size of D2C trades in CDX.IG that are disseminated with capped notional amounts exceeds the cap by USD 141.17 ($= 511 \times (9,843 - 6,433)/(0.212 \times 58,222)$) million, on average (511 is the number of trading days in the sample period). Most of these trades are capped at USD 110 million, suggesting that, conditional on being capped, the average size of D2C trades in CDX.IG is approximately USD 250 million. Similarly, conditional on being capped, the average size of D2D trades in CDX.IG is approximately USD 200 million. For CDX.HY, most trades are capped at USD 100 million and, conditional on being capped, the average sizes of D2C and D2D trades in CDX.HY are approximately USD 225 million and USD 160 million, respectively.

The vast majority of trades are in the five-year contract tenor and around 90% of trades are in on-the-run index CDSs. Almost all trades have standardized contract terms and are centrally cleared.[28] The fact that there are virtually no D2D block trades, whereas about 20% of D2C trades are blocks is consistent with D2D trades occurring on order books. This is because block-sized trades that are executed on order books do not qualify as block trades. Outright trades account for most of the trading volume and, among package transactions, index rolls are most popular.

Table 2 shows the fraction of trades and trading volume that are outright trades in five-year on-the-run index CDSs, other outright trades, index rolls between five-year on-the-run and immediate off-the-run index CDSs, and other package transactions. Outright trades in five-year on-the-run CDX.IG and CDX.HY account for 88.5% and 84.6%, respectively, of D2C trading volume and 67.2% and 63.8% of D2D trading volume. These trades are the main focus of the paper. Index rolls between five-year on-the-run and immediate off-the-run CDX.IG and CDX.HY account for 4.9% and 8.8%, respectively, of D2C trading volume and 12.0% and 18.2% of D2D trading volume. These trades provide an interesting comparison to outright trades as they are predominantly liquidity motivated and do not involve a significant risk transfer.[29]

[Table 2 about here.]

[28] Loon and Zhong (2016) find that bespoke contract terms, central clearing, and a counterparty that qualifies as an end-user are trade characteristics that significantly affect transaction costs of index CDSs. These characteristics cannot be a main driver of potential transaction cost differences between D2C and D2D trades because the vast majority of both D2C and D2D trades are non-bespoke and centrally cleared, and there are few end-user exempt trades in our sample (there are no such trades after February 10, 2014).

[29] The Internet Appendix provides an analysis of outright trades in five-year immediate off-the-run index CDSs. Results are consistent with those of outright trades in five-year on-the-run index CDSs. For all other trade types there are too few trades to reliably measure transaction costs.

4 Transaction Costs

We now compare transaction costs and price impacts across D2C and D2D trades. We also control for differences in trade characteristics and market conditions under which trades are executed.

4.1 Transaction Cost Decomposition

We measure the cost of a transaction by the effective half-spread with respect to Markit's intraday mid-quote. We further decompose the effective half-spread into a price impact and a realized half-spread, which measures the transaction cost taking price impact into account. Specifically,

$$\underbrace{q_t(p_t - m_t)}_{=\text{EffcSprd}_t} = \underbrace{q_t(p_t - m_{t+\Delta})}_{=\text{RlzdSprd}_t} + \underbrace{q_t(m_{t+\Delta} - m_t)}_{=\text{PrcImp}_t}, \tag{1}$$

where p_t is the transaction price of the t-th trade in a given index CDS, m_t is the latest mid-quote in the 15-minute period prior to trade execution, and $m_{t+\Delta}$ is the first mid-quote in the 15-minute period that follows trade execution by 15 minutes. In case of index rolls, p_t is the difference in transaction prices of the involved on-the-run and immediate off-the-run index CDSs,[30] and m_t ($m_{t+\Delta}$) is the corresponding difference in mid-quotes.[31] Trade direction, q_t, is inferred by the Lee and Ready (1991) algorithm and equals +1 (−1) in case of protection-buyer-initiated (protection-seller-initiated) trades.

[30] Following market convention, p_t is the par spread of the on-the-run index CDS minus the par spread of the immediate off-the-run index CDS.

[31] Specifically, m_t is the corresponding difference in the latest mid-quotes prior to trade execution, with the later of the two quotes occurring in the 15-minute period prior to trade execution and the earlier of the two quotes occurring within 15 minutes from the later. Similarly, $m_{t+\Delta}$ is the corresponding difference in mid-quotes that occur after trade execution, with the later of the two quotes being the first quote on either of the two index CDSs that occurs in the 15-minute period that follows trade execution by 15 minutes and the earlier of the two quotes being the latest quote on the other index CDS that occurs within 15 minutes from the later of the two quotes.

4.2 Transaction Costs Across Market Segments

Figure 2 shows weekly averages of effective half-spreads, realized half-spreads, and price impacts of outright D2C and D2D trades. Panels A and B show that, for both indices, D2C trades have consistently higher effective half-spreads than D2D trades. Panels E and F show that D2C trades also have consistently higher price impacts than D2D trades, suggesting that transaction cost differentials reflect differences in price impacts. Panels C and D support this conjecture as there are no systematic differences between the realized half-spreads of D2C and D2D trades.

[Figure 2 about here.]

Table 3 displays average effective half-spreads, realized half-spreads, and price impacts of outright trades and index rolls. For outright trades the results confirm the impression from Figure 2. In case of CDX.IG, average effective half-spreads for D2C and D2D trades are 0.137 bps and 0.088 bps, respectively, with the difference of 0.049 bps being statistically significant. The corresponding numbers for CDX.HY are 0.674 bps and 0.402 bps, respectively, with the difference of 0.273 bps again being statistically significant.

[Table 3 about here.]

These transaction cost differentials are largely due to D2C trades having higher price impacts than D2D trades. For CDX.IG, average price impacts for D2C and D2D trades are 0.106 bps and 0.063 bps, respectively, with the difference of 0.043 bps being statistically significant. The corresponding numbers for CDX.HY are 0.508 bps and 0.246 bps, respectively, with the difference of 0.262 bps again being statistically significant. After taking price impact into account, there is no significant difference in transaction costs (as captured by realized half-spreads) of D2C and D2D trades.

As explained in Section 3.3, index rolls are liquidity motivated. Consistent with a non-informational motive for trade, Table 3 shows that index rolls have lower average effective half-spreads and price impacts than outright trades. For index rolls there are also no significant differences in average transaction costs and price impacts across D2C and D2D trades.

Table 4 focuses on outright trades only and displays average effective half-spreads, realized half-spreads, and price impacts by quartiles of the trade size distribution. In case of both indices and regardless of the quartile of the trade size distribution, effective half-spreads and price impacts of D2C trades are significantly higher than those of D2D trades.

Effective half-spreads of D2C trades in both indices increase with trade size which is in contrast to evidence from other dealer markets, such as the corporate and municipal bond markets, where transaction costs typically decrease with trade size; see, e.g., Bessembinder, Maxwell, and Venkataraman (2006), Edwards et al. (2007), Goldstein et al. (2007), Harris and Piwowar (2006), and Green, Hollifield, and Schürhoff (2007). This reflects structural differences between the two markets: the index CDS market is purely institutional with professional investors trading in large sizes; in contrast, there is significant retail participation in bond markets with relatively unsophisticated market participants trading in small sizes with dealers who can exert market power.

[Table 4 about here.]

Price impact of D2C trades in both indices tends to increase with trade size as well but only up to the third quartile of the trade size distribution. The decrease of price impact for block-sized trades in the fourth quartile of the trade size distribution is consistent with block trade provisions that aim at mitigating the price impact of block-sized trades.

4.3 Accounting for Trade Characteristics and Market Conditions

The evidence thus far does not account for the possibility that different trade characteristics (other than trade size) of D2C and D2D trades and potentially different market conditions under which trades are executed can explain the observed differences in average effective half-spreads and price impacts. In order to rule out such possibilities (or selection biases), we estimate selection-bias-corrected averages from trade-by-trade regressions that control for trade characteristics and market conditions, and we analyze pairs of trades with matching trade characteristics that are executed at around the same time.

Trade-By-Trade Regressions

We estimate the following trade-by-trade regressions:[32]

$$y_t = \alpha_{\text{D2C}}\text{D2C}_t + \alpha_{\text{D2D}}\text{D2D}_t + \beta' X_t + \epsilon_t, \tag{2}$$

where y_t is the either the effective half-spread, the realized half-spread, or the price impact of the t-th trade in a given index CDS, D2C_t and D2D_t are dummy variables for D2C and D2D trades, respectively, and X_t is a vector of control variables. Continuous control variables are stated in deviations from their sample means for ease of interpretation.

The continuous control variables include the bid-ask spread of the latest quote, the corresponding mid-quote, and the end-of-day implied volatility of the at-the-money three-month option on the index CDS. In addition, we include a set of dummy variables for trades with sizes in the second, third, and fourth quartile of the trade size distribution, and a dummy variable for trades with transaction prices at which reference levels of index swaption and

[32] Comparing regression-based methods for addressing selection biases, Bessembinder (2003) concludes "...while it is important to control for selection biases, the specific method of control has little practical effect on inference regarding market quality. In particular, the simple technique of including in a regression framework economic variables that are known to be related to trade execution costs appears to provide selectivity bias corrections that work as well as more complex two-stage methods."

tranche swap trades tend to be set. The continuous control variables proxy for the prevailing market conditions at trade execution. We account for trade size because Table 4 shows that transaction costs and price impacts vary with trade size, and we include a reference level dummy to account for potentially unidentified index swaption and tranche swap delta hedges.[33]

Due to demeaned continuous control variables, α_{D2C} and α_{D2D}, respectively, estimate average effective half-spreads (or, depending on the dependent variable used, realized half-spreads or price impacts) of outright D2C and D2D trades that have trade sizes in the first quartile of the trade size distribution and non-reference-level transaction prices, and that are executed when average market conditions prevail. Note that the estimates are directly comparable with those reported in Table 3 because the latter correspond to coefficient estimates of a restricted version of Equation (2) which excludes control variables.

Table 5 displays regression results. Accounting for trade characteristics and market conditions does not materially change the conclusions from Table 3. For CDX.IG, the difference in effective half-spreads of D2C and D2D trades is 0.033 bps (in comparison to 0.049 bps in Table 3) and statistically significant. For CDX.HY, the difference is 0.219 bps (in comparison to 0.273 bps in Table 3) and statistically significant. For both indices, the estimated regression coefficients show that transaction costs increase with trade size, quoted bid-ask spread, and implied volatility. In addition, trades with reference level transaction prices are more expensive.

[Table 5 about here.]

Again, transaction cost differentials are largely due to differences in price impact. For CDX.IG, the difference in price impacts of D2C and D2D trades is a statistically significant

[33] One reason for unidentified delta hedges is that we only identify delta hedges of on-SEF index swaption and tranche swap trades, but neither swaptions nor tranche swaps have to be traded on SEFs. Nevertheless, the delta hedges of off-SEF index swaption and tranche swap trades would typically be executed on SEFs in order to satisfy other regulatory requirements.

0.027 bps, accounting for most of the 0.033 bps difference in effective half-spreads. For CDX.HY, the difference in price impacts is a statistically significant 0.218 bps, accounting for almost the entire 0.219 bps difference in effective half-spreads. It follows that there are no significant differences in realized half-spreads of D2C and D2D trades. For both indices, the estimated regression coefficients show that price impacts increase with quoted bid-ask spread and implied volatility. Price impacts tend to increase with trade size; however, for CDX.HY, block-sized trades have lower price impacts than large-sized trades in the third quartile of the trade size distribution.

Matched Pair Analysis

Alternatively, trade characteristics and market conditions can be controlled for by focussing on pairs of D2C and D2D trades with matching trade characteristics that are executed relatively close in time. To this end, we focus on those outright D2D trades for which we are able to find at least one matching outright D2C trade with trade size in the same quartile of the trade size distribution (or, in one analysis, with exactly the same trade size) that occurs within a 15-minute window bracketing the execution of the D2D trade. In case of more than one matching D2C trade, the match is a hypothetical trade with effective half-spread, realized half-spread, and price impact corresponding to the average value among matching D2C trades.[34]

Table 6 shows the results. In case of CDX.IG, 52.7% of D2D trades have a matching D2C trade with trade size in the same quartile of the trade size distribution, and 38.0% of D2D trades have a matching D2C trade with exactly the same trade size. The pairs of trades with exactly matching trade sizes consist of D2C trades with an average effective half-spread of 0.124 bps (which is slightly less than in the full sample) and D2D trades with an average

[34] Similar matching methods have, for instance, been used by Lee (1993) to construct a sample of New York Stock Exchange trades that match the characteristics of a given set of OTC and regional exchange equity trades.

effective half-spread of 0.097 bps (which is slightly more than in the full sample). The average paired difference in effective half-spreads of matching D2C and D2D trades is 0.027 bps and statistically significant. The average paired difference in price impacts is 0.019 bps and also statistically significant. There is no significant difference in realized half-spreads of matching D2C and D2D trades.

[Table 6 about here.]

Most of these observations carry over to pairs of matched trades with trade sizes in the same quartile of the trade size distribution and to pairs of matched trades in CDX.HY. Overall, the results of the matched pair analysis are consistent with those of the trade-by-trade regressions, both in terms of the magnitude of differences between D2C and D2D trades and in terms of inference.

The Internet Appendix shows that our results are robust to using an alternative mid-quote from Credit Market Analysis. It also shows that our results are robust to using both shorter and longer time periods over which to compute realized half-spreads and price impacts and to using alternative time windows when constructing pairs of trades with matching characteristics.

5 The Dynamics of Trades and Quotes

We now consider a VAR model in the spirit of Hasbrouck (1991a, 1991b). The model provides a robust and flexible tool to study the dynamic interaction between D2C and D2D trades, to distinguish permanent (information-driven) and transitory (inventory-driven) price impacts, and to quantify the relative importance of D2C and D2D trades for price discovery.

5.1 VAR Model and Estimation

Specifically, we estimate an event-time VAR model for mid-quote changes, Δm_t, and D2C- and D2D-trade-related variables, x_t^{D2C} and x_t^{D2D}, respectively; that is,

$$\Delta m_t = \sum_{j=1}^{20} \alpha_j \Delta m_{t-j} + \sum_{j=0}^{20} \beta_j x_{t-j}^{\text{D2C}} + \sum_{j=0}^{20} \gamma_j x_{t-j}^{\text{D2D}} + \epsilon_t^{\Delta m}, \tag{3a}$$

$$x_t^{\text{D2C}} = \sum_{j=1}^{20} \delta_j \Delta m_{t-j} + \sum_{j=1}^{20} \zeta_j x_{t-j}^{\text{D2C}} + \sum_{j=1}^{20} \eta_j x_{t-j}^{\text{D2D}} + \epsilon_t^{\text{D2C}}, \tag{3b}$$

$$x_t^{\text{D2D}} = \sum_{j=1}^{20} \kappa_j \Delta m_{t-j} + \sum_{j=0}^{20} \lambda_j x_{t-j}^{\text{D2C}} + \sum_{j=1}^{20} \rho_j x_{t-j}^{\text{D2D}} + \epsilon_t^{\text{D2D}}, \tag{3c}$$

where t indexes the t-th quote revision and x_t^{D2C} (x_t^{D2D}) is the number of signed D2C (D2D) trades that occur between the $t-1$-th and t-th quote revision (i.e., x_t^{D2C} and x_t^{D2D} are sums of the above trade direction indicators, q_u, with u between the calendar time of the $t-1$-th and t-th quote revision). The error terms, $\epsilon_t^{\Delta m}$, ϵ_t^{D2C}, and ϵ_t^{D2D}, are uncorrelated because we resolve contemporaneous effects by including contemporaneous trade-related variables in Equations (3a) and (3c).[35] Intuitively, the D2C-trade-related variable may contemporaneously affect the D2D-trade-related variable when dealers immediately offload inventory in the interdealer market, and D2C- and D2D-trade-related variables may contemporaneously affect mid-quote revisions when dealers immediately adjust quotes in response to trades.[36]

The vector moving average (VMA) representation of the VAR model is used to characterize price impact and study the role of trades in the price discovery process. In our case,

[35] Moreover, error terms are assumed to be serially uncorrelated and homoscedastic.

[36] Our results are not sensitive to how we resolve contemporaneous effects between D2C- and D2D-trade-related variables. Virtually identical results obtain when we instead allow the D2D-trade-related variable to contemporaneously affect the D2C-trade-related variable.

it is given by

$$\Delta m_t = \epsilon_t^{\Delta m} + a_1\epsilon_{t-1}^{\Delta m} + \cdots + b_0\epsilon_t^{\text{D2C}} + b_1\epsilon_{t-1}^{\text{D2C}} + \cdots + c_0\epsilon_t^{\text{D2D}} + c_1\epsilon_{t-1}^{\text{D2D}} + \cdots, \tag{4a}$$

$$x_t^{\text{D2C}} = d_1\epsilon_{t-1}^{\Delta m} + d_2\epsilon_{t-2}^{\Delta m} + \cdots + \epsilon_t^{\text{D2C}} + z_1\epsilon_{t-1}^{\text{D2C}} + \cdots + h_1\epsilon_{t-1}^{\text{D2D}} + h_2\epsilon_{t-2}^{\text{D2D}} + \cdots, \tag{4b}$$

$$x_t^{\text{D2D}} = k_1\epsilon_{t-1}^{\Delta m} + k_2\epsilon_{t-2}^{\Delta m} + \cdots + l_0\epsilon_t^{\text{D2C}} + l_1\epsilon_{t-1}^{\text{D2C}} + \cdots + \epsilon_t^{\text{D2D}} + r_1\epsilon_{t-1}^{\text{D2D}} + \cdots. \tag{4c}$$

It immediately follows from Equation (4a), that the permanent price impact of a single protection-buyer-initiated D2C trade is

$$\Lambda_{\text{D2C}} = \lim_{n\to\infty}\sum_{j=0}^{n}\mathbb{E}\Big[\Delta m_{t+j}\Big|\Omega_t^{\text{D2C}}\Big] = \lim_{n\to\infty}\sum_{j=0}^{n} b_j = \sum_{j=0}^{\infty} b_j, \tag{5}$$

where $\Omega_t^{\text{D2C}} = \Big\{\epsilon_t^{\text{D2C}} = 1,\ \epsilon_t^{\Delta m} = \epsilon_t^{\text{D2D}} = 0,\ \epsilon_s^{\Delta m} = \epsilon_s^{\text{D2C}} = \epsilon_s^{\text{D2D}} = 0,\ s < t\Big\}$ denotes the event of an isolated unit-sized shock of the D2C-trade-related variable.[37] Similarly, the permanent price impact of a single protection-buyer-initiated D2D trade is $\Lambda_{\text{D2D}} = \sum_{j=0}^{\infty} c_j$.

Moreover, the VAR model is consistent with a fairly general unobserved component model. Accordingly, $m_t = \bar{p}_t + s_t$, where $\bar{p}_t$ is the (unobservable) efficient price and s_t is (unobservable) microstructure noise. The former is assumed to follow a random walk (which is, for instance, consistent with $\bar{p}_t$ being the conditional expectation of some future payoff), while the latter is a generic covariance stationary process with mean zero (which is, for instance, consistent with the transient nature of most microstructure effects, such as inventory-control-driven price pressure). Hasbrouck (1991b) shows that the variance of efficient price innovations, $\sigma^2_{\Delta\bar{p}}$, can be explicitly expressed in terms of error term variances

[37] A single protection-buyer-initiated D2C trade is an event in Ω_t^{D2C} but obviously not the only event that gives rise to a unit-sized shock of the D2C-trade-related variable. For instance, occurrence of two protection-buyer-initiated D2C trades and one protection-seller-initiated D2C trade between the $t-1$-th and t-th quote revision also result in a unit-sized shock of the D2C-trade-related variable.

and VMA-representation parameters; that is,

$$\sigma^2_{\Delta\bar{p}} = \left(\sum_{j=0}^{\infty} a_j\right)^2 \sigma^2_{\Delta m} + \left(\sum_{j=0}^{\infty} b_j\right)^2 \sigma^2_{\text{D2C}} + \left(\sum_{j=0}^{\infty} c_j\right)^2 \sigma^2_{\text{D2D}}, \tag{6}$$

where $\sigma^2_{\Delta m} = \mathbb{V}\left(\epsilon_t^{\Delta m}\right)$, $\sigma^2_{\text{D2C}} = \mathbb{V}\left(\epsilon_t^{\text{D2C}}\right)$, and $\sigma^2_{\text{D2D}} = \mathbb{V}\left(\epsilon_t^{\text{D2D}}\right)$. Equation (6) reflects a decomposition of efficient price innovations into three mutually orthogonal components: a trade-unrelated component with variance given by the first term on the right hand side of Equation (6), and two trade-related components with variances given by the second and third term. The first trade-related component is associated with D2C trades and the second one is associated with D2D trades. Equation (6) is the basis of our price discovery metric, Hasbrouck's (1991b) R^2, that expresses each component's variance as a fraction of $\sigma^2_{\Delta\bar{p}}$.

We estimate the VAR model using all quote changes between 7:00 a.m. and 5:30 p.m., New York time, during the sample period. We exclude a few intraday periods during which quotes are stale.[38] Finally, we winsorize quote changes at the 0.1% and 99.9% quantile of their distribution.

5.2 Results

Panels A1 and A2 of Table 7 display VAR coefficient estimates for CDX.IG and CDX.HY, respectively. The results are similar for both indices and, therefore, our discussion focuses on CDX.IG. The significant coefficients of contemporaneous trade-related variables in Equation (3a) suggest that dealers immediately raise mid-quotes by 0.009 bps and 0.003 bps in response to single protection-buyer-initiated D2C and D2D trades, respectively. Mid-quotes tend to be raised further in subsequent revisions due to the generally positive and significant coefficients of lagged variables in the equation (the table reports sums of coefficients of lagged

[38] During these periods there are typically neither quotes for CDX.IG nor for CDX.HY suggesting technical disruptions to the quote generation process.

variables and the corresponding t-statistics).

The generally positive and significant coefficients of lagged D2C-trade-related (D2D-trade-related) variables in Equation (3b) (Equation (3c)) indicate positively autocorrelated trades.[39] In contrast, standard models of inventory control in which an individual dealer sets quotes to elicit client trades in the direction of inventory (i.e., who reduces quotes to elicit protection-buyer-initiated client trades when being a net protection buyer, and vice versa, when being a net protection seller) would imply negative autocorrelation in D2C trades—at least at longer lags.

Coefficients of contemporaneous and lagged D2C-trade-related variables in Equation (3c) are generally positive and significant. Furthermore, Granger causality tests show that D2C- and D2D-trade-related variables are characterized by one-way Granger causality from D2C trades to D2D trades. This is consistent with inventory management taking place in the interdealer market.[40]

[Table 7 about here.]

Figure 3 displays the model-implied price impacts of D2C and D2D trades. Specifically, the figure traces out the cumulative quote revision following a single protection-buyer-initiated D2C or D2D trade. The price impact is higher for D2C trades, consistent with the findings in Section 4. Panel B of Table 7 displays the permanent price impact of a trade—i.e., the long-run limit of the cumulative quote revisions exhibited in Figure 3—and formally rejects the hypothesis of identical permanent price impacts of D2C and D2D trades.

Asymmetric information entails a permanent price impact, while inventory control entails a transient price impact. We do not observe a transient component in the price impact of D2C trades, which again suggests that the quote revision mechanism in standard inventory

[39] Persistence in order flow has been found to characterize trade of many financial securities after Hasbrouck and Ho (1987) provided initial evidence for equities.

[40] In addition, dealers have told us that CDX positions are hedged using other liquid instruments, such as S&P 500 futures.

control models is largely absent in this market—at least for the horizon that is captured by the VAR model.[41] Instead, the permanent price impact of D2C trades points to clients trading on information. This could, for instance, be private information about the credit risk of certain index constituents (see, e.g., Acharya and Johnson (2007) and Ivashina and Sun (2011)) or stem from an advantage over dealers in interpreting public information in relation to the aggregate credit risk in the economy. The lower permanent price impact of D2D trades is consistent with dealers mainly using the interdealer market for inventory management. From this it also follows that D2C trades play a more important role in the price discovery process than D2D trades. Indeed, Panel C of Table 7 shows that D2C trades account for a much greater proportion of the efficient price variance than D2D trades.

[Figure 3 about here.]

The Internet Appendix shows that our results are robust to different specifications of the VAR model that take the size of trades into account, and to estimating price impacts via robust impulse responses (see, e.g., Jordà (2005)).

6 Trading Protocols and Dealer Profits

While trading in the D2C segment is almost exclusively via RFQ, a number of different trading protocols are used in the interdealer market. Of particular interest are mid-market matching and workup—two size-discovery mechanisms where orders are crossed at a fixed price (see, e.g., Duffie and Zhu (2017)). We use unique order-book data from the main IDB SEF, the GFI Swaps Exchange, to investigate how transaction costs and price impacts vary

[41] This result holds true for reasonable variations in the lag-length of the VAR model. We have also estimated ultra-long VAR models with up to 5000 lags (roughly one month in calendar time) using polynomial distributed lags. These VAR models do indicate a low-frequent transient component in the price impact of D2C trades. However, the associated confidence intervals are so wide that the partial quote reversal is not statistically significant. As discussed in Hasbrouck (1988) and Hasbrouck and Sofianos (1993), it is difficult to statistically detect inventory-control effects that operate at low frequencies because the effects are dwarfed by the arrival of new information.

across trading protocols.[42] We also use this data to estimate dealer profits from liquidity provision to clients.

6.1 Size-Discovery Trading Protocols

In addition to a standard limit order book and RFQ (both of which are price-discovery mechanisms), the GFI Swaps Exchange offers both matching and workup protocols. For five-year on-the-run index CDSs, the predominant matching protocol is continuous mid-market matching.[43] The mid-market level is set by a GFI broker and is usually somewhere between the best bid and offer on the limit order book but does not necessarily have to coincide with the mid-point implied by the best bid and offer. The mid-market level is displayed on the trading screen that shows the limit order book, and the color in which the mid-market level is displayed informs market participants about whether there is interest for matching or not. Market participants are not informed about the direction and size of unfilled interests but they know that interests must be at least of a minimum size.[44] Any opposing interests for matching at the mid-market level immediately result in a trade.

A workup session on the GFI Swaps Exchange is initiated by a trade in the limit order book. During the session, the parties to the initiating trade and other market participants joining the trade can work up the size of the trade by submitting size orders that, in case of a match, result in trades at the transaction price of the initiating trade. The aggressor and liquidity provider of the initiating trade are privileged by means of a 10-second exclusivity period during which they are the only market participants who can work up trade size. The exclusivity period is followed by a public period during which other market participants can

[42]Focusing on trades executed on the GFI Swaps Exchange is not restrictive because it is the IDB SEF facilitating most of the D2D trading volume (see Table 1). Other IDB SEFs also offer matching and workup, but the data are not readily available.

[43]For less frequently traded index CDSs, GFI offers periodic matching sessions which are briefly discussed in the Internet Appendix.

[44]Current minimum sizes are USD 25 million for CDX.IG and USD 10 million for CDX.HY.

join the trade. The public period lasts for at least 30 seconds with any workup during the public period automatically extending the workup session. The session terminates 40 seconds after the initiating trade or, in case of being extended, 30 seconds after the last workup. In contrast to continuous mid-market matching, market participants are informed about the direction and size of unfilled interests.

6.2 Data and Identification of Mid-Market Matches and Workups

The GFI data consist of the best bid and offer quotes that rest on the limit order book as well as the mid-market levels that GFI brokers set for mid-market matching. The average number of bid/offer quotes (mid-market levels) per day for CDX.IG and CDX.HY are 1,136 and 872 (130 and 149), respectively. From this data, we identify trades in the limit order book, mid-market matches, and workups (the Internet Appendix contains the details). Trades that are not identified as belonging to any of the three categories are subsumed into their own category. Some of these trades are voice-brokered RFQs.

Table 8 shows the fraction of volume that is executed via the different trading protocols. We separately report volume shares for outright trades in five-year on-the-run index CDSs, other outright trades, index rolls between five-year on-the-run and immediate off-the-run index CDSs, and other package transactions (excluding index swaption and tranche swap delta hedges for which we are unable to identify the trading protocol). For outright trades in five-year on-the-run CDX.IG and CDX.HY, mid-market matches account for 52.2% and 58.6%, respectively, trades in the limit order book account for 19.2% and 15.8%, and the ensuing workups account for 19.9% and 15.6%.[45] Together, size-discovery trading protocols account for more than 70% of trading volume.[46]

[45] About half of the five-year on-the-run trades in the limit order book are subsequently worked up.

[46] The distribution of volume across trading protocols varies by quartiles of the trade size distribution. For block-sized trades, the mode category is for trades with unidentified trading protocol, which includes voice-brokered RFQs. Voice brokers' ability to match dealers with offsetting inventory imbalances may explain why block-sized D2D trades have essentially no price impact (see Table 4).

[Table 8 about here.]

6.3 Transaction Costs Across Trading Protocols

In order to compare effective half-spreads, realized half-spreads, and price impacts across trading protocols, we estimate trade-by-trade regressions similar to those in Equation (2). As before, we focus on outright trades in five-year on-the-run index CDSs and, for comparability with previous results, we continue to compute half-spreads and price impacts with respect to Markit's intraday mid-quote. Specifically, we estimate

$$y_t = \alpha + \beta_{\text{MTCH}}\text{MTCH}_t + \beta_{\text{WRKUP}}\text{WRKUP}_t + \beta_{\text{UNID}}\text{UNID}_t + \gamma' X_t + \epsilon_t, \quad (7)$$

where y_t and X_t are defined as before and MTCH_t, WRKUP_t, and UNID_t are dummy variables for mid-market matches, workups, and trades with unidentified trading protocol. Thus, α estimates the average effective half-spread (or, depending on the dependent variable used, realized half-spread or price impact) of a trade in the limit order book that has a trade size in the first quartile of the trade size distribution and is executed when average market conditions prevail, and βs estimate effective half-spread differences with respect to a trade in the limit order book.

Table 9 displays regression results. First, compare trades in the limit order book with mid-market matches. Effective half-spreads are significantly lower for mid-market matches. This is unsurprising as the mid-market level is usually somewhere between the best bid and offer resting on the limit order book. More importantly, price impacts are significantly lower for mid-market matches, and there are no significant differences in realized half-spreads. That is, we observe a partial segmentation of the order flow, with a higher proportion of uninformed trades being executed via mid-market matching. This is consistent with Zhu's (2014) model of strategic venue selection by informed and liquidity traders. In his model, traders optimally choose between sending orders to a mid-point dark pool (roughly equivalent to continuous mid-market matching) and executing against limit orders. Sending an order to a dark pool involves a trade-off between potential price improvement and the risk of no execution. In equilibrium, liquidity traders prefer the dark pool, while informed traders prefer the certainty of executing against limit orders.

[Table 9 about here.]

Next, compare trades in the limit order book with workups. There are no significant differences in effective half-spreads. This is by design because a workup is executed at the price of the initiating trade in the limit order book. There are also no significant differences in price impacts. Because the duration of a workup session is much shorter than the 15-minute period over which price impact is measured, the price impact of a workup will include most of the price impact of the initiating trade in the limit order book. The result, therefore, indicates that the additional price impact of a workup is close to zero.

To further explore differences in price impacts across trading protocols, we extend the VAR model in Section 5 with separate D2D-trade-related variables for trades in the limit order book, mid-market matches, workups, and the category of trades with unidentified trading protocol. These trade-related variables are defined for trades that are executed on the GFI Swaps Exchange only; other D2D trades are ignored in this analysis. Mid-quote changes and the D2C-trade-related variable are the same as before. We resolve contemporaneous effects between D2D-trade-related variables by allowing mid-market matches to contemporaneously affect the other D2D-trade-related variables, trades in the limit order book to contemporaneously affect workups and trades with unidentified trading protocol, and workups to contemporaneously affect trades with unidentified trading protocol. With the exception of the relation between trades in the limit order book and workups, for which the ordering is determined by the design of the workup trading protocol, the causal ordering imposed among the remaining pairs of D2D-trade-related variables is inconsequential because their contemporaneous correlations are negligible.

Figure 4 shows cumulative quote revisions in response to buyer-initiated trades. A trade in the limit order book has the highest price impact; indeed, the price impact is similar to that of a D2C trade. A mid-market match has significantly lower price impact than a trade in the limit order book consistent with Table 9. Finally, the VAR model separates the price impact of a workup from that of the initiating trade in the limit order book, and the figure shows that a workup has very low price impact (essentially zero in case of CDX.IG). That a workup has significantly lower price impact than the initiating trade in the limit order book is consistent with evidence from the interdealer Treasury market reported in Fleming and Nguyen (2015).

[Figure 4 about here.]

Overall, our results show that the average transaction cost and price impact of D2D trades reported in Sections 4 and 5 mask significant heterogeneity across trading protocols, with size-discovery protocols attracting liquidity-motivated trading.

6.4 Estimates of Dealer Profits from Liquidity Provision

We use the GFI data to estimate dealer profits from liquidity provision to clients in five-year on-the-run index CDSs. Specifically, we assume that dealers provide immediacy on D2C SEFs and close their positions on the GFI Swaps Exchange. For each index we compute, day by day, the trade-size-weighted average profits from all D2C trades and multiply them by the aggregate trading volumes on D2C SEFs (from Clarus FT). Our estimates of profits from liquidity provision are sample means of daily profits computed in this way. In computing per trade profits, we consider two scenarios: first, that liquidity providers are able to immediately close D2C trades at the mid-market level that prevails at trade execution. Second, that liquidity providers are able to immediately close protection-buyer-initiated (protection-seller-initiated) D2C trades at the best offer (bid) that prevails at trade execution on the order book.[47]

In the first scenario, estimated daily profits are USD 0.433 million for CDX.IG and USD 0.808 million for CDX.HY, or USD 1.241 million in total. However, this presumes that the quoted mid-market level is executable, which is only the case if there are opposing interests for matching. In the second scenario that uses executable bid and offer quotes, estimated profits are negative.[48] This suggests that dealers only make profits through their willingness to bear inventory risk (see, e.g., Grossman and Miller (1988)).

The results show that clients who value immediacy would typically not be able to reduce transaction costs by executing trades on the limit order book of the GFI Swaps Exchange. Indeed, 96.0% and 96.6% of the D2C trades in CDX.IG and CDX.HY, respectively, are executed at prices that are strictly more favorable than the contemporaneous best bid or offer.[49] Transaction costs can only be reduced at the cost of execution risk either through limit orders or through mid-market matching.

[47] We require mid-market levels and quotes to come from within 15 minutes prior to trade execution. Therefore, per trade profits cannot be computed for a few trades and we drop these trades from the computation of daily trade-size-weighted profits. Similarly, when assuming that trades are closed at the best bid or offer, we drop trades for which the side of the order book at which the trade would be closed is empty at trade execution.

[48] Trades can be closed at the prevailing best bid or offer provided that there is sufficient depth. We abstract from this issue when computing per trade profits because GFI data does not include the depth available at the best bid and offer.

[49] In the computation of fractions, trades are signed based on Markit intraday mid-quotes. A more robust approach is to consider only D2C trades for which, based on the latest order-book quote from within 15 minutes prior to trade execution, neither side of the order book is empty at trade execution and report the fraction of D2C trades with transaction prices that are strictly within the bid-offer spread. The corresponding fractions are 95.7% and 96.4% for CDX.IG and CDX.HY, respectively.

7 Conclusion

Using transaction data, we study the market structure and transaction costs of index CDSs after the implementation of the Dodd-Frank Act. The market exhibits a two-tiered structure, and we identify dealer-to-client (D2C) trades and interdealer (D2D) trades. Average transaction costs and price impact are higher for D2C trades and increase with trade size, quoted bid-ask spread, and volatility. Price impact is largely permanent; the relatively high permanent price impact of D2C trades is suggestive of clients trading on information. D2C trades Granger-cause D2D trades consistent with the interdealer market being used for managing inventory risk. Unique order-book data show that D2D transaction costs and price impacts vary across trading protocols with mid-market matching and workup attracting liquidity-motivated trades. D2C prices are typically better than those available on the main interdealer limit order book, indicating that clients who value immediacy could not get better execution by trading in the interdealer market. This may explain the endurance of the two-tiered market structure.

Appendices

A Dodd-Frank Act Implementation Timeline

Jul 21, 2010	President Obama signs the Dodd-Frank Wall Street Reform and Consumer Protection Act (the "Dodd-Frank Act") into law.
Jan 9, 2012	The CFTC publishes the final rules for real-time public reporting of swap transaction data.
Nov 28, 2012	The CFTC announces mandatory central clearing of certain swaps in three implementation phases. In the first phase, swap dealers and private funds active in swap markets (so-called Category 1 Entities) are required to clear their swap trades. In the second phase, financial entities other than Category 3 Entities (so-called Category 2 Entities) are required to clear their swap trades. In the third phase, investment managers and pension plans (so-called Category 3 Entities) are required to clear their swap trades. End-users, i.e., non-financial entities hedging commercial risks, are exempt from mandatory central clearing.
Dec 31, 2012	Real-time public reporting of index CDS trades becomes mandatory for swap dealers.
Feb 28, 2014	Real-time public reporting of index CDS trades becomes mandatory for major swap market participants.
Mar 11, 2013	Central clearing becomes mandatory for Category 1 Entities trading CDX.IG or CDX.HY (for trades in the five-year contract tenor, mandatory central clearing applies to series 11 and all subsequent series).
Apr 10, 2013	Real-time public reporting of index CDS trades becomes mandatory for any swap market participant.
May 31, 2013	The CFTC publishes the final block-trade rules.[50]
Jun 4, 2013	The CFTC publishes the final rules for SEF compliance and mandatory trade execution on SEFs. These specify: (i) the (electronic) trading platforms that are required to be registered as SEFs and the prescribed methods of execution for trades in swaps that are subject to mandatory trade execution on SEFs (either execution against orders resting on a SEF's order book or execution against a response to a RFQ facilitated by a SEF and transmitted to at least three other SEF participants) and (ii) the process that SEFs can initiate (via so-called made available to trade determinations demonstrating sufficiently liquid trading) to get CFTC approval for mandatory trade execution of certain swaps on SEFs.[51]

[50] Block trades are exempt from the trade execution requirement and may be publicly disseminated with delay.
[51] Swaps eligible for made available to trade determinations have to be subject to mandatory central clearing.

Jun 10, 2013	Central clearing becomes mandatory for Category 2 Entities trading CDX.IG or CDX.HY.
Jul 30, 2013	Block trade rules become effective, with index CDS trades of notional amounts exceeding certain par-spread- and contract-tenor-dependent minimum block sizes being defined as block trades (note that minimum block sizes defining block trades do not necessarily coincide with the sizes at which publicly disseminated notional amounts are being capped).[52]
Aug 5, 2013	Closing date for applications to become a CFTC-registered SEF according to (i) from above. Temporarily registered SEFs are free to initiate made available to trade determinations that are subject to CFTC approval as set forth in (ii) from above.
Sep 9, 2013	Central clearing becomes mandatory for Category 3 Entities trading CDX.IG or CDX.HY.
Oct 2, 2013	The first temporarily registered SEFs start operating.
Jan 28, 2014	The CFTC approves a made available to trade determination for on-the-run and immediate off-the-run index CDSs on CDX.IG, CDX.HY, iTraxx Europe, and iTraxx Europe Crossover with five-year contract tenors.
Feb 26, 2014	The approved made available to trade determination becomes effective and all trades in the above-mentioned index CDSs (not qualifying as block trades or not being end-user exempt) must be executed on SEFs.

[52] Prior to July 30, 2013, all index CDS trades were publicly disseminated with delay and for trades of notional amounts exceeding USD 100 million, the disseminated notional amounts were capped at USD 100 million.

References

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

Atkeson, Andrew G., Andrea L. Eisfeldt, and Pierre-Olivier Weill, 2013, The market for OTC derivatives, Working paper, UCLA.

Babus, Ana, and Cecilia Parlatore, 2016, Strategic fragmented markets, Working paper, New York University.

Barclay, Michael J., Terrence Hendershott, and Kenneth Kotz, 2006, Automation versus intermediation: Evidence from Treasuries going off the run, *Journal of Finance* 61, 2395–2414.

Benos, Evangelos, Richard Payne, and Michalis Vasios, 2016, Centralized trading, transparency and interest rate swap market liquidity: Evidence from the implementation of the Dodd-Frank Act, Staff report, Bank of England.

Bessembinder, Hendrik, 2003, Selection biases and cross-market trading cost comparisons, Working paper, University of Utah.

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.

Biswas, Gopa, Stanislava Nikolova, and Christof W. Stahel, 2015, The transaction costs of trading corporate credit, Working paper, University of Nebraska–Lincoln.

Comerton-Forde, Carole, and Tālis J. Putniņš, 2015, Dark trading and price discovery, *Journal of Financial Economics* 118, 70–92.

Duffie, Darrell, and Haoxiang Zhu, 2017, Size discovery, *Review of Financial Studies* 30, 1095–1150.

Dunne, Peter G., Harald Hau, and Michael J. Moore, 2015, Dealer intermediation between markets, *Journal of the European Economic Association* 13, 770–804.

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

Erlandsson, Ulf, Arup Ghosh, and Graham Rennison, 2008, Systematic CDS index trading, Barclays Capital Quantitative Credit Strategy Research.

Fleming, Michael, and Giang Nguyen, 2015, Order flow segmentation and the role of dark pool trading in the price discovery of U.S. treasury securities, Staff report, Federal Reserve Bank of New York.

Giancarlo, J. Christopher, 2015, Pro-reform reconsideration of the CFTC swaps trading rules: Return to Dodd-Frank, White paper, U.S. Commodity Futures Trading Commission.

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

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

Grossman, Sanford J., and Merton H. Miller, 1988, Liquidity and market structure, *Journal of Finance* 43, 617–633.

Harris, Lawrence E., 2015, Transaction costs, trade throughs, and riskless principal trading in corporate bond markets, Working paper, University of Southern California.

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

Hasbrouck, Joel, 1988, Trades, quotes, inventories, and information, *Journal of Financial Economics* 22, 229–252.

———, 1991a, Measuring the information content of stock trades, *Journal of Finance* 46, 179–207.

———, 1991b, The summary informativeness of stock trades: An econometric analysis, *Review of Financial Studies* 4, 571–595.

Hasbrouck, Joel, and Thomas S. Y. Ho, 1987, Order arrival, quote behavior, and the return-generating process, *Journal of Finance* 42, 1035–1048.

Hasbrouck, Joel, and George Sofianos, 1993, The trades of market makers: An empirical analysis of NYSE specialists, *Journal of Finance* 48, 1565–1593.

Hendershott, Terrence, Dmitry Livdan, and Norman Schürhoff, 2015, Are institutions informed about news?, *Journal of Financial Economics* 117, 249–287.

Hendershott, Terrence, and Ananth Madhavan, 2015, Click or call? Auction versus search in the over-the-counter market, *Journal of Finance* 70, 419–447.

Hünseler, Michael, 2013, *Credit portfolio management* (Palgrave Macmillan, London, Greater London).

International Swaps and Derivatives Association, 2014, Dispelling myths: End-user activity in OTC derivatives, Research study.

Ivashina, Victoria, and Zheng Sun, 2011, Institutional stock trading on loan market information, *Journal of Financial Economics* 100, 284–303.

Jordà, Òscar, 2005, Estimation and inference of impulse responses by local projections, *American Economic Review* 95, 161–182.

Lee, Charles M. C., 1993, Market integration and price execution for NYSE-listed securities, *Journal of Finance* 48, 1009–1038.

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

Loon, Yee Cheng, and Zhaodong Ken Zhong, 2016, Does Dodd-Frank affect OTC transaction costs and liquidity? Evidence from real-time CDS trade reports, *Journal of Financial Economics* 119, 645–672.

Managed Funds Association, 2015, Why eliminating post-trade name diclosure will improve the swaps market, Position paper.

Newey, Whitney K., and Kenneth D. West, 1987, A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix, *Econometrica* 55, 703–708.

Reiss, Peter C., and Ingrid M. Werner, 1998, Does risk sharing motivate interdealer trading?, *Journal of Finance* 53, 1657–1703.

Wang, Chaojun, 2016, Core-periphery trading networks, Working paper, Stanford University.

Zhu, Haoxiang, 2014, Do dark pools harm price discovery?, *Review of Financial Studies* 27, 747–789.

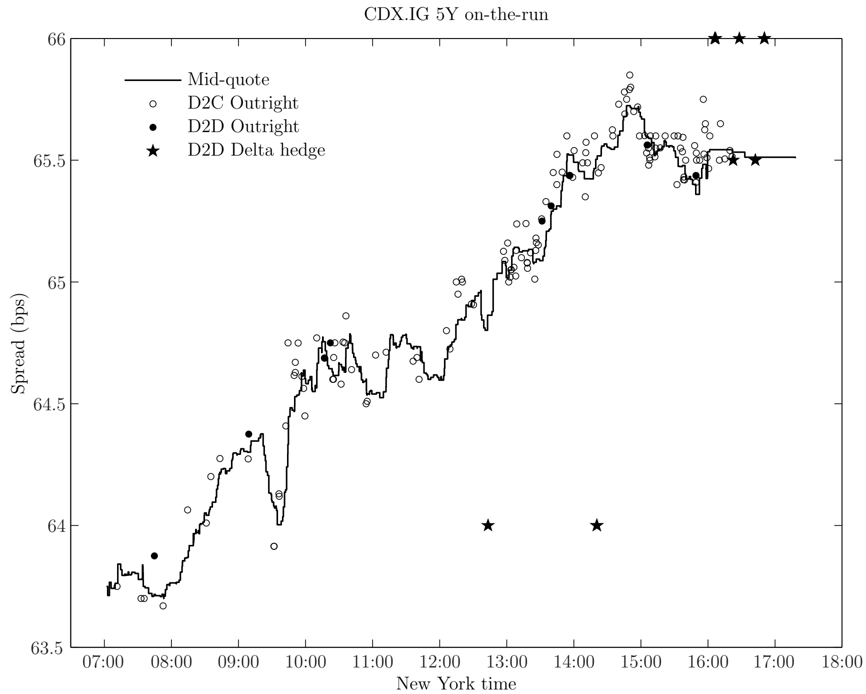

Figure 1: CDX.IG Trades and Quotes on May 6, 2015.

The figure shows transaction prices of all dealer-to-client (D2C) and dealer-to-dealer (D2D) trades in five-year on-the-run index CDSs on CDX.IG and the corresponding composite mid-quote on May 6, 2015. Circles indicate trades that are identified as being outright and stars indicate trades that are identified as being delta hedges of index swaptions. Unfilled symbols indicate D2C trades and filled symbols indicate D2D trades. Both transaction prices and quotes are in terms of par spreads and expressed in basis points (bps). Series 24 of CDX.IG was on-the-run on May 6, 2015.

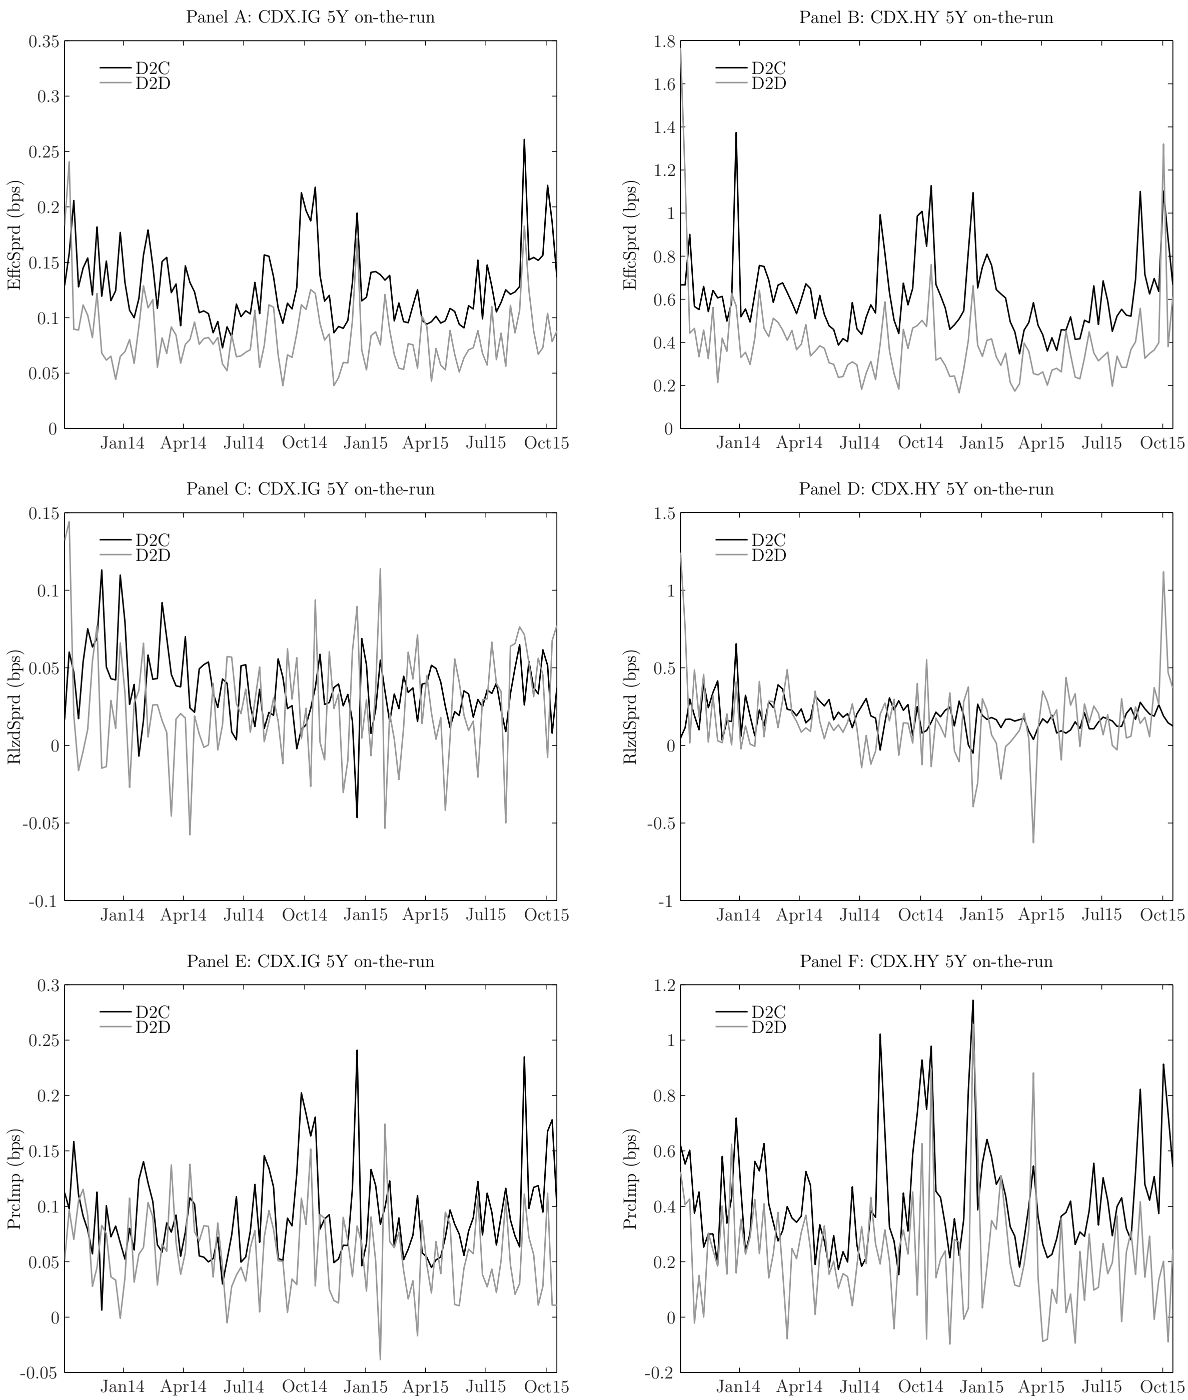

Figure 2: Weekly Average Effective Half-Spreads, Realized Half-Spreads, and Price Impacts. Panels A and B, Panels C and D, and Panels E and F, show weekly sample means of effective half-spreads (EffcSprd), realized half-spreads (RlzdSprd), and price impacts (PrcImp) of outright dealer-to-client (D2C) and dealer-to-dealer (D2D) trades in five-year on-the-run index CDSs on CDX.IG and CDX.HY, respectively. EffcSprd is defined as $q_t \times (p_t - m_t)$, where p_t is the transaction price and m_t is the latest mid-quote in the 15-minute period prior to trade execution. RlzdSprd is defined as $q_t \times (p_t - m_{t+\Delta})$, where $m_{t+\Delta}$ is the first mid-quote in the 15-minute period that follows trade execution by 15 minutes. PrcImp is defined as $q_t \times (m_{t+\Delta} - m_t)$. Both transaction prices and quotes are in terms of par spreads and expressed in basis points (bps). Trade direction, q_t, is inferred by the Lee and Ready (1991) algorithm. The sample period is October 2, 2013 to October 16, 2015 and comprises 50,126 (8,881) and 71,697 (10,219) outright D2C (D2D) trades in five-year on-the-run index CDSs on CDX.IG and CDX.HY, respectively.

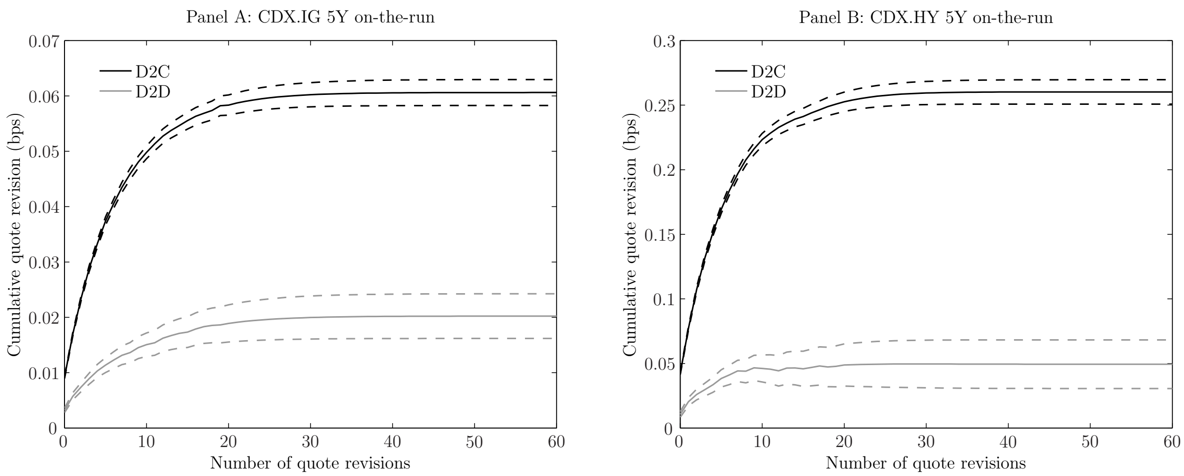

Figure 3: Price Impact.

The panels show cumulative quote revisions in response to either a single protection-buyer-initiated dealer-to-client (D2C; solid black lines) trade or a single protection-buyer-initiated dealer-to-dealer (D2D; solid light gray lines) trade. The trades are outright five-year on-the-run index CDS trades in CDX.IG (Panel A) and CDX.HY (Panel B). Cumulative quote revisions are implied by event-time vector autoregressive models for mid-quote revisions, the sum of signed D2C trades that occur between quote revisions, and the sum of signed D2D trades that occur between quote revisions. Dashed lines mark 95% confidence intervals based on OLS standard errors. Quotes are in terms of par spreads and expressed in basis points (bps). The sample period is October 2, 2013 to October 16, 2015 and comprises 216,280 and 187,871 quote revisions for CDX.IG and CDX.HY, respectively.

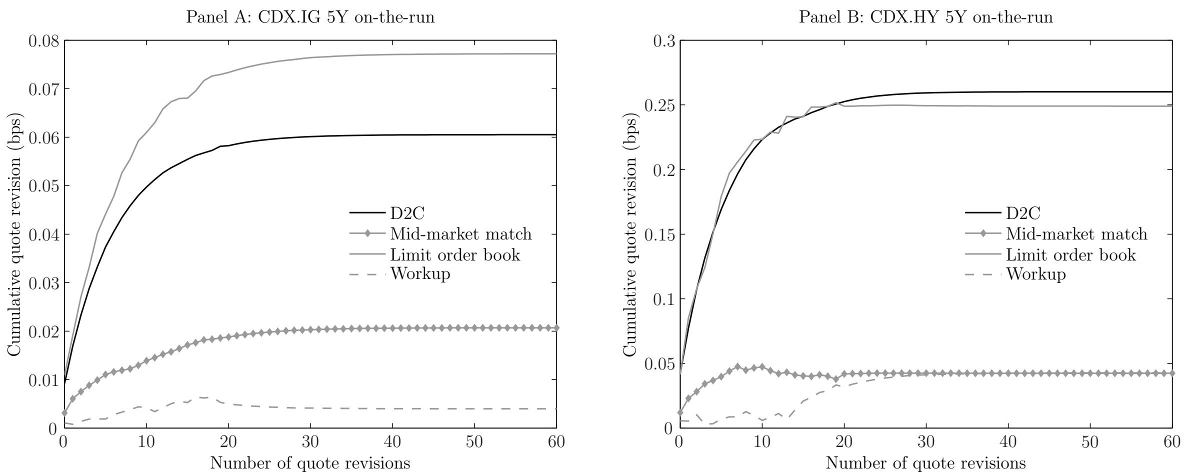

Figure 4: Price Impact.

The panels show cumulative quote revisions in response to either a single protection-buyer-initiated dealer-to-client (D2C) trade (solid black lines) or a single protection-buyer-initiated dealer-to-dealer (D2D) mid-market match (marked solid light gray lines), trade in the limit order book (solid light gray lines), or workup (dashed light gray lines). The trades are outright five-year on-the-run index CDS trades in CDX.IG (Panel A) and CDX.HY (Panel B). Cumulative quote revisions are implied by event-time vector autoregressive models for mid-quote revisions, the sum of signed D2C trades that occur between quote revisions, and sums of signed D2D mid-market matches, trades in the limit order book, workups, and trades with unidentified trading protocol that occur between quote revisions. D2D trades other than those that are executed on the GFI Swaps Exchange are ignored. Quotes are in terms of par spreads and expressed in basis points (bps). The sample period is October 2, 2013 to October 16, 2015 and comprises 216,280 and 187,871 quote revisions for CDX.IG and CDX.HY, respectively.

		USD MM		% of Trds						% of Vlm			
SEF	Trds	Sz	Vlm (ActVlm)	5Y	OTR	Bspk	Clrd	Blck	Cppd	Crv	Rll	Swptn	Trnch
Panel A1: CDX.IG *Dealer-To-Client*													
Bloomberg SEF	95	50	5,394 (7,681)	99.7	93.5	0.0	100.0	20.1	19.4	0.1	4.5	0.0	—
ICE Swap Trade	3	44	154 (318)	98.0	90.1	0.0	84.1	27.2	32.5	0.1	0.0	16.9	6.5
MarketAxess SEF	5	37	281 (533)	99.1	89.8	0.0	100.0	25.0	26.1	0.2	13.1	0.6	—
TW SEF	11	50	605 (1,310)	98.9	84.8	0.0	100.0	28.4	31.5	0.0	7.3	—	—
Total	114	50	6,433 (9,843)	99.6	92.4	0.0	99.6	21.3	21.2	0.1	5.0	0.4	0.2
Panel A2: CDX.IG *Dealer-To-Dealer*													
GFI Swaps Exchange	17	50	773 (808)	96.1	91.6	0.7	99.3	0.0	3.9	4.7	16.9	3.2	1.9
ICAP SEF	1	25	58 (71)	75.8	71.2	0.0	94.9	0.0	9.7	0.0	0.0	7.4	51.8
tpSEF	6	50	351 (445)	94.8	86.9	0.0	96.0	2.2	14.1	4.7	21.9	1.7	0.0
Tradition SEF	0	79	19 (30)	78.7	63.8	0.0	72.4	0.0	27.6	0.0	0.0	61.6	35.0
Total	24	50	1,201 (1,354)	94.8	89.3	0.5	98.0	0.5	6.8	4.4	17.3	3.9	4.3
Panel B1: CDX.HY *Dealer-To-Client*													
Bloomberg SEF	140	10	2,583 (2,840)	100.0	93.9	0.0	100.0	17.8	1.4	0.0	7.8	0.0	—
ICE Swap Trade	3	5	68 (77)	99.7	87.9	0.0	91.0	28.3	7.2	0.0	0.1	9.9	6.5
MarketAxess SEF	6	10	138 (150)	100.0	90.8	0.0	100.0	24.5	3.9	0.0	12.1	0.2	—
TW SEF	15	16	434 (639)	100.0	87.0	0.0	100.0	35.8	9.9	0.0	15.9	—	—
Total	164	10	3,224 (3,705)	100.0	93.0	0.0	99.8	19.9	2.3	0.0	8.9	0.2	0.1
Panel B2: CDX.HY *Dealer-To-Dealer*													
GFI Swaps Exchange	17	10	209 (211)	99.9	94.1	0.7	99.2	0.0	0.9	0.0	21.1	4.2	0.5
ICAP SEF	1	10	17 (25)	98.6	68.9	0.0	87.8	0.0	0.7	0.0	0.0	13.1	45.8
tpSEF	8	10	147 (157)	99.6	89.1	0.0	96.2	2.7	2.6	0.0	26.1	1.0	0.0
Tradition SEF	0	20	6 (8)	95.0	68.3	0.0	79.2	0.0	1.0	0.0	0.0	49.6	36.8
Total	27	10	380 (402)	99.7	91.5	0.4	97.8	0.9	1.4	0.0	21.7	4.1	3.0

Table 1: Descriptive Statistics of On-SEF Index CDS Trades.

The table shows descriptive statistics of on-SEF dealer-to-client (D2C) and dealer-to-dealer (D2D) index CDS trades in CDX.IG and CDX.HY by SEF. Trds is the number of trades per day computed as the total number of trades divided by the number of trading days in the sample period, 511. Sz is median trade size. Vlm is daily volume computed as the aggregate notional amount divided by the number of trading days in the sample period (ActVlm is actual daily volume computed equivalently using daily volumes reported by SEFs). 5Y (OTR) is the percentage of trades in five-year (on-the-run) index CDSs. Bspk is the percentage of trades with bespoke contract terms. Clrd is the percentage of cleared trades. Blck is the percentage of trades qualifying as block trades. Cppd is the percentage of trades that are disseminated with capped notional amounts. Crv (Rll) is the percentage of the aggregate notional amount that is identified as being part of curve trades (index rolls). Swptn (Trnch) is the percentage of the aggregate notional amount that is identified as being index swaption (index tranche swap) delta hedges. The sample period is October 2, 2013 to October 16, 2015 and comprises 58,222 (12,396) and 83,771 (13,585) D2C (D2D) trades in CDX.IG and CDX.HY, respectively.

	% of Trds		% of Vlm	
Trade Type	D2C	D2D	D2C	D2D
Panel A: CDX.IG				
Outright				
5Y on-the-run	90.0	79.9	88.5	67.2
Other	6.0	2.1	5.8	3.0
Package				
Roll 5Y on-the-run/immediate off-the-run	3.4	7.6	4.9	12.0
Other	0.6	10.4	0.8	17.8
Panel B: CDX.HY				
Outright				
5Y on-the-run	91.2	85.5	84.6	63.5
Other	5.3	3.5	6.1	7.7
Package				
Roll 5Y on-the-run/immediate off-the-run	3.2	6.4	8.8	18.2
Other	0.3	4.7	0.5	10.6

Table 2: Percentages of On-SEF Index CDS Trades and Volumes by Trade Type.
Panels A and B show percentages of on-SEF dealer-to-client (D2C) and dealer-to-dealer (D2D) index CDS trades and volumes in CDX.IG and CDX.HY, respectively, by trade type. The sample period is October 2, 2013 to October 16, 2015 and comprises 58,222 (12,396) and 83,771 (13,585) D2C (D2D) trades in CDX.IG and CDX.HY, respectively.

Trade Type	*Dealer-To-Client* Effc Sprd	Rlzd Sprd	Prc Imp	*Dealer-To-Dealer* Effc Sprd	Rlzd Sprd	Prc Imp	*D2C-D2D* Effc Sprd	Rlzd Sprd	Prc Imp
Panel A: CDX.IG									
Outright	0.137	0.031	0.106	0.088	0.025	0.063	0.049**	0.006	0.043**
Index roll	0.048	0.020	0.028	0.050	0.027	0.023	-0.002	-0.007	0.005
Panel B: CDX.HY									
Outright	0.674	0.166	0.508	0.402	0.155	0.246	0.273**	0.011	0.262**
Index roll	0.392	0.239	0.153	0.354	0.131	0.223	0.038	0.108*	-0.070

Table 3: Effective Half-Spreads, Realized Half-Spreads, and Price Impacts by Trade Type. Panels A and B show sample means of effective half-spreads (EffcSprd), realized half-spreads (RlzdSprd), and price impacts (PrcImp) of dealer-to-client (D2C) and dealer-to-dealer (D2D) trades in CDX.IG and CDX.HY, respectively. Sample means are separately computed for outright trades in five-year on-the-run index CDSs and for index rolls between five-year on-the-run and immediate off-the-run index CDSs. EffcSprd is defined as $q_t \times (p_t - m_t)$, where p_t is the transaction price (the difference between on-the-run and immediate off-the-run transaction prices for index rolls) and m_t is the latest mid-quote (the difference between the latest on-the-run and immediate off-the-run mid-quotes for index rolls) in the 15-minute period prior to trade execution. RlzdSprd is defined as $q_t \times (p_t - m_{t+\Delta})$, where $m_{t+\Delta}$ is the first mid-quote in the 15-minute period that follows trade execution by 15 minutes. PrcImp is defined as $q_t \times (m_{t+\Delta} - m_t)$. Both transaction prices and quotes are in terms of par spreads and expressed in basis points. Trade direction, q_t, is inferred by the Lee and Ready (1991) algorithm. ** and * denote rejection of a regression-based t test for the null hypothesis that D2C and D2D sample means are identical at the 1% and 5% level, respectively, with inference based on the Newey and West (1987) estimate of the covariance matrix of coefficient estimates. The sample period is October 2, 2013 to October 16, 2015 and comprises 50,126 (8,881) and 71,697 (10,219) outright D2C (D2D) trades in five-year on-the-run index CDSs on CDX.IG and CDX.HY, respectively and 943 (338) and 1,094 (329) D2C (D2D) index rolls between five-year on-the-run and immediate off-the-run index CDSs on CDX.IG and CDX.HY, respectively.

	Dealer-To-Client			*Dealer-To-Dealer*			*D2C-D2D*		
Trade Size	Effc Sprd	Rlzd Sprd	Prc Imp	Effc Sprd	Rlzd Sprd	Prc Imp	Effc Sprd	Rlzd Sprd	Prc Imp
				Panel A: CDX.IG					
≤ 25	0.121	0.031	0.090	0.082	0.017	0.065	0.039**	0.015**	0.025**
25–50	0.131	0.022	0.109	0.095	0.022	0.073	0.036**	0.000	0.036**
50–100	0.143	0.022	0.121	0.090	0.053	0.037	0.053**	-0.031**	0.084**
> 100	0.169	0.051	0.118	0.125	0.153	-0.028	0.044**	-0.102**	0.146**
				Panel B: CDX.HY					
≤ 5	0.603	0.169	0.434	0.383	0.108	0.275	0.220**	0.061	0.159**
5–10	0.636	0.120	0.516	0.413	0.154	0.259	0.223**	-0.034	0.257**
10–25	0.700	0.118	0.582	0.394	0.204	0.190	0.306**	-0.086*	0.392**
> 25	0.800	0.287	0.513	0.468	0.478	-0.011	0.332**	-0.191	0.523**

Table 4: Effective Half-Spreads, Realized Half-Spreads, and Price Impacts by Trade Size. Panels A and B show sample means of effective half-spreads (EffcSprd), realized half-spreads (RlzdSprd), and price impacts (PrcImp) of outright dealer-to-client (D2C) and dealer-to-dealer (D2D) trades in five-year on-the-run index CDSs on CDX.IG and CDX.HY, respectively. Sample means are separately computed for quartiles of the trade size distribution. EffcSprd is defined as $q_t \times (p_t - m_t)$, where p_t is the transaction price and m_t is the latest mid-quote in the 15-minute period prior to trade execution. RlzdSprd is defined as $q_t \times (p_t - m_{t+\Delta})$, where $m_{t+\Delta}$ is the first mid-quote in the 15-minute period that follows trade execution by 15 minutes. PrcImp is defined as $q_t \times (m_{t+\Delta} - m_t)$. Both transaction prices and quotes are in terms of par spreads and expressed in basis points. Trade size is in USD million. Trade direction, q_t, is inferred by the Lee and Ready (1991) algorithm. ** and * denote rejection of a regression-based t test for the null hypothesis that D2C and D2D sample means are identical at the 1% and 5% level, respectively, with inference based on the Newey and West (1987) estimate of the covariance matrix of coefficient estimates. The sample period is October 2, 2013 to October 16, 2015 and comprises 50,126 (8,881) and 71,697 (10,219) outright D2C (D2D) trades in five-year on-the-run index CDSs on CDX.IG and CDX.HY, respectively.

	CDX.IG			CDX.HY		
	EffcSprd	RlzdSprd	PrcImp	EffcSprd	RlzdSprd	PrcImp
D2C	0.121**	0.027**	0.093**	0.609**	0.154**	0.455**
	(67.10)	(11.84)	(29.95)	(74.40)	(13.63)	(30.27)
D2D	0.087**	0.021**	0.066**	0.390**	0.153**	0.236**
	(31.66)	(5.16)	(14.51)	(28.62)	(6.00)	(9.48)
MDM	0.008**	-0.007**	0.015**	0.014*	-0.040**	0.054**
	(5.45)	(-2.76)	(5.24)	(1.98)	(-3.40)	(4.51)
LRG	0.015**	-0.004	0.020**	0.062**	-0.041**	0.103**
	(8.39)	(-1.39)	(6.12)	(8.16)	(-3.36)	(7.65)
BLCK	0.044**	0.023**	0.020**	0.188**	0.122**	0.067**
	(17.70)	(6.33)	(4.84)	(19.72)	(7.13)	(3.91)
RFRNC	0.021**	0.023**	-0.002	0.111**	0.178**	-0.067**
	(8.20)	(4.98)	(-0.44)	(6.13)	(6.17)	(-2.61)
BAS	0.445**	0.034	0.410**	0.345**	0.069	0.276**
	(8.22)	(0.53)	(4.23)	(10.91)	(1.38)	(4.58)
SPRD/100	0.022	0.089*	-0.067	0.066*	0.015	0.051
	(0.61)	(2.05)	(-1.00)	(2.03)	(0.43)	(0.90)
VLTLTY	0.199**	-0.166**	0.365**	1.220**	-0.503*	1.722**
	(5.94)	(-3.83)	(5.72)	(7.34)	(-2.39)	(5.36)
N	59,007	59,007	59,007	81,916	81,916	81,916
D2C − D2D	0.033	0.006	0.027	0.219	0.001	0.218
p-value	<0.01	0.12	<0.01	<0.01	0.96	<0.01

Table 5: Regressions Controlling for Outright Trade Characteristics and Market Conditions. The table shows OLS estimates of regression specifications that control for selection bias in the comparison of effective half-spreads (EffcSprd), realized half-spreads (RlzdSprd), and price impacts (PrcImp) of outright dealer-to-client (D2C) and dealer-to-dealer (D2D) trades in five-year on-the-run index CDSs on CDX.IG and CDX.HY (t-statistics based on Newey and West (1987) standard errors are shown in parenthesis). EffcSprd is defined as $q_t \times (p_t - m_t)$, where p_t is the transaction price and m_t is the latest mid-quote in the 15-minute period prior to trade execution. RlzdSprd is defined as $q_t \times (p_t - m_{t+\Delta})$, where $m_{t+\Delta}$ is the first mid-quote in the 15-minute period that follows trade execution by 15 minutes. PrcImp is defined as $q_t \times (m_{t+\Delta} - m_t)$. Both transaction prices and quotes are in terms of par spreads and expressed in basis points (bps). Trade direction, q_t, is inferred by the Lee and Ready (1991) algorithm. The explanatory variables include dummy variables for D2C trades (D2C), for D2D trades (D2D), for medium-sized trades (MDM; USD 25–50 MM for CDX.IG and USD 5–10 MM for CDX.HY), for large-sized trades (LRG; USD 50–100 MM for CDX.IG and USD 10–25 MM for CDX.HY), for block-sized trades (BLCK; +USD 100 MM for CDX.IG and +USD 25 MM for CDX.HY), and for trades with transaction prices at typical reference levels (RFRNC; par spread multiples 0.5 bps for CDX.IG and price multiples of 0.125% for CDX.HY), the bid-ask spread of the latest quote for the five-year on-the-run index CDS (BAS), the corresponding mid-quote (SPRD), and the implied volatility of three-month at-the-money swaptions on the five-year on-the-run index CDS (VLTLTY). Continuous explanatory variables are demeaned. The prior to last row shows the difference between D2C and D2D coefficient estimates and the last row shows the p-value of a Wald test for the null hypothesis that D2C and D2D coefficients are identical. ** and * denote statistical significance at the 1% and 5% level, respectively. The sample period is October 2, 2013 to October 16, 2015 and comprises 50,126 (8,881) and 71,697 (10,219) outright D2C (D2D) trades in five-year on-the-run index CDSs on CDX.IG and CDX.HY, respectively.

Matching Trade Size	*Dealer-To-Client* Effc Sprd	Rlzd Sprd	Prc Imp	*Dealer-To-Dealer* Effc Sprd	Rlzd Sprd	Prc Imp	*D2C-D2D* Effc Sprd	Rlzd Sprd	Prc Imp
Panel A: CDX.IG									
≤ 25	0.117	0.036	0.080	0.086	0.017	0.069	0.031**	0.019*	0.011
25–50	0.125	0.029	0.096	0.102	0.016	0.086	0.022**	0.013	0.010
50–100	0.124	0.010	0.114	0.099	0.055	0.044	0.025**	-0.045*	0.070**
> 100	0.153	0.097	0.056	0.114	0.163	-0.049	0.039	-0.066	0.105**
Exact	0.124	0.026	0.098	0.097	0.019	0.078	0.027**	0.007	0.019*
Panel B: CDX.HY									
≤ 5	0.575	0.154	0.420	0.384	0.071	0.313	0.191**	0.083	0.108*
5–10	0.580	0.117	0.464	0.448	0.179	0.269	0.132**	-0.063	0.195**
10–25	0.621	0.137	0.484	0.412	0.212	0.200	0.210**	-0.075	0.284**
> 25	0.690	0.140	0.550	0.377	0.392	-0.016	0.313**	-0.253	0.566**
Exact	0.596	0.109	0.488	0.432	0.149	0.283	0.164**	-0.041	0.205**

Table 6: Effective Half-Spreads, Realized Half-Spreads, and Price Impacts of Matched Pairs. Panels A and B show sample means of effective half-spreads (EffcSprd), realized half-spreads (RlzdSprd), and price impacts (PrcImp) of matched pairs of outright dealer-to-client (D2C) and dealer-to-dealer (D2D) trades in five-year on-the-run index CDSs on CDX.IG and CDX.HY, respectively. Sample means are separately computed for quartiles of the trade size distribution. EffcSprd is defined as $q_t \times (p_t - m_t)$, where p_t is the transaction price and m_t is the latest mid-quote in the 15-minute period prior to trade execution. RlzdSprd is defined as $q_t \times (p_t - m_{t+\Delta})$, where $m_{t+\Delta}$ is the first mid-quote in the 15-minute period that follows trade execution by 15 minutes. PrcImp is defined as $q_t \times (m_{t+\Delta} - m_t)$. Both transaction prices and quotes are in terms of par spreads and expressed in basis points. Trade size is in USD million. Trade direction, q_t, is inferred by the Lee and Ready (1991) algorithm. A pair consists of a D2D trade and matching D2C trade in the same index CDS and with trade size in the same quartile of the trade size distribution (or with exactly matching trade size) that occur within a 15-minute window bracketing the execution of the D2D trade. In case of more than one matching D2C trade, the EffcSprd, RlzdSprd, and PrcImp of the D2C trade of the pair are averages of the matching D2C trades. ** and * denote rejection of a regression-based t test for the null hypothesis that the mean of the distribution of paired differences is zero at the 1% and 5% level, respectively, with inference based on the Newey and West (1987) estimate of the covariance matrix of coefficient estimates. The sample period is October 2, 2013 to October 16, 2015 and comprises 4,683 (3,372) and 6,463 (5,115) (exactly) matched pairs of outright D2C and D2D trades in five-year on-the-run index CDSs on CDX.IG and CDX.HY, respectively.

	Coefficient Estimates					Granger Causality Tests		
	$\sum_{j=1}^{20} \Delta m_{t-j}$	x_t^{D2C}	$\sum_{j=1}^{20} x_{t-j}^{\mathrm{D2C}}$	x_t^{D2D}	$\sum_{j=1}^{20} x_{t-j}^{\mathrm{D2D}}$	Δm	x^{D2C}	x^{D2D}
				Panel A1: CDX.IG				
Δm_t	0.351	0.009	0.016	0.003	0.006		4802.0	141.4
	(51.79)	(77.39)	(34.31)	(14.81)	(6.71)		[<0.01]	[<0.01]
x_t^{D2C}	1.742		0.266		0.035	1106.8		19.2
	(14.07)		(30.40)		(2.24)	[<0.01]		[0.51]
x_t^{D2D}	-0.222	0.023	0.030		0.146	35.3	113.4	
	(-3.29)	(19.90)	(6.34)		(17.23)	[0.02]	[<0.01]	
				Panel A2: CDX.HY				
Δm_t	0.262	0.042	0.074	0.011	0.022		5315.8	95.0
	(32.66)	(81.90)	(36.48)	(9.53)	(5.02)		[<0.01]	[<0.01]
x_t^{D2C}	0.321		0.338		-0.019	880.7		21.2
	(8.92)		(37.24)		(-0.94)	[<0.01]		[0.38]
x_t^{D2D}	-0.075	0.023	0.029		0.130	56.1	127.1	
	(-4.49)	(21.23)	(6.84)		(14.18)	[<0.01]	[<0.01]	

	Panel B: Permanent Price Impact					
	CDX.IG			CDX.HY		
	D2C	D2D	D2C − D2D	D2C	D2D	D2C − D2D
Λ	0.061	0.020	0.040	0.260	0.049	0.211
	(50.57)	(9.83)	(17.17)	(53.99)	(5.15)	(19.68)

	Panel C: Price Discovery					
	CDX.IG			CDX.HY		
	D2C	D2D	Trade-Unrelated	D2C	D2D	Trade-Unrelated
R^2	29.57	0.97	69.46	38.89	0.30	60.81

Table 7: VAR Estimates.

The table shows coefficient estimates of event-time vector autoregressive (VAR) models for mid-quote revisions (Δm), the sum of signed dealer-to-client (D2C) trades that occur between quote revisions (x^{D2C}), and the sum of signed dealer-to-dealer (D2D) trades that occur between quote revisions (x^{D2D}). Panels A1 and A2 show VAR coefficient estimates (t-statistics based on OLS standard errors are shown in parenthesis) and Wald test statistics (p-values are shown in brackets) for the null hypothesis that the column variable does not Granger-cause the row variable. Coefficient estimates of contemporaneous variables are separated from coefficient estimates of lagged variables and sums of the latter are reported in columns that show sums of lagged variables. Panel B shows permanent price impact estimates (Λ; t-statistics based on OLS standard errors are shown in parenthesis) as captured by the model-implied long-run cumulative quote revision (in basis points) in response to either a single protection-buyer-initiated D2C trade or a single protection-buyer-initiated D2D trade, as well as the difference in price impacts of D2C and D2D trades. Panel C shows a model-implied variance decomposition of efficient price innovations into trade-related and trade-unrelated components (in percent of the variance of efficient price innovations). Quotes are in terms of par spreads and trade direction used to sign trades is inferred by the Lee and Ready (1991) algorithm. The sample period is October 2, 2013 to October 16, 2015 and comprises 216,280 and 187,871 quote revisions for CDX.IG and CDX.HY, respectively.

	Mid-Market Matching	Limit Order Book	Workup	Unidentified Protocol
Panel A: CDX.IG				
Outright				
5Y on-the-run	52.2	19.2	19.9	8.8
Other	13.6	8.8	18.4	59.1
Package				
Roll 5Y on-the-run/immediate off-the-run	20.1	17.2	27.9	34.8
Other	7.9	23.7	27.2	41.1
Total	42.6	19.4	21.6	16.4
Panel B: CDX.HY				
Outright				
5Y on-the-run	58.6	15.8	15.6	10.0
Other	6.5	7.3	14.9	71.3
Package				
Roll 5Y on-the-run/immediate off-the-run	31.2	19.2	16.1	33.6
Other	5.1	24.1	28.4	42.4
Total	49.3	16.5	16.2	18.0

Table 8: GFI Swaps Exchange Volume Shares by Trading Protocol.

Panels A and B show percentages of GFI Swaps Exchange trading volumes of index CDS trades in CDX.IG and CDX.HY, respectively, by trading protocol. Row values add to 100% and delta hedges of index swaption and tranche swap trades are excluded from the computation. The sample period is October 2, 2013 to October 16, 2015 and comprises 8,253 and 8,199 (non-delta-hedge) trades for CDX.IG and CDX.HY, respectively.

	CDX.IG			CDX.HY		
	EffcSprd	RlzdSprd	PrcImp	EffcSprd	RlzdSprd	PrcImp
CONST	0.102**	-0.014*	0.116**	0.479**	-0.049	0.528**
	(30.09)	(-1.99)	(15.56)	(12.80)	(-0.84)	(9.79)
MTCH	-0.041**	0.004	-0.045**	-0.155**	0.014	-0.168**
	(-9.94)	(0.59)	(-5.83)	(-5.08)	(0.26)	(-3.42)
WRKUP	0.004	0.003	0.001	0.011	-0.020	0.030
	(0.95)	(0.35)	(0.15)	(0.55)	(-0.50)	(0.79)
UNID	0.039*	0.098**	-0.059**	0.351*	0.581**	-0.230**
	(2.00)	(3.38)	(-3.55)	(2.11)	(3.26)	(-3.34)
MDM	0.010**	0.008	0.002	0.001	0.050	-0.048
	(3.20)	(1.32)	(0.25)	(0.05)	(1.13)	(-1.25)
LRG	0.007	0.019	-0.012	0.002	0.033	-0.030
	(1.11)	(1.95)	(-1.25)	(0.10)	(0.66)	(-0.60)
BLCK	-0.001	0.063*	-0.064**	-0.025	0.346	-0.371*
	(-0.08)	(2.23)	(-2.73)	(-0.34)	(1.95)	(-2.43)
RFRNC	0.015**	0.003	0.012	0.060	0.092	-0.032
	(2.90)	(0.22)	(1.02)	(1.33)	(1.29)	(-0.54)
BAS	0.312**	0.151	0.161	0.216**	0.079	0.137
	(4.24)	(1.50)	(1.21)	(4.96)	(0.87)	(1.51)
SPRD/100	0.119	-0.020	0.140	0.306	0.378	-0.072
	(1.95)	(-0.19)	(1.95)	(1.63)	(1.73)	(-0.78)
VLTLTY	0.139**	-0.017	0.156	0.485	-1.346*	1.831**
	(3.22)	(-0.20)	(1.58)	(1.43)	(-1.98)	(2.97)
N	6,623	6,623	6,623	6,844	6,844	6,844

Table 9: Regressions Controlling for D2D Trade Characteristics and Market Conditions.
The table shows OLS estimates of regression specifications that control for selection bias in the comparison of effective half-spreads (EffcSprd), realized half-spreads (RlzdSprd), and price impacts (PrcImp) of order-book trades, mid-market matches, workups, and trades with unidentified trading protocol in five-year on-the-run index CDSs on CDX.IG and CDX.HY (t-statistics based on Newey and West (1987) standard errors are shown in parenthesis). EffcSprd is defined as $q_t \times (p_t - m_t)$, where p_t is the transaction price and m_t is the latest mid-quote in the 15-minute period prior to trade execution. RlzdSprd is defined as $q_t \times (p_t - m_{t+\Delta})$, where $m_{t+\Delta}$ is the first mid-quote in the 15-minute period that follows trade execution by 15 minutes. PrcImp is defined as $q_t \times (m_{t+\Delta} - m_t)$. Both transaction prices and quotes are in terms of par spreads and expressed in basis points (bps). Trade direction, q_t, is inferred by the Lee and Ready (1991) algorithm. The explanatory variables include a constant (CONST), dummy variables for mid-market matches (MTCH), for workups (WRKUP), for trades with unidentified trading protocol (UNID), for medium-sized trades (MDM; USD 25–50 MM for CDX.IG and USD 5–10 MM for CDX.HY), for large-sized trades (LRG; USD 50–100 MM for CDX.IG and USD 10–25 MM for CDX.HY), for block-sized trades (BLCK; +USD 100 MM for CDX.IG and +USD 25 MM for CDX.HY), and for trades with transaction prices at typical reference levels (RFRNC; par spread multiples 0.5 bps for CDX.IG and price multiples of 0.125% for CDX.HY), the bid-ask spread of the latest Markit intraday quote for the five-year on-the-run index CDS (BAS), the corresponding mid-quote (SPRD), and the implied volatility of three-month at-the-money swaptions on the five-year on-the-run index CDS (VLTLTY). Continuous explanatory variables are demeaned. ** and * denote statistical significance at the 1% and 5% level, respectively. The sample period is October 2, 2013 to October 16, 2015 and comprises 1,280 (3,640) [1,243] {460} and 1,050 (4,261) [1,084] {449} outright trades in the limit order book (mid-market matches) [workups] {trades with unidentified trading protocol} in five-year on-the-run index CDSs on CDX.IG and CDX.HY, respectively.