

Essays on the Market Structure and Pricing of Credit Derivatives

THÈSE N° 7322 (2016)

PRÉSENTÉE LE 4 NOVEMBRE 2016

AU COLLÈGE DU MANAGEMENT DE LA TECHNOLOGIE
CHAIRE DU PROF. ASSISTANT TENURE-TRACK TROLLE
PROGRAMME DOCTORAL EN FINANCE

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

POUR L'OBTENTION DU GRADE DE DOCTEUR ÈS SCIENCES

PAR

Jan Benjamin JUNGE

acceptée sur proposition du jury:

Prof. L. Mancini, président du jury
Prof. A. Trolle, directeur de thèse
Prof. D. Duffie, rapporteur
Prof. D. Lando, rapporteur
Prof. P. Collin-Dufresne, rapporteur



ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Suisse
2016

Acknowledgements

First of all, I would like to thank my supervisor, Anders B. Trolle, for his continued guidance, patience, and support throughout the past five years. I would also like to thank Pierre Collin-Dufresne for his insight, support, and vision, Darrell Duffie and David Lando for their interest in my research and their comments and advice, and Lorian Mancini for charing the thesis jury. Finally, I would like to thank my family and friends for always being there and, in particular, my uncle, Georg Junge, for being role model, mentor, and sparring partner throughout my undergraduate and graduate studies.

Lausanne, September 2016

B. J.

Abstract

This thesis analyzes the interrelation between market structure and price formation in credit derivatives markets. Traditionally, credit derivatives are traded in relatively opaque over-the-counter markets in which trading is segmented and subject to many imperfections from which illiquidity may arise. Recent regulatory reforms have brought transparency to some credit derivatives markets without affecting their segmented structures.

The first chapter, which is joint work with Anders B. Trolle, analyzes whether liquidity risk is priced in the cross section of returns on credit default swaps (CDSs). The analysis is based on a factor pricing model and a tradable liquidity factor that is constructed from returns on index arbitrage strategies. The underlying presumption is that violations of simple no-arbitrage relations between different CDS contracts reflect constraints on the risk-bearing capacity of CDS market intermediaries and, in broad terms, CDS market illiquidity. The analysis reveals priced liquidity risk in that credit protection sellers earn higher expected excess returns on CDS contracts with higher liquidity exposures. The liquidity risk premium is significant and accounts for 24% of CDS spreads, on average. CDS risk premia correlate negatively with proxies for the risk-bearing capacity of CDS market intermediaries, which is consistent with intermediary frictions affecting the pricing of CDSs.

The second chapter, which is joint work with Pierre Collin-Dufresne and Anders B. Trolle, analyzes transaction costs in the dealer-to-customer (D2C) and dealer-to-dealer (D2D) segments of the post-Dodd-Frank index CDS market. Dodd-Frank regulations that made all-to-all trading possible had the potential to break up the market's segmented structure but failed to do so. This led to a controversy with some market participants arguing that the segmented structure is optimal and other market participants arguing that dealers maintain the segmented structure in order to limit competition by alternative liquidity providers. The analysis reveals that D2C trades indeed have larger transaction costs than D2D trades but that the differences in transaction costs reflect differences in price impacts rather than differences in profits from liquidity provision. D2C trades are even competitive relative to executable bids and offers in the D2D segment, suggesting that the market structure delivers favorable prices for customers who value immediacy.

The third chapter documents a decline of transaction costs and profits from liquidity provision in the index CDS market over a two-and-a-half-year period during which Dodd-Frank regulations were implemented. Transaction costs and profits from liquidity provision de-

Abstract

clined around the introduction of so-called swap execution facilities (SEFs); i.e., regulated trading platforms that offer pre-trade transparent methods of trade execution. Trades that are executed on SEFs have lower transaction costs and are less profitable from a liquidity provider's perspective in comparison to bilaterally negotiated trades, which is consistent with better comparison shopping and stronger price competition on SEFs. Dodd-Frank regulations mandating on-SEF trade execution that were implemented after the introduction of SEFs did not affect transaction costs and profits from liquidity provision, suggesting that there is no incremental effect associated with mandatory pre-trade transparency.

Key words: Credit Default Swap; Dodd-Frank Act; Index Credit Default Swap; Liquidity Risk; Over-The-Counter Markets; Swap Execution Facility; Transaction Costs

Zusammenfassung

Thema dieser Dissertation sind Kreditderivate und insbesondere die Beziehung zwischen Preisbildung und Marktstruktur bei deren Bewertung. Kreditderivate werden üblicherweise in außerbörslichen und dezentralisierten Märkten gehandelt, die für Außenstehende nur schwer einsehbar sind. In jüngster Vergangenheit haben Finanzmarktregulierungen zu mehr Transparenz in einigen Märkten für Kreditderivate geführt, jedoch ohne deren dezentralisierte Marktstruktur zu beeinflussen.

Das erste Kapitel, welches auf gemeinsamer Arbeit mit Anders B. Trolle basiert, geht der Fragestellung nach, ob in den Renditen von Credit Default Swaps (CDSs) ein Liquiditätsrisiko eingepreist ist. Die Untersuchung basiert auf einem Faktormodell und einem eigens für die Untersuchung konstruierten Liquiditätsfaktor. Letzterer basiert auf einer Index Arbitrage Strategie. Die hinter der Konstruktion stehende Prämisse ist, dass Arbitragemöglichkeiten dann auftreten, wenn diejenigen Akteure, die für gewöhnlich den Markt machen, nicht genügend Kapital aufbringen können, um selbst von den Arbitragemöglichkeiten zu profitieren und somit generell wenig Liquidität im Markt zur Verfügung steht. Die Untersuchung ergibt, dass Liquiditätsrisiko in den Renditen von CDSs eingepreist ist insofern, dass vom Gesichtspunkt des Verkäufers der Kreditversicherung, CDSs mit höherem Liquiditätsrisiko eine höhere erwartete Rendite haben. Die Risikoprämie ist signifikant und macht im Durchschnitt 24% der CDS Prämie aus.

Das zweite Kapitel, welches auf gemeinsamer Arbeit mit Pierre Collin-Dufresne und Anders B. Trolle basiert, untersucht Transaktionskosten in den zwei Segmenten des amerikanischen Index CDS Marktes, in denen Kreditderivate-Händler mit ihren institutionellen Kunden (D2C) bzw. untereinander (D2D) handeln. Finanzmarktregulierungen, die im Rahmen des Dodd-Frank Act implementiert wurden, haben Marktbedingungen geschaffen, die eine solche, zweigeteilte Marktstruktur nicht unbedingt vorsehen. Da die zweigeteilte Marktstruktur aber weiterhin fortbesteht, entbrannte eine Kontroverse darüber, ob dies die optimale Marktstruktur sei oder ob sie, wie von einigen Marktteilnehmern behauptet, nur daher fortbesteht, weil sich Kreditderivate-Händler einem Wandel zu einer zentralisierten Marktstruktur widersetzen, um die Konkurrenz von alternativen Marktmachern zu unterbinden. Die Untersuchung ergibt, dass Transaktionskosten im D2C Segment höher sind als im D2D Segment was jedoch nicht auf höhere Gewinnmargen von Händlern zurückzuführen ist, sondern darauf, dass D2C Transaktionen einen höheren Preiseffekt haben als D2D Transaktionen.

Zusammenfassung

Das dritte Kapitel geht der Fragestellung nach, warum, wie in diesem Kapitel dokumentiert, Transaktionskosten und Gewinnmargen im amerikanischen Index CDS Markt über einen zweieinhalbjährigen Zeitraum hinweg gefallen sind, der die Implementierung der oben erwähnten Finanzmarktregulierungen einschließt. Ein Grund scheint der Handel auf regulierten Handelsplattformen (SEFs) zu sein, die im Rahmen des Dodd-Frank Act eingeführt wurden. Transaktionen, die auf SEFs ausgeführt wurden, haben geringere Transaktionskosten und Gewinnmargen als Transaktionen, die außerbörslich abgeschlossen wurden. Dies ist konsistent damit, dass SEFs, im Vergleich zu außerbörslichem Handel, bessere Möglichkeiten bieten, die Preise unterschiedlicher Händler miteinander zu vergleichen und somit zu direktem Preiswettbewerb unter Händlern führen.

Stichwörter: Credit Default Swap; Dodd-Frank Act; Index Credit Default Swap; Liquiditätsrisiko; Over-The-Counter Märkte; Swap Execution Facility; Transaktionskosten

Contents

Acknowledgements	i
Abstract (English/Deutsch)	iii
List of figures	xi
List of tables	xiii
1 Liquidity Risk in Credit Default Swap Markets	1
1.1 Introduction	1
1.2 Measuring CDS Market Illiquidity	5
1.2.1 Credit Indices	5
1.2.2 Index Replication	7
1.2.3 Data	8
1.2.4 CDS Market Illiquidity Measure	10
1.2.5 Determinants of CDS Market Illiquidity	12
1.2.6 Tradable Liquidity Factor	15
1.3 Pricing of Liquidity Risk	16
1.3.1 Asset Pricing Model	16
1.3.2 Data and Portfolio Construction	19
1.3.3 Results	21
1.3.4 Robustness Checks	28
1.4 Conclusion	32
2 Market Structure and Transaction Costs of Index CDSs	35
2.1 Introduction	35
2.2 The Index CDS Market	40
2.2.1 Index Credit Default Swaps	40
2.2.2 Pre-Dodd-Frank Market Structure	40
2.2.3 The Dodd-Frank Act and Current Market Structure	41
2.3 Data and Identification Algorithms	42
2.3.1 Data	42
2.3.2 Identification of SEFs	43
2.3.3 Identification of Package Transactions	44
	vii

Contents

2.3.4	SEF Order Flow	45
2.4	Transaction Cost Comparison	49
2.4.1	Transaction Cost Decomposition	49
2.4.2	Descriptive Statistics	51
2.4.3	Accounting for Trade Characteristics and Market Conditions	53
2.5	The Dynamics of Trades and Quotes	58
2.5.1	VAR Framework and Model Estimation	59
2.5.2	Results	61
2.6	Why is Trade in the Interdealer Market Cheaper?	63
2.6.1	Size-Discovery Trading Mechanisms	64
2.6.2	Data and Identification of Mid-Market Matches and Workups	64
2.6.3	Transaction Costs Across Trading Protocols	65
2.6.4	Estimates of Profits from Liquidity Provision	66
2.7	Conclusion	68
3	Index CDS Trading Costs around the Introduction of SEFs	69
3.1	Introduction	69
3.2	The Dodd-Frank Act	73
3.2.1	Minimum Trading Functionality and Trade Execution Requirements	73
3.2.2	Potential Impact of Requirements on Profits from Liquidity Provision	75
3.3	Institutional Details and Data	75
3.3.1	The Index CDS Market	75
3.3.2	Data and Descriptive Statistics	77
3.4	Dodd-Frank Regime Trading Costs	80
3.4.1	Difference-in-Differences Analysis	82
3.4.2	On-SEF and Off-SEF Trading Cost Comparison	88
3.4.3	Trades at Prices Outside the Quoted Bid-Ask Spread	96
3.4.4	Robustness	99
3.5	Conclusion	103
A	Appendix to Chapter 1	105
A.1	Explanatory Variables of CDS Market Illiquidity	105
A.2	Excess Return Computation	106
A.2.1	Realized CDS Excess Return	107
A.2.2	Expected CDS Excess Return	109
A.2.3	Portfolio Excess Returns	109
A.2.4	Realized Credit Index Excess Return	110
A.3	CDS Spread Decomposition	111
A.4	Robustness Checks: Factor Constructions	111
A.5	Standard Error Computation	112
A.5.1	Standard Errors of Factor Price of Risk Estimates	112
A.5.2	Standard Error of the Cross-Sectional R^2	118
A.6	Additional Figures and Tables	120

B Appendix to Chapter 2	125
B.1 Dodd-Frank Act Implementation Timeline	125
B.2 Data Processing	126
B.2.1 On-SEF Trade Report History	127
B.2.2 Identification of SEFs	128
B.2.3 Assessing the Identification Algorithm's Performance	133
B.2.4 Identification of Package Transactions	134
B.2.5 Trade Size Aggregation	138
B.2.6 Cleaning Transaction Prices	139
B.2.7 Inference of Capped Trade Sizes in case of ICETV Trade Reports	141
B.3 Additional Data Sources	142
B.3.1 Clarus FT SEF Volumes	142
B.3.2 Credit Market Analysis Intraday Quotes	143
B.3.3 GFI Market Data	143
B.3.4 Markit Index Swaptions	144
B.3.5 Markit Intraday Quotes	144
B.4 Trading Protocol Identification for GFI Swaps Exchange Trades	145
B.5 Trade Size Weighting	146
B.6 Outright Immediate Off-The-Run Index CDS Trades	147
B.7 Robustness Checks	148
B.7.1 Robustness of Results to Alternative Mid-Quote	150
B.7.2 5- and 30-Minute Realized Half-Spreads and Price Impacts	154
B.7.3 Time Window in Matched Pair Analysis	158
B.8 Standard Error Computation	158
B.8.1 Standard Errors of Cumulative Impulse Responses	158
B.8.2 Standard Error of Price Impact	162
B.9 VAR Models in Trade Size	162
B.10 Additional Figures and Tables	166
 C Appendix to Chapter 3	 169
C.1 Sample Construction	169
 Bibliography	 181
 Curriculum Vitae	 183

List of Figures

1.1	Credit Index Levels, CDS-Implied Index Levels, and Index-to-CDS Bases	9
1.2	CDS Market Illiquidity Measure	12
1.3	Liquidity Factor	17
1.4	Factor Prices and Intermediary Equity Capital	26
1.5	CDS Spread Decomposition of Bid-Ask-Spread-Sorted Portfolios	27
2.1	CDX.IG Trades and Mid-Quotes on May 6, 2015	44
2.2	Average Effective Half-Spreads, Realized Half-Spreads, and Price Impacts	50
2.3	VAR-Model-Implied Price Impact	63
3.1	Effective and Realized Half-Spreads and Fractions of On-SEF Volume	81
3.2	Effective and Realized Half-Spreads by Quartiles of the Trade Size Distribution	95
3.3	Effective and Realized Half-Spreads by Quartiles of the Trade Size Distribution when Excluding Interdealer Trades	102
A.1	Explanatory Variables of CDS Market Illiquidity	121
A.2	CDS Spread Decomposition of Price-Impact-Sorted Portfolios	122
B.1	Non-Block Notional Amount Identified	134
B.2	VAR-Model-Implied Price Impact	166
B.3	Mid-Quote Changes as a Function of Time	167

List of Tables

1.1	Descriptive Statistics of Credit Index Levels and Index-to-CDS Bases	10
1.2	Determinants of CDS Market Illiquidity	14
1.3	Descriptive Statistics of Index Arbitrage Returns	16
1.4	Descriptive Statistics of Factors	21
1.5	Descriptive Statistics of Bid-Ask-Spread-Sorted Portfolios	22
1.6	Results of Time-Series Regressions	23
1.7	Results of Cross-Sectional Regressions	24
1.8	Decompositions of Expected Excess Returns	25
1.9	Robustness of Cross-Sectional Regressions	30
2.2	Percentages of On-SEF Index CDS Trades by Trade Type	48
2.3	Effective Half-Spreads, Realized Half-Spreads, and Price Impacts by Trade Type	52
2.4	Effective Half-Spreads, Realized Half-Spreads, and Price Impacts by Trade Size	53
2.5	Regressions Controlling for Trade Characteristics and Market Conditions	55
2.6	Regressions Controlling for Roll Characteristics and Market Conditions	57
2.7	Effective Half-Spreads, Realized Half-Spreads, and Price Impacts of Pairs	58
2.8	VAR Estimates	62
2.9	GFI Swaps Exchange Volume Shares by Trading Protocol	65
2.10	Regressions Controlling for Trade Characteristics and Market Conditions	67
3.1	Descriptive Statistics Trade and Quote Data	78
3.2	Difference-in-Differences Regression Specifications for Effective Half-Spreads	85
3.3	Difference-in-Differences Regression Specifications for Realized Half-Spreads	87
3.4	Effective and Realized Half-Spreads	89
3.5	Probit Regressions for the Choice of On- and Off-SEF Trade Execution	91
3.6	Effective Half-Spreads in Choice Model for On- and Off-SEF Trade Execution	93
3.7	Regression Specifications for Realized Half-Spreads of On- and Off-SEF Trades	96
3.8	Trades at Prices Outside the Quoted Bid-Ask Spread	97
3.9	Probit Regressions for Trades at Prices Outside the Quoted Bid-Ask Spread	100
3.10	Effective and Realized Half-Spreads when Excluding Interdealer Trades	101
A.1	Descriptive Statistics of Price-Impact-Sorted Portfolios	120
A.2	Summary of Index Rules	123

List of Tables

B.1	Swap Execution Facilities' Reporting Formats	132
B.2	Effective Half-Spreads, Realized Half-Spreads, and Price Impacts	147
B.3	Effective Half-Spreads, Realized Half-Spreads, and Price Impacts by Trade Type	148
B.4	Regressions Controlling for Trade Characteristics and Market Conditions	149
B.5	Effective Half-Spreads, Realized Half-Spreads, and Price Impacts by Trade Type	150
B.6	Effective Half-Spreads, Realized Half-Spreads, and Price Impacts by Trade Size	151
B.7	Regressions Controlling for Trade Characteristics and Market Conditions	152
B.8	Regressions Controlling for Roll Characteristics and Market Conditions	153
B.9	Effective Half-Spreads, Realized Half-Spreads, and Price Impacts of Pairs	154
B.10	Effective Half-Spreads, Realized Half-Spreads, and Price Impacts by Trade Size	155
B.11	Regressions Controlling for Trade Characteristics and Market Conditions	156
B.12	Regressions Controlling for Trade Characteristics and Market Conditions	157
B.13	Effective Half-Spreads, Realized Half-Spreads, and Price Impacts of Pairs	159
B.14	VAR Estimates	163
B.15	VAR Estimates	165

1 Liquidity Risk in Credit Default Swap Markets

This chapter is based on joint work with Anders B. Trolle in which we show that liquidity risk is priced in the cross section of returns on credit default swaps (CDSs). We measure CDS market illiquidity by aggregating deviations of credit index levels from their no-arbitrage values implied by the index constituents' CDS spreads, and we construct a tradable liquidity factor from returns on index arbitrage strategies. CDS contracts with higher liquidity exposures have higher expected excess returns for sellers of credit protection; on average, liquidity risk accounts for 24% of CDS spreads. Illiquidity and risk premia correlate negatively with proxies for the risk-bearing capacity of CDS market intermediaries.

1.1 Introduction

A recent literature starting with Pástor and Stambaugh (2003) and Acharya and Pedersen (2005) has shown that liquidity risk is priced within a variety of asset classes including stocks, Treasuries, corporate bonds, hedge funds, and private equity.¹ In this paper, we study whether liquidity risk is priced in credit default swaps (CDSs). This issue is important for several reasons. First, CDS contracts are exposed to many potential sources of illiquidity as they trade in a relatively opaque, dealer-dominated, and decentralized market.² In particular, illiquidity stemming from funding and capital constraints of financial intermediaries is likely to play an important role for the pricing of CDS contracts. Second, in contrast to the asset classes listed above, CDS contracts are in zero net supply, implying that the sign on any liquidity risk premium is not given a priori. Third, from a practical perspective, liquidity risk is important

¹ As discussed by Acharya and Pedersen (2005), liquidity risk can be defined in several ways. The notion of liquidity risk used in this paper (covariation between returns and a market-wide liquidity factor) has been shown to be priced in stocks (see Sadka (2006) and Korajczyk and Sadka (2008) in addition to Pástor and Stambaugh (2003) and Acharya and Pedersen (2005)), Treasuries (see Li, Wang, Wu, and He (2009)), corporate bonds (see Lin, Wang, and Wu (2011)), hedge funds (see Sadka (2010) and Hu, Pan, and Wang (2013)), and private equity (see Franzoni, Nowak, and Phalippou (2012)).

² Recent regulatory reform has brought more transparency to the market for credit index contracts, the most liquid of which must now be traded on so-called swap execution facilities. However, single-name CDS contracts continue to trade with little transparency.

Chapter 1. Liquidity Risk in Credit Default Swap Markets

for the trading, pricing, hedging, and risk-management of CDS contracts—recently illustrated by J.P. Morgan’s six billion dollar trading loss on relatively illiquid CDS market strategies.³ Finally, from a regulatory perspective, liquidity risk is important given the potential systemic nature of the CDS market.

Measuring the liquidity of the CDS market is challenging given the over-the-counter (OTC) market structure and the absence of publicly available transaction data. In this paper we propose to measure CDS market illiquidity by the extent to which simple no-arbitrage relations in CDS markets are violated. Such violations reflect not only the direct transaction costs associated with exploiting the arbitrage opportunities but also constraints on the risk-bearing capacity of arbitrageurs and financial intermediaries; therefore, illiquidity in this paper is meant in broad terms.⁴

Specifically, we consider the law-of-one-price relation between a credit index and a basket of single-name CDSs that replicates the cash flow of the index. We denote the difference between the level of the index and its CDS-implied level as the *index-to-CDS basis*, and we construct a market-wide CDS illiquidity measure as a weighted average of absolute values of index-to-CDS bases. The average is taken over the four most representative and liquid indices of investment-grade and high-yield credit risk in North America and Europe. These indices cover a substantial part of the overall CDS market.⁵

We find time-varying index-to-CDS bases across all four credit indices during our sample period from September 20, 2006 to February 1, 2012. In particular, bases widened significantly in the period following the collapse of Lehman Brothers and AIG. For instance, bases of credit indices referencing North American investment-grade and high-yield names dropped to -61 basis points (bps) and -452 bps, respectively (corresponding to -25% and -38%, respectively, of the index levels). Our illiquidity measure suggests that CDS market liquidity was relatively high and stable in the early part of the sample period, deteriorated somewhat around the time of the collapse of two Bear Stearns structured-credit hedge funds in late June 2007, deteriorated significantly in the aftermath of the Lehman Brothers default and the AIG bailout in September 2008, and then recovered substantially since early 2009—although not reaching the level of liquidity that prevailed prior to the crisis. We show that the illiquidity measure correlates not only with bid-ask spreads and price impact measures in the CDS market but also with funding costs and the equity capital of the main dealers in the CDS market.

A key advantage of our illiquidity measure is that its innovations can be closely tracked by a tradable liquidity factor. For each credit index, we consider a trading strategy consisting

³ See “‘London Whale’ Rattles Debt Market,” *Wall Street Journal*, April 6, 2012.

⁴ Our notion of illiquidity is similar to that of Hu, Pan, and Wang (2013) who measure Treasury market illiquidity by the extent to which Treasury yields deviate from fair-value. An important difference between the illiquidity measures—apart from the fact that they pertain to different markets—is that their measure relies on a model to obtain fair-value yields, while our measure is completely model independent.

⁵ In constructing the illiquidity measure, we use five-year on-the-run index series. Very similar results are obtained when including additional maturities, off-the-run series, or sub-indices in the construction of the illiquidity measure.

of positions in both the credit index and its replicating basket that profits from a narrowing of the index-to-CDS basis. The tradable liquidity factor is the excess return on a portfolio of these credit index arbitrage trades. The correlation between the tradable liquidity factor and innovations to the CDS market illiquidity measure is -0.80 .

Next, we study whether exposure to our liquidity factor—i.e., liquidity risk—is priced in the cross section of returns on CDSs. This is motivated by the connection between our liquidity factor and changes in the risk-bearing capacity of financial intermediaries, and the recent literature on intermediary asset pricing, which predicts that proxies for the risk-bearing capacity of financial intermediaries show up as pricing factors (see, e.g., Brunnermeier and Pedersen (2009), He and Krishnamurthy (2013), and Kondor and Vayanos (2014)). Models of intermediary asset pricing seem particularly relevant for the CDS market because it is highly concentrated around a group of global credit derivatives dealers (often referred to as the G14 dealers), who participate in virtually all transactions.⁶ Given that dealers in aggregate are net sellers of credit protection to end-users⁷ we expect that, from a protection seller's perspective, (i) realized excess returns on CDS contracts correlate positively with the liquidity factor (since CDS spreads widen when dealers become more constrained); (ii) expected excess returns are higher on contracts with higher liquidity exposures; and (iii) risk premia increase when intermediary constraints tighten.

We estimate a factor pricing model which, in its basic formulation, has two systematic factors: a default factor and the tradable liquidity factor.⁸ The underlying data set is a large panel of single-name CDS contracts referencing 666 North American and European entities. These contracts are sorted into portfolios that exhibit variation in credit quality and the level of illiquidity. Across all portfolios, unconditional expected excess returns are positive from a credit protection seller's perspective, ranging from 0.35% per year for a portfolio of the most liquid high-credit-quality CDSs to 5.80% per year for a portfolio of the most illiquid low-credit-quality CDSs.⁹ Factor exposures (i.e., betas) are estimated from time-series regressions of realized excess returns on the factors. Consistent with prediction (i), liquidity betas are

⁶ Adrian, Etula, and Muir (2014) find empirical support for intermediary asset pricing models in the context of the stock market.

⁷ The aggregate positions of dealers and end-users can be obtained from the Depository Trust & Clearing Corporation's Trade Information Warehouse (TIW). On the first date that TIW data is available (October 31, 2008), the gross notional amount of credit protection sold by dealers to end-users was 171 billion USD larger than the gross notional amount bought by dealers from end-users. At the end of our sample period, the difference was 102 billion USD.

⁸In principle, counterparty risk could also be a determinant of CDS returns. However, Arora, Gandhi, and Longstaff (2012) find that the effect of counterparty risk on CDS spreads is negligible, which is consistent with the widespread use of collateralization and netting agreements. Indeed, from the annual "Margin Surveys" by the International Swaps and Derivatives Association (ISDA), we infer that the fraction of credit derivatives trades covered by collateral agreements averaged more than 80% over the sample period. Hence, we do not take counterparty risk into account in our factor pricing model.

⁹ Sample means of realized excess returns on CDSs are very imprecise estimates of expected excess returns because of the short sample period and the peso problem that arises when computing returns on securities that are subject to credit risk (as credit events are rare and have a dramatic impact on returns when they occur). Therefore, we follow Bongaerts, de Jong, and Driessen (2011) in obtaining forward-looking estimates of conditional expected excess returns by using Moody's KMV Expected Default Frequencies to calculate expected default losses.

Chapter 1. Liquidity Risk in Credit Default Swap Markets

positive (realized excess returns for credit protection sellers correlate positively with the liquidity factor).

Factor prices of risk are estimated from a cross-sectional regression of unconditional expected excess returns on betas.¹⁰ Consistent with prediction (ii), liquidity risk is priced, and sellers of credit protection earn higher expected excess returns on CDS contracts with higher liquidity exposures. The price of liquidity risk is not only statistically significant but also economically important. For instance, considering the expected excess return differential of 5.45% per year between the portfolio of the most illiquid low-credit-quality CDSs and the portfolio of the most liquid high-credit-quality CDSs, 2.21% per year is due to liquidity risk, while 3.30% is due to default risk (the remainder is a pricing error). Alternatively, considering the average expected excess return across all portfolios, 0.59% per year is due to liquidity risk, while 1.08% is due to default risk. We also decompose CDS spreads instead of expected excess returns. Averaging the relative contributions across portfolios, liquidity risk accounts for 24% of model-implied CDS spreads, while default risk and expected default losses account for 47% and 29%, respectively.

We also consider a conditional version of the factor pricing model in which time-varying factor prices of risk are estimated from cross-sectional regressions of conditional expected excess returns on betas. Consistent with prediction (iii), we find that factor prices of risk increase when intermediary equity capital decreases. Specifically, prices of default and liquidity risk have correlations of -0.71 and -0.32, respectively, with a proxy for intermediary equity capital. Factor prices of risk are particularly high in the aftermath of the default of Lehman Brothers and the bailout of AIG. Other periods during which factor prices of risk are high include the initial phase of the financial crisis as well as the last part of the sample period when the European sovereign debt crisis intensified.

We conduct a range of robustness checks; we control for the contract-specific level of illiquidity, consider an alternative construction of the tradable liquidity factor, and add additional risk factors to the asset pricing model. These include a factor that correlates with the availability of arbitrage capital, corporate bond and stock market illiquidity factors, as well as stock market and volatility factors. Across the robustness checks, the contribution of liquidity risk to the expected excess return differential mentioned above ranges from 1.59% to 2.22% per year.

A number of studies have documented liquidity effects in the pricing of CDS contracts but mainly focus on the effect of the contract-specific level of illiquidity. Most studies find that CDS spreads increase with contract-specific illiquidity; see, e.g., Bühler and Trapp (2009) who infer the liquidity component in CDS spreads via an intensity-based pricing model and Tang and Yan (2007), Qiu and Yu (2012), Lesplingart, Majois, and Petitjean (2012), and Pires, Pereira, and Martins (2014) who run panel regressions of CDS spreads on various illiquidity proxies.¹¹ Our

¹⁰ Standard errors are adjusted for errors-in-variables, heteroscedastic and autocorrelated errors, and potential model misspecification as in Kan, Robotti, and Shanken (2013). They show, along with other recent papers, that taking into account potential model misspecification is crucial for reliable statistical inference.

¹¹ Chen, Fabozzi, and Sverdløve (2010) also use an intensity-based pricing model to estimate the magnitude of the liquidity component in CDS spreads; however, the sign of the liquidity component is hard-wired into their

paper differs from these papers by instead focussing on the pricing of systematic liquidity risk. We find that once systematic liquidity risk is accounted for, the effect of the contract-specific level of illiquidity on the pricing of CDS contracts becomes much less pronounced.¹²

Bongaerts et al. (2011) investigate the effect of both the contract-specific level of transaction costs and liquidity risk on the expected excess return on CDS contracts. However, they use a very different definition of liquidity risk; they focus on covariation between innovations to transaction costs on individual CDS contracts and systematic default risk, while we focus on covariation between CDS returns and market-wide CDS liquidity. They find that their notion of liquidity risk is not significantly priced in the cross section of CDS returns, leading them to attribute most of the variation in expected excess returns to variation in the levels of transaction costs. In contrast, we find that our notion of liquidity risk is significantly priced in the cross section of CDS returns and largely subsumes the effect of the levels of transaction costs.¹³

Finally, our finding that risk premia increase when intermediary equity capital decreases is consistent with recent work by Siriwardane (2015) who shows in a post-crisis sample that capital losses for sellers of credit protection lead to wider CDS spreads. In a similar vein, Froot and O'Connell (1999) show that capital losses for sellers of catastrophe insurance lead to higher insurance premia, and Greenwood and Vayanos (2014) find that capital losses for arbitrageurs in the Treasury market lead to higher bond risk premia.

The paper proceeds as follows: Section 1.2 describes the construction of the CDS market illiquidity measure and the tradable liquidity factor, Section 1.3 investigates the pricing of liquidity risk, and Section 1.4 concludes.

1.2 Measuring CDS Market Illiquidity

This section presents the construction of the CDS market illiquidity measure and the tradable liquidity factor. Furthermore, it explores determinants of CDS market illiquidity and it briefly describes credit indices and the replication argument on which index arbitrage is based.

1.2.1 Credit Indices

Credit indices are standardized credit derivatives that provide insurance against any defaults

model.

¹² Tang and Yan (2007) and Lesplingart et al. (2012) also provide tentative results on the effect of systematic liquidity risk by augmenting their panel regressions with the betas that appear in Acharya and Pedersen's (2005) liquidity-adjusted CAPM. However, their results are somewhat inconclusive. Moreover, the Acharya and Pedersen (2005) model is based on the assumption that assets are in positive net supply, making it not directly applicable to CDS contracts that are in zero net supply.

¹³ We verify that Bongaerts et al.'s (2011) notion of liquidity risk is also not priced in our more recent and broader sample of CDSs. We also argue that the frequent marking to market of CDS contracts makes their notion of liquidity risk less relevant in case of the CDS market.

Chapter 1. Liquidity Risk in Credit Default Swap Markets

among their constituents. They allow investors to gain or reduce credit risk exposure in certain segments of the market. Due to their widespread use and standardized terms, credit indices are more liquid than both single-name CDSs and corporate bonds.¹⁴ Credit indices trade in OTC markets for maturities between one and ten years. The five-year maturity is typically the most liquid and is the focus of our empirical analysis.¹⁵

Each credit index is a separate CDS contract with a specific maturity, fixed spread, and underlying basket of reference entities. Over the life of the contract, the seller of protection on the index provides default protection on each index constituent, with the notional amount of the contract divided evenly among the index constituents. In return, the seller of index protection earns the fixed spread. In case of default, the seller of index protection pays the loss-given-default and the notional amount of the contract is reduced accordingly. If the quoted level of the index differs from its fixed spread, counterparties initially exchange an upfront payment equal to the contract's present value.

As a clarifying example, suppose that on September 21, 2007 an investor sells a 10 million USD notional amount of protection on the main North American investment-grade credit index (CDX.NA.IG.9) with a maturity of five years and a fixed spread of 60 bps.¹⁶ On that date the index traded at 50 bps which translates into a 46,183 USD upfront charge for the seller of protection. Over the next three quarters he receives quarterly spread payments each being approximately equal to $1/4 \times 0.0060 \times 10,000,000 = 15,000$ USD (for the purpose of illustration, we abstract from the actual day-count convention). On September 7, 2008, Fannie Mae and Freddy Mac, both reference names of the CDX.NA.IG.9, were placed into conservatorship by their regulator. Creditors recovered 91.51 cents and 94 cents per dollar of senior unsecured debt issued by Fannie Mae and Freddy Mac, respectively. Thus, the seller of index protection compensates the losses incurred, paying $1/125 \times (1 - 0.9151) \times 10,000,000 + 1/125 \times (1 - 0.94) \times 10,000,000 = 11,592$ USD.¹⁷ Due to the credit events, the spread payment on September 20, 2008 is reduced to $1/4 \times 123/125 \times 0.0060 \times 10,000,000 = 14,760$ USD. Until expiry of the index on December 20, 2012, another two credit events occurred: first, the default of Washington Mutual on September 27, 2008 triggers a $1/125 \times (1 - 0.57) \times 10,000,000 = 34,400$ USD payout and reduces subsequent spread payments to $1/4 \times 122/125 \times 0.0060 \times 10,000,000 = 14,640$ USD. Second, the Chapter 11 filing of CIT Group on November 1, 2009 triggers a $1/125 \times (1 - 0.68125) \times 10,000,000 = 25,500$ USD payout and reduces successive spread payments to

¹⁴ For example, the Depository Trust & Clearing Corporation's "Market Activity Report" for the three-month period from June 20, 2011 to September 19, 2011, shows that the average daily notional amount of trades is 29 million USD, on average, across single-name CDSs referencing corporate names that belong to the 1000 most actively traded single-name CDSs. In contrast, the average daily notional amount of untranching index transactions is approximately one billion USD.

¹⁵ Using CDS transaction data, Chen, Fleming, Jackson, Li, and Sarkar (2011) find that 84% of all index transactions are in the five-year maturity.

¹⁶ The number following the index name is referred to as the index's series and uniquely identifies the underlying basket of reference names.

¹⁷ In this example, we assume that cash settlement, the standard settlement method of credit index transactions, applies. Furthermore, we ignore accrual payments on default and the fact that recovery values are determined in credit event auctions that usually do not take place on the credit event dates.

$$1/4 \times 121/125 \times 0.0060 \times 10,000,000 = 14,520 \text{ USD.}$$

Twice a year, on the so-called index roll dates in March and September, a new series of each credit index is launched, with the basket of reference entities revised according to credit rating and liquidity criteria. Entities that fail to maintain a credit rating within a specified range, due to either an upgrade or a downgrade, and entities whose CDS contracts have deteriorated significantly in terms of their liquidity are replaced by the most liquid reference names meeting the credit rating requirements. Liquidity is typically concentrated in the most recently launched series, which are referred to as the on-the-run series. Consequently, these are the subject of our empirical analysis.

In case of a credit event, a new version of the index series starts trading, with the entity that triggered the event having been removed from the index. Because triggered CDSs usually continue to trade in the market until their recovery values are determined, multiple versions of the same index series can trade at the same time. In such cases, we focus on the most liquid version.

All the credit indices considered in this paper are administrated by Markit. It sets the rules and procedures that govern the index revisions on the roll dates. In addition, it determines a group of licensed dealers, who actively make markets for credit indices. Based on their spread quotes, Markit computes index levels that are published on a daily basis.

1.2.2 Index Replication

Investors can gain credit risk exposure either by selling protection on the index contract or by selling protection on a basket of single-name CDSs that replicates the cash flow of the index contract. Thus, an alternative index level can be implied from single-name CDS quotes on the index constituents. This gives rise to what we call the index-to-CDS basis, defined as the difference between the index level and the CDS-implied level. In perfect capital markets, index arbitrage will keep index-to-CDS bases close to zero.

Suppose that on date t an investor wants to sell index protection with a five-year maturity, fixed spread C , and notional amount A . This involves an initial upfront payment equal to the contract's present value.¹⁸ Instead of selling index protection, the investor can sell protection on the index constituents via single-name CDSs. In particular, to replicate the payments of the index contract, the investor must sell protection on each of the I_t index constituents that, prior to the inception of trade, have not triggered a credit event. Each single-name CDS must have a five-year maturity, fixed spread C , and notional amount A/I , where I denotes the number of reference entities at the launch of the index's series. As for credit indices, upfront payments are necessary when trading single-name CDSs at off-par spreads. Hence, the investor faces costs

¹⁸In addition, there will be an accrual payment. The seller of index protection is entitled to a full spread payment on the first payment date after inception of trade, regardless of the actual time of opening his position. Therefore, he has to compensate the buyer of protection for the fixed spread accrued between the last spread payment date and the inception of trade. We abstract from these accrual payments in our discussion of the index replication.

Chapter 1. Liquidity Risk in Credit Default Swap Markets

equal to the aggregate amount of all upfront charges from the single-name CDS transactions.

Until the earlier of the maturity date and the first credit event by one of the remaining index constituents, the seller of index protection earns quarterly spread payments of $d/360 \times C \times I_t/I \times A$, while the seller of protection via single-name CDSs receives quarterly spread payments of $\sum_{i=1}^{I_t} d/360 \times C \times A/I$. Here, $I_t/I \times A$ is the index's adjusted notional amount, and $d/360$ denotes the accrual time during a given quarter determined by ACT/360 day-count convention. Obviously both payment streams are identical.

In case that one of the remaining reference names, say i^* , defaults prior to maturity, the seller of index protection has to make a payment of $1/I \times (1 - R_{i^*}) \times A$, where R_{i^*} is the recovery per dollar of notional on i^* 's debt. This payment coincides with the one that the seller of protection via single-name CDSs has to make.¹⁹

Following the credit event, the notional amount of the index is adjusted to $(I_t - 1)/I \times A$ and quarterly spread payments earned by the seller of index protection decrease to $d/360 \times C \times (I_t - 1)/I \times A$. Because there is also one single-name CDS less in the basket, the seller of protection via single-name CDSs collects quarterly spread payments of $\sum_{i=1}^{I_t-1} d/360 \times C \times A/I$. Thus, payments coincide in this case as well.

Because the same reasoning applies to any possible credit event that may occur prior to maturity, it follows that the cash flows for the seller of index protection and the seller of protection via single-name CDSs are identical. The CDS-implied index level, C_t^{CDS} , can be thought of as that fixed spread on the single-name CDSs that makes the replicating basket have zero net present value.²⁰ The index-to-CDS basis, B_t , of a credit index is then defined as $B_t = C_t^{IDX} - C_t^{CDS}$, where C_t^{IDX} denotes the index level as of date t .

1.2.3 Data

The credit index data are obtained from Markit and comprise index levels, CDS-implied levels, and the corresponding upfront amounts. In addition, the number of licensed dealers that submit spread quotes for the computation of the index level is reported. We use the four most representative and liquid indices of investment-grade and high-yield credit risk in North America and Europe, always focusing on the five-year maturity. For each index, we splice together on-the-run series to create continuous time series for the period from September 20, 2006 to February 1, 2012. Whenever multiple versions of the on-the-run series trade simultaneously, we choose the version with the largest number of contributing dealers.

¹⁹Upon default, both the seller of index protection and the seller of protection via single-name CDSs will receive an accrual payment. This payment compensates for the protection they provided on the defaulted reference name since the last spread payment date prior to the credit event.

²⁰Because the ISDA CDS Standard Model is used to convert between upfront amounts and index levels, the CDS-implied index level is the par spread on a hypothetical single-name CDS contract whose upfront amount equals that of the replicating basket of single-name CDSs. The contract terms of the hypothetical single-name CDS and the recovery rate used for conversion are specified in the index's contract terms.

1.2. Measuring CDS Market Illiquidity



Figure 1.1: Credit Index Levels, CDS-Implied Index Levels, and Index-to-CDS Bases.

The figure displays daily observations of credit index levels of the five-year on-the-run series (thin black lines, left hand scales), CDS-implied index levels (thick gray lines, left hand scales), and index-to-CDS bases (light gray shaded areas, right hand scales) from September 20, 2006 to February 1, 2012. Index levels and bases are in basis points and dashed vertical lines correspond to index roll dates.

The four credit indices are CDX.NA.IG, CDX.NA.HY, iTraxx Eur, and iTraxx Xover. CDX.NA.IG and iTraxx Eur each comprise 125 investment-grade reference names from North America and Europe, respectively. CDX.NA.HY comprises 100 high-yield reference names from North America, while iTraxx Xover comprises up to 50 high-yield reference names from Europe. Table A.2 at the end of Appendix A summarizes index rules and contract terms for these indices.

Figure 1.1 displays time series of the on-the-run index levels (thin black lines). Each of the indices increased shortly before the March 2008 roll date when Bear Stearns was on the brink of bankruptcy and, after a short period of relief, peaked in the aftermath of the September 2008 credit events of Fannie Mae, Freddy Mac, Lehman Brothers, and Washington Mutual. The iTraxx indices again sharply increased in July 2011 as the European sovereign debt crisis intensified.²¹ Descriptive statistics of the credit index levels are reported in Panel A of Table 1.1.

²¹ This month saw a sharp sell-off in non-core European sovereign bonds, partly triggered by downgrades of the sovereign debt of Portugal and Ireland to non-investment-grade status (see "Italy Fears Jolt Markets," *Wall Street*

Chapter 1. Liquidity Risk in Credit Default Swap Markets

Panel A: Credit Index Levels				
	CDX.NA.IG	CDX.NA.HY	iTraxx Eur	iTraxx Xover
Mean	106.5	630.9	97.6	514.0
Standard Deviation	49.7	312.6	48.4	217.8
Minimum	28.9	208.5	20.1	170.8
Maximum	279.7	1893.6	215.9	1150.3
N	1338	1337	1357	1356
Panel B: Index-to-CDS Bases				
	CDX.NA.IG	CDX.NA.HY	iTraxx Eur	iTraxx Xover
Mean	-4.9	2.4	-3.8	3.6
Standard Deviation	11.6	68.3	9.0	16.5
Minimum	-61.1	-451.9	-58.5	-106.1
Maximum	12.2	172.4	13.9	49.9
$\text{Corr}(C_t^{IDX}, \sigma_t(B))$	0.85	0.91	0.73	0.87

Table 1.1: Descriptive Statistics of Credit Index Levels and Index-to-CDS Bases.

The table displays descriptive statistics of credit index levels and index-to-CDS bases. Panel A provides descriptive statistics of credit index levels of the five-year on-the-run series. Panel B provides descriptive statistics of the corresponding index-to-CDS bases. Mean, standard deviation, minimum, and maximum are in basis points, and $\text{Corr}(C_t^{IDX}, \sigma_t(|B|))$ denotes the time-series correlation between the index level and the conditional volatility of the index-to-CDS basis' absolute value. For each index, the conditional volatility is inferred from a GARCH(1,1) model for the conditional variance of the error term in an ARMA(1,1) specification of the absolute value of the index-to-CDS basis. The sample period is from September 20, 2006 to February 1, 2012. N denotes the number of daily observations.

1.2.4 CDS Market Illiquidity Measure

In addition to index levels, Figure 1.1 also displays CDS-implied levels (thick gray lines) and the corresponding index-to-CDS bases (light gray shaded areas). Nonzero index-to-CDS bases frequently arise; in particular, between the September 2008 index roll and the next index roll in March 2009, i.e., at the height of the financial crisis, bases are wide and very volatile. Bases of the investment-grade indices CDX.NA.IG and iTraxx Eur drop to -61.1 bps and -58.6 bps, respectively, while bases of the high-yield indices CDX.NA.HY and iTraxx Xover drop to -451.9 bps and -106.2 bps, respectively.²² These numbers correspond to -25.3%, -33.4%, -38.0%, and -9.8% of the index levels. Descriptive statistics of the index-to-CDS bases are reported in Panel B of Table 1.1. Sample means of bases are negative for investment-grade indices and positive for high-yield indices, and standard deviations of bases are higher for the high-yield indices than for the investment-grade indices.

Journal, July 12, 2011).

²² One explanation for the negative index-to-CDS bases at the height of the financial crisis is the following: because index contracts traded significantly above their fixed spreads, a seller of credit protection via an index contract would have received a large upfront payment. On the other hand, because most single-name CDSs were typically executed at par spreads during that time, a seller of credit protection via a portfolio of single-name CDS contracts would often not have received an upfront payment. This created an incentive for funding- and capital-constrained dealers to sell protection via index contracts, helping to push index-to-CDS bases deep into negative territory.

As explained below, we measure illiquidity from absolute values of index-to-CDS bases. We observe a strong relation between index levels and volatilities of absolute bases. The table shows that the time-series correlations between index levels and conditional volatilities of absolute bases range from 0.73 to 0.91 across indices. Furthermore, the cross-sectional correlation between average index levels and unconditional volatilities of absolute bases is 0.74.

To measure market-wide CDS illiquidity, we aggregate absolute values of index-to-CDS bases across indices. We use absolute values because positive and negative bases are equally informative about illiquidity in the CDS market. Given the significant cross-sectional variation in the volatilities of absolute bases, taking an equally weighted average of absolute bases would cause variation in the CDS market illiquidity measure to be driven mostly by the bases of high-yield indices. One option would be to weight absolute bases by the inverse of their conditional volatilities; however, this has the disadvantage that the sample period is shortened by the window over which the initial conditional volatilities are estimated. Instead, we opt to weight absolute bases by the inverse of the index levels, exploiting the strong relation between the index levels and the volatilities of the absolute bases. Thus, the CDS market illiquidity measure, $CDSILLIQ_t$, is given by

$$CDSILLIQ_t = \sum_{i=1}^{n_t} w_{i,t} |B_{i,t}|, \quad (1.1)$$

where $w_{i,t} = (1/C_{i,t}^{IDX}) / (\sum_{j=1}^{n_t} 1/C_{j,t}^{IDX})$ and n_t is the number of indices with available data on date t .

Figure 1.2 shows the time series of the CDS market illiquidity measure at a weekly frequency. As can be seen from the figure, the illiquidity measure is very persistent with a 0.93 first-order autocorrelation. The measure suggests that liquidity was relatively high and stable until June 2007. The deterioration in liquidity towards the end of that month coincided with the high-profile collapse of two Bear Stearns structured-credit hedge funds, which was followed by further turmoil in credit and funding markets.²³ Liquidity deteriorated significantly in the aftermath of the Lehman Brothers default and the AIG bailout in September 2008, and the illiquidity measure peaked at 79 bps at the end of December 2008. Since then, liquidity has recovered substantially, but, within our sample period, did not reach pre-crisis levels.

The construction of our CDS market illiquidity measure is very robust. We have considered three alternative constructions that use a larger number of indices or index series. First, we include sub-indices (on-the-run series with five-year maturities) of the four credit indices, in which case the average in Equation (1.1) is taken over ten indices.²⁴ Second, we include the

²³The two funds—the High Grade Structured Credit Strategies Fund, and the High Grade Structured Credit Strategies Enhanced Leverage Fund—were largely invested in collateralized debt obligations tied to subprime mortgages (see “Two Big Funds At Bear Stearns Face Shutdown,” *Wall Street Journal*, June 20, 2007).

²⁴The sub-indices are CDX.NA.IG.HVOL, CDX.NA.HY.BB, CDX.NA.HY.B, iTraxx Eur HiVol, iTraxx Eur Sr Finls, and iTraxx Eur Sub Finls. CDX.NA.IG.HVOL (iTraxx Eur HiVol) comprises the 30 reference names from the CDX.NA.IG (iTraxx Eur) with the widest five-year CDS spreads. iTraxx Eur Sr Finls comprises the 25 financial sector reference




Figure 1.2: CDS Market Illiquidity Measure.

The figure displays the CDS market illiquidity measure (in basis points). The time series consists of 281 weekly observations from September 20, 2006 to February 1, 2012. Dotted vertical lines correspond to (from left to right) the collapse of two Bear Stearns structured-credit hedge funds on June 20, 2007, the Bear Stearns near-bankruptcy on March 17, 2008, the default of Lehman Brothers on September 15, 2008, and July 1, 2011 marking the beginning of a month in which a sharp sell-off in non-core European government bonds intensified the European sovereign debt crisis.

full term structure of on-the-run series of the four credit indices, in which case the average in Equation (1.1) is taken over 16 index series. Third, we include the immediate off-the-run series, i.e., the series that most recently became off-the-run, of the four credit indices (with five-year maturities), in which case the average in Equation (1.1) is taken over eight index series. The original CDS market illiquidity measure is very highly correlated with these alternative constructions both in levels (correlations between 0.97 and 0.99) and weekly changes (correlations between 0.87 and 0.98). We prefer the original construction because it is more parsimonious.

1.2.5 Determinants of CDS Market Illiquidity

An advantage of our illiquidity measure is that it captures many dimensions of illiquidity in the CDS market, including constraints on the risk-bearing capacity of financial intermediaries. Here, we investigate the relation between our illiquidity measure and alternative measures of CDS market illiquidity as well as measures of intermediary constraints. Exact definitions of all variables are provided in Appendix A.1 and their time-series dynamics are exhibited in

names from the iTraxx Eur. iTraxx Eur Sub Finls comprises the same reference names as the iTraxx Eur Sr Finls, but reference obligations are subordinated. CDX.NA.HY.BB and CDX.NA.HY.B comprise, respectively, BB- and B-rated reference names from the CDX.NA.HY.

Figure A.1 at the end of Appendix A.

We consider three alternative measures of CDS market illiquidity. The first measure is the average bid-ask spread of single-name CDSs. When average bid-ask spreads are wider, index arbitrage is more expensive, and index-to-CDS bases can drift further away from zero before index arbitrage becomes profitable. The second measure is the absolute spread change per contributed quote, averaged across single-name CDSs. To the extent that volume can be proxied by the number of contributors, this captures the price impact of CDS trades much like the Amihud (2002) illiquidity measure. The third measure is the absolute change in the index level per contributed quote, averaged across on-the-run credit indices. This captures the price impact of index trades. We expect higher CDS market illiquidity the higher the price impact of trade for single-name and index contracts.

We consider several measures of funding and capital constraints of financial intermediaries. Unsecured funding costs are proxied by the LIBOR-OIS spread (see, e.g., Filipović and Trolle (2013)), secured funding costs are proxied by the spread between Agency MBS and Treasury general collateral repo rates (see, e.g., Bartolini, Hilton, Sundaresan, and Tonetti (2011)), and intermediary equity capital is proxied by the market capitalization of the financial institutions that make up the G14 group of major credit derivatives dealers. We also include indirect measures of the risk-bearing capacity of the intermediary sector including market volatility proxied by the VIX index (see, e.g., Brunnermeier and Pedersen (2009)), the Hu, Pan, and Wang (2013) “Noise” measure of deviations of Treasury yields from a smooth yield curve, and the CDS-bond basis averaged across U.S. investment-grade bonds (see, e.g., Duffie (2010) and Mitchell and Pulvino (2012)).

We run univariate regressions of monthly changes in the CDS market illiquidity measure on monthly changes in the explanatory variables. We run regressions in first differences to avoid spurious results due to persistence of the dependent and explanatory variables (unit root tests are available upon request). For those measures that are available at a daily frequency, we obtain the monthly time series by averaging daily observations within each month. Panel A of Table 1.2 shows the regression results with Newey and West (1987) t -statistics given in brackets, and Panel B reports correlations of monthly changes in the explanatory variables.²⁵ All slope coefficients have the expected sign. The CDS market illiquidity measure is significantly related to bid-ask spreads (t -stat of 4.29, R^2 of 0.36), the price impact of credit index trades (t -stat of 3.04, R^2 of 0.18), unsecured funding costs (t -stat of 2.06, R^2 of 0.08), intermediary equity capital (t -stat of -2.42, R^2 of 0.12), the VIX index (t -stat of 2.68, R^2 of 0.12), the “Noise” measure (t -stat of 6.27, R^2 of 0.46), and the CDS-bond basis (t -stat of -3.24, R^2 of 0.26).²⁶ This confirms the multidimensional nature of our CDS market illiquidity measure, including its relation to

²⁵ As can be seen from Panel B of Table 1.2, many of the explanatory variables are relatively highly correlated. Thus, a multivariate regression including all variables will be subject to multicollinearity. Collectively, all variables together explain 67% of the time-series variation of the CDS market illiquidity measure.

²⁶ That the CDS market illiquidity measure is not significantly related to secured funding costs indicates that index arbitrage traders primarily fund their trades in unsecured interbank markets, which is consistent with the fact that CDSs cannot be used as collateral in repo transactions.

Panel A: Regression Results										
	$\Delta Bid-Ask$	$\Delta ILLIQ^{CDS}$	$\Delta ILLIQ^{IDX}$	$\Delta LIB-OIS$	$\Delta Repo$	$\Delta Capital$	ΔVIX	$\Delta Noise$	$\Delta CDS-Bond$	
Intercept	-0.16 [-0.31]	0.09 [0.12]	0.06 [0.08]	0.06 [0.08]	0.10 [0.13]	-0.23 [-0.35]	0.05 [0.07]	0.08 [0.26]	0.06 [0.12]	
Slope	1.34 [4.29]	1.04 [1.03]	6.49 [3.04]	5.27 [2.06]	9.18 [1.63]	-1.98 [-2.42]	0.34 [2.68]	3.05 [6.27]	-12.12 [-3.24]	
R^2	0.36	0.05	0.18	0.08	0.06	0.12	0.12	0.46	0.26	
Panel B: Correlations										
	$\Delta Bid-Ask$	$\Delta ILLIQ^{CDS}$	$\Delta ILLIQ^{IDX}$	$\Delta LIB-OIS$	$\Delta Repo$	$\Delta Capital$	ΔVIX	$\Delta Noise$	$\Delta CDS-Bond$	
$\Delta ILLIQ^{CDS}$	0.52									
$\Delta ILLIQ^{IDX}$	0.45	0.48								
$\Delta LIB-OIS$	0.27	0.24	0.30							
$\Delta Repo$	0.38	0.11	0.26	0.75						
ΔVIX	-0.63	-0.37	-0.45	-0.12	-0.24					
$\Delta Capital$	0.57	0.41	0.65	0.61	0.68	-0.57				
$\Delta Noise$	0.72	0.28	0.24	0.31	0.50	-0.50	0.54			
$\Delta CDS-Bond$	-0.64	-0.45	-0.24	-0.56	-0.46	0.29	-0.49	-0.63		
$\Delta CDSILLIQ$	0.60	0.22	0.42	0.28	0.25	-0.35	0.35	0.68	-0.51	

Table 1.2: Determinants of CDS Market Illiquidity.

The table displays regression results and correlations between explanatory variables and between explanatory variables and the CDS market illiquidity measure. Panel A displays results from regressing monthly changes in the CDS market illiquidity measure on monthly changes in the average bid-ask spread of single-name CDSs ($Bid-Ask$), the average absolute spread change per quote contributed across single-name CDSs ($ILLIQ^{CDS}$), the average absolute change in the index level per quote contributed across on-the-run credit indices ($ILLIQ^{IDX}$), the spread between three-month LIBOR and OIS rates ($LIB-OIS$), the spread between three-month Agency MBS and Treasury general collateral repo rates ($Repo$), the aggregate market capitalization of financial institutions that make up the G14 group of major credit derivatives dealers ($Capital$), the VIX index (VIX), the Hu, Pan, and Wang (2013) "Noise" measure ($Noise$), and the average CDS-bond basis across U.S. investment-grade bonds ($CDS-Bond$). Reported are intercepts and slope coefficients, their respective t -statistics in brackets, and R^2 's. t -statistics are based on Newey and West (1987) heteroscedasticity and autocorrelation consistent standard errors with three lags. Panel B displays correlations between monthly changes in the explanatory variables and between monthly changes in the explanatory variables and monthly changes in the CDS market illiquidity measure. $CDS-Bond$, $LIB-OIS$, $Repo$, and VIX are in %. $Bid-Ask$, $CDSILLIQ$, and $Noise$ are in basis points. $ILLIQ^{CDS}$ and $ILLIQ^{IDX}$ are in basis points per contributed quote. $Capital$ is in 100 billion USD. The time series consist of 64 monthly observations from October 2006 to January 2012.

the risk-bearing capacity of the intermediary sector.

1.2.6 Tradable Liquidity Factor

An additional advantage of our illiquidity measure is that its innovations can be closely tracked by a tradable liquidity factor based on index arbitrage strategies. For each index, we consider a trading strategy that profits from a narrowing of the index-to-CDS basis. If the index trades above its CDS-implied level, the strategy sells protection on the index contract and buys protection via the replicating basket of single-name CDSs. If the index trades below its CDS-implied level, the strategy is the reverse trade. As shown in Section 1.2.2, if held to the index's maturity, this strategy is an arbitrage in a textbook sense. However, for a shorter holding period, the strategy is risky. For some index i , the holding period excess return on the strategy is given by

$$\text{sgn}(B_{i,t-1}) \left(r_{i,t}^{IDX} - r_{i,t}^{CDS} \right), \quad (1.2)$$

where $r_{i,t}^{IDX}$ and $r_{i,t}^{CDS}$ denote holding period excess returns from selling protection on the index contract and via its replicating basket of single-name CDSs, respectively, and $B_{i,t-1}$ is the index-to-CDS basis at the beginning of the holding period.²⁷ Because excess returns on the strategy are positive when index-to-CDS bases narrow, excess returns should be negatively correlated with changes in the absolute basis.

We construct the tradable liquidity factor, LIQ_t , by aggregating the excess returns on the individual index arbitrage strategies using the same weighting scheme as for $CDSILLIQ_t$; that is,

$$LIQ_t = \sum_{i=1}^{n_t} w_{i,t-1} \text{sgn}(B_{i,t-1}) \left(r_{i,t}^{IDX} - r_{i,t}^{CDS} \right), \quad (1.3)$$

where the weights, $w_{i,t-1}$, are given in Section 1.2.4. We use a one-week holding period because the asset pricing model in Section 1.3 is estimated at a weekly frequency. Descriptive statistics of the returns on the individual index arbitrage trades are given in Panel A of Table 1.3. There is considerable variation in the means and standard deviations of excess returns and the annualized Sharpe ratios, using Lo's (2002) correction for non-i.i.d. excess returns, are between 1.15 and 2.42. These Sharpe ratios are not directly realizable for an index arbitrageur because we ignore transaction costs.

Figure 1.3 displays the time-series evolution of the tradable liquidity factor. Its correlation with changes in the CDS market illiquidity measure is -0.80. The factor's annualized mean and standard deviation are 2.65% and 1.20%, respectively, and its annualized Sharpe ratio, using Lo's (2002) correction for non-i.i.d. excess returns, is 2.52. The high Sharpe ratio of the factor reflects the diversification that comes from the moderate correlations between the excess

²⁷We compute holding period excess returns from upfront amounts on credit index contracts and their replicating baskets of single-name CDSs, see Appendix A.2 for details.

Chapter 1. Liquidity Risk in Credit Default Swap Markets

Panel A: Return Descriptive Statistics				
	CDX.NA.IG	CDX.NA.HY	iTraxx Eur	iTraxx Xover
Mean	4.74	12.90	2.88	12.31
Standard Deviation	21.60	72.95	16.30	44.71
Sharpe Ratio	2.42	1.49	1.59	1.15
Skewness	0.67	0.25	0.84	0.70
Kurtosis	13.61	6.00	9.45	6.14
ρ_1	-0.11	-0.06	-0.01	0.04
N	277	268	279	277

Panel B: Pairwise Correlations				
	CDX.NA.IG	CDX.NA.HY	iTraxx Eur	iTraxx Xover
CDX.NA.IG		0.18	0.43	0.09
CDX.NA.HY			0.23	0.14
iTraxx Eur				0.28

Table 1.3: Descriptive Statistics of Index Arbitrage Returns.

The table displays descriptive statistics of one-week excess returns on the trading strategies underlying the construction of the tradable liquidity factor. Mean and standard deviation are in basis points per week, the Sharpe ratio is annualized using Lo's (2002) correction for non-i.i.d. excess returns, and ρ_1 denotes first-order autocorrelation. The sample period is from October 4, 2006 to February 1, 2012. N denotes the number of weekly observations.

returns on the individual arbitrage trades, see Panel B of Table 1.3.

1.3 Pricing of Liquidity Risk

This section investigates whether exposure to our liquidity factor is priced in the cross section of returns on CDS contracts. Recent models of intermediary asset pricing predict that proxies for the risk-bearing capacity of financial intermediaries show up as pricing factors; see, e.g., Brunnermeier and Pedersen (2009), He and Krishnamurthy (2013), and Kondor and Vayanos (2014). As argued above, our liquidity factor reflects fluctuations in the risk-bearing capacity of CDS market intermediaries, so we view it as a plausible candidate for a priced risk factor. In particular, given that dealers/intermediaries are mostly net sellers of credit protection on single-name CDSs, we expect that, from a protection seller's perspective, liquidity betas are positive (realized excess returns on CDS contracts correlate positively with the liquidity factor), the price of liquidity risk is positive (expected excess returns are higher on contracts with higher liquidity exposures), and factor prices of risk increase when intermediary constraints tighten.

1.3.1 Asset Pricing Model

We apply a parsimonious factor pricing model in which two systematic risk factors determine CDS returns: default risk and liquidity risk. This may seem overly simplistic in light of the multitude of risk factors brought forward in the empirical asset pricing literature in recent




Figure 1.3: Liquidity Factor.

The figure displays one-week excess returns (in %) on the tradable liquidity factor. The time series consists of 279 weekly observations from October 4, 2006 to February 1, 2012. Dotted vertical lines correspond to (from left to right) the collapse of two Bear Stearns structured-credit hedge funds on June 20, 2007, the Bear Stearns near-bankruptcy on March 17, 2008, the default of Lehman Brothers on September 15, 2008, and July 1, 2011 marking the beginning of a month in which a sharp sell-off in non-core European government bonds intensified the European sovereign debt crisis.

decades, but we show below that our results are robust to adding a range of additional risk factors, including a broad stock market factor, a volatility factor, and several liquidity factors from other markets.²⁸

Realized excess returns on the CDS contract referencing entity i , $r_{i,t}^e$, are given by

$$r_{i,t}^e = \alpha_i + \beta_i^{DEF} DEF_t + \beta_i^{LIQ} LIQ_t + \epsilon_{i,t}, \quad (1.4)$$

where DEF_t denotes a default factor, LIQ_t is the tradable liquidity factor from Section 1.2.6, β_i s denote factor exposures, and $\epsilon_{i,t}$ denotes nonsystematic risk. DEF_t is the realized excess return from selling protection on the CDX.NA.IG and iTraxx Eur indices, with equal weights on the two indices. Unconditional expected excess returns are given by

$$E[r_{i,t}^e] = \beta_i^{DEF} \lambda_{DEF} + \beta_i^{LIQ} \lambda_{LIQ}, \quad (1.5)$$

where λ s denote factor prices of risk.

The factor pricing model is estimated by the standard two-step methodology in which full-sample betas are estimated in the first step and market prices of risk are estimated in the

²⁸In addition, we find that the Fama and French (1993) size and book-to-market equity factors are irrelevant for the pricing of CDSs.

Chapter 1. Liquidity Risk in Credit Default Swap Markets

second step from a single cross-sectional regression of expected excess returns on full-sample betas.²⁹ Specifically, in the first step, we estimate full-sample betas from (1.4) using weekly realized excess returns. These are computed from a protection seller's perspective assuming that CDS contracts are covered by collateral agreements and marked to market on a weekly basis. In this case, each realized excess return is the weekly change in the mark-to-market value of the contract relative to the collateral amount posted at the beginning of the weekly period. We assume that this collateral amount equals the notional of the contract resulting in an "unlevered" return.³⁰ It is standard practice in the CDS literature to work with such "unlevered" returns; see, e.g., Berndt and Obreja (2010), Bongaerts et al. (2011), and Bao and Pan (2013). It also has the advantage of making returns comparable in magnitude to returns on corporate bonds. Details of the return computation are given in Appendix A.2.

In the second step, we estimate factor prices of risk (and potentially an intercept c) from the sample counterpart to (1.5); that is,

$$\widehat{E}[r_{i,t}^e] = c + \widehat{\beta}_i^{DEF} \lambda_{DEF} + \widehat{\beta}_i^{LIQ} \lambda_{LIQ} + u_i, \quad (1.6)$$

where $\widehat{\beta}_i$ s denote the estimated factor exposures, $\widehat{E}[r_{i,t}^e]$ denotes an estimate of the unconditional expected excess return, and u_i denotes the pricing error. In the empirical asset pricing literature, unconditional expected excess returns are typically estimated by sample means of realized excess returns. However, because of the short sample period and the fact that credit events are rare and have a dramatic impact on returns when they occur, sample means of realized excess returns are very imprecise estimates of unconditional expected excess returns on CDSs. Instead, we follow Bongaerts et al. (2011) in obtaining forward-looking estimates of conditional expected excess returns by using Moody's KMV Expected Default Frequencies (EDFs) to calculate expected default losses, see Appendix A.2 for details.³¹ Unconditional expected excess returns are then estimated by the sample means of conditional expected excess returns. That is, $\widehat{E}[r_{i,t}^e]$ in regression (1.6) is given by

$$\widehat{E}[r_{i,t}^e] = \frac{1}{T} \sum_{s=1}^T \widehat{E}_s[r_{i,s+1}^e], \quad (1.7)$$

where T denotes the sample size.

We emphasize that the accuracy of the conditional expected excess return estimates depends on EDFs being accurate estimates of conditional default probabilities. EDFs are unbiased estimates of average default rates that are based on a structural model in spirit of Merton

²⁹Kan et al. (2013) give a recent exposition of the two-step methodology and note that "some studies allow $\widehat{\beta}$ to change throughout the sample period.... It has become more customary in recent decades to use full-period beta estimates for portfolios formed by ranking ... on various characteristics." We follow this approach.

³⁰In reality, most counterparties post collateral amounts that are smaller than contract notionals.

³¹Our approach is similar in spirit to Campello, Chen, and Zhang (2008) who estimate factor pricing models for equity returns. They estimate betas using realized returns and run cross-sectional regressions using forward-looking estimates of conditional expected excess returns.

(1974), book values of debt, and market values of equity.³² As shown in Duffie, Saita, and Wang (2007), there exist econometric specifications of conditional default probabilities that have marginally higher predictive power than EDFs. However, EDFs have the advantage of being readily available for reference names in our sample and widely used in practice. As such, they are part of the information set of most market participants. Compared to credit ratings, EDFs adjust faster to new information and, consequently, have superior predictive power. Recent studies including Korablev and Qu (2009) and Crossen and Zhang (2011) confirm the performance of EDFs for predicting defaults during both the financial crisis and the pre-crisis period.

The empirical setup necessitates several adjustments to the standard errors in regression (1.6). First, an errors-in-variables (EIV) adjustment arising from betas being estimated. Second, an adjustment for heteroscedastic and autocorrelated errors. Third, an adjustment for potential model misspecification, arising from the possibility that, even in population, there is no combination of λ s such that Equation (1.5) is satisfied. To make these three adjustments, we use the approach of Kan et al. (2013). Adjusting standard errors for potential model misspecification allows one to draw inference on the relation between betas and expected excess returns in cases where betas do not explain the entire cross-sectional variation in expected excess returns. As shown in Kan et al. (2013), ignoring potential model misspecification typically leads to an overly positive assessment of the performance of an asset pricing model and the significance with which risk factors are priced (see also Gospodinov, Kan, and Robotti (2014) and the discussion in Ludvigson (2013)). Details of the standard error computation are provided in Appendix A.5. For comparison, we also report t -statistics based on generalized method of moments standard errors, which adjust for EIV as well as heteroscedastic and autocorrelated errors but not for potential model misspecification.

1.3.2 Data and Portfolio Construction

Data

The daily data that we use in the construction of our sample come from Markit, Bloomberg, and Moody's Analytics and extend from June 1, 2006 to February 1, 2012. From Markit, we collect five-year composite mid CDS spreads, consensus expected recovery rates, and the average credit rating by Moody's and S&P for all companies domiciled in North America and Europe. We focus on CDS contracts written on senior unsecured debt and denominated in either EUR or USD.³³ From Bloomberg, we obtain composite bid and ask CDS spreads, with the matching of CDS contracts from the two sources based on the reference entities'

³²Berndt, Douglas, Duffie, Ferguson, and Schranz (2005) and Vassalou and Xing (2004) give more detailed accounts of the methodology on which EDFs are based.

³³We select those contract terms that, on a given date, were the market standard. That is, for EUR denominated contracts, we select the modified-modified restructuring clause and for USD denominated contracts referencing high-yield names, we select the no restructuring clause. For USD denominated contracts referencing investment-grade names, we select the modified restructuring clause prior to 2009 and the no restructuring clause thereafter, in order to account for a change in the market standard resulting from ISDA's "Big Bang" Protocol.

Chapter 1. Liquidity Risk in Credit Default Swap Markets

six-digit Reference Entity Database (RED) codes and the currency denominations. From Moody's Analytics, we obtain one-year and five-year EDFs for all public companies that are contained in the Markit database. Thus, our sample consists of North American and European reference names with data coverage by each of the three providers. Credit events in our sample are identified from settlement auctions of CDSs, and we collect credit event data from the corresponding settlement protocols and auction results.³⁴

Because the key ingredients to our asset pricing tests, namely realized and expected CDS excess returns, are inferred from mid CDS spreads, we filter those for stale quotes. A quote is classified as stale, once it does not change over five or more consecutive trading days. In this case, only the spread quotation on the first of the consecutive days is retained in the sample, while the remaining ones are excluded.

From the collected data, we compute weekly time series of realized and conditional expected excess returns, bid-ask spreads, and price impact measures. Due to a considerable number of missing bid-ask spreads, we use weekly averages of bid-ask spreads instead of end-of-period observations. The price impact measure is constructed as in Section 1.2.5, with the exception that we average absolute spread changes per contributed quote over one-week as opposed to one-month periods. Weekly observations are sampled on Wednesdays and we exclude all entities with less than fifty joint observations. This leaves a sample of 666 reference entities, of which 426 are domiciled in North America and 240 in Europe, and a total of 144,163 joint observations.

Portfolio Construction

Because individual-asset betas are usually very imprecisely estimated, we conduct our analysis on a set of 40 equally-weighted portfolios rather than at the level of individual CDSs. Portfolios are rebalanced at a quarterly frequency and formed such that they exhibit variation across the default risk and liquidity dimensions.

The portfolio formation is as follows: on month-ends of March, June, September, and December of a given year, we first sort reference names from best to worst credit quality according to the average issuer credit rating over the previous quarter, and then group them into five credit rating categories: AAA-AA, A, BBB, BB, and B-CCC. Subsequently, we sort reference names within a given default risk group from most liquid to least liquid either according to the average bid-ask spread over the previous quarter or according to the average price impact over the previous quarter. In both cases, we group reference names into illiquidity quartiles.

Because the first quarter of data is used for portfolio formation, this procedure yields portfolio time series from October 11, 2006 to February 1, 2012. During this period, we find two weeks in which only a small number of North American reference names have quoted bid-ask spreads.

³⁴Creditex and Markit administrate credit event auctions and publish auction results on www.creditfixings.com. Settlement protocols are published by the ISDA.

	DEF	LIQ
Mean	0.46	5.13
Standard Deviation	41.78	16.74
Skewness	-0.43	0.98
Kurtosis	4.93	14.85
ρ_1	-0.13	0.01

Table 1.4: Descriptive Statistics of Factors.

The table displays descriptive statistics for the default and liquidity factors. Mean and standard deviation are in basis points per week, and ρ_1 denotes first-order autocorrelation. Factor time series consist of 276 weekly observations from October 11, 2006 to February 1, 2012.

We exclude the corresponding portfolio observations from the analysis, leaving a total of 276 one-week periods during the sample period.

1.3.3 Results

Descriptive Statistics

Table 1.4 displays descriptive statistics for the two factors. During our sample period, the average realized excess return on the default factor is positive, but not statistically significant. In contrast, the average realized excess return on the liquidity factor is positive and significant, with a Newey and West (1987) t -statistic of 5.06. The two factors are virtually uncorrelated, with a correlation coefficient of -0.01.

Table 1.5 displays descriptive statistics for the 20 portfolios formed by first sorting CDS contracts according to credit ratings and then according to bid-ask spreads. Descriptive statistics for the remaining 20 portfolios are provided in Table A.1 at the end of Appendix A. Sample means of expected excess returns are positive across portfolios and strongly significant with Newey and West (1987) t -statistics between 4.13 and 11.14.³⁵ This reflects the fact that risk neutral default probabilities, on average, exceed physical default probabilities. Expected excess returns tend to increase with portfolio illiquidity and deteriorating credit quality. For instance, among the bid-ask-spread-sorted portfolios, we observe a difference of 5.45% per year in the expected excess return between a portfolio consisting of the most illiquid low-credit-quality CDSs (B-CCCQ4) and a portfolio consisting of the most liquid high-credit-quality CDSs (AAA-AAQ1). Sample means of realized excess returns are not significantly different from zero, underscoring the importance of using forward-looking information when estimating expected

³⁵Overall, the portfolios exhibit ex-post the properties they were chosen to reflect ex-ante with CDS spreads (bid-ask spreads) of the bid-ask-spread-sorted portfolios increasing from 45 bps (4 bps) for the portfolio consisting of the most liquid high-credit-quality CDSs to 1710 bps (109 bps) for the portfolio consisting of the most illiquid low-credit-quality CDSs. Similarly, CDS spreads (price-impact measures) of the price-impact-sorted portfolios increase from 42 bps (0.12 bps per contributed quote) for the portfolio consisting of the most liquid high-credit-quality CDSs to 1701 bps (8.86 bps per contributed quote) for the portfolio consisting of the most illiquid low-credit-quality CDSs.

Chapter 1. Liquidity Risk in Credit Default Swap Markets

Credit Rating	Expected Excess Returns				Realized Excess Returns			
	Bid-Ask Spread				Bid-Ask Spread			
	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4
AAA-AA	0.35 [6.75]	0.40 [7.54]	0.40 [6.22]	0.64 [6.27]	-0.74 [-1.06]	-0.25 [-0.33]	-0.49 [-0.52]	0.22 [0.14]
A	0.38 [7.70]	0.44 [7.17]	0.53 [5.81]	0.97 [4.13]	-0.62 [-0.80]	-0.59 [-0.60]	-0.18 [-0.12]	0.85 [0.25]
BBB	0.53 [9.61]	0.69 [7.27]	0.92 [6.42]	1.58 [4.95]	-0.73 [-0.78]	-0.43 [-0.29]	0.38 [0.17]	2.22 [0.48]
BB	1.32 [9.96]	1.75 [9.50]	2.19 [8.84]	3.07 [8.48]	-1.34 [-0.59]	0.50 [0.17]	2.12 [0.47]	6.50 [0.91]
B-CCC	2.98 [9.89]	2.94 [11.14]	4.17 [6.92]	5.80 [5.19]	-2.54 [-0.45]	6.72 [0.85]	6.67 [0.66]	25.38 [1.41]
Credit Rating	CDS Spreads				Standard Deviations			
	Bid-Ask Spread				Bid-Ask Spread			
	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4
AAA-AA	0.45	0.52	0.74	1.20	1.33	1.75	3.11	4.38
A	0.53	0.67	0.86	1.97	1.55	1.99	2.57	5.51
BBB	0.74	1.00	1.35	2.68	1.98	2.57	3.33	5.76
BB	1.96	2.64	3.54	5.69	4.74	5.81	7.32	9.56
B-CCC	4.70	6.06	9.33	17.10	9.82	18.01	17.74	27.63

Table 1.5: Descriptive Statistics of Bid-Ask-Spread-Sorted Portfolios.

The table displays descriptive statistics for the 20 portfolios formed by first sorting CDS contracts according to credit ratings and then according to bid-ask spreads. The upper part of the table reports sample means of conditional expected excess returns (in % per year) and realized excess returns (in % per year). In brackets are *t*-statistics based on Newey and West (1987) heteroscedasticity and autocorrelation consistent standard errors with 24 lags. The lower part of the table reports sample means of average five-year CDS spreads across portfolio constituents (in % per year) and standard deviations of realized excess returns (in % per year). Portfolio time series consist of 276 weekly observations from October 11, 2006 to February 1, 2012.

excess returns. Standard deviations of realized excess returns also increase with portfolio illiquidity and deteriorating credit quality, and the resulting unconditional and forward-looking annualized Sharpe ratios lie in a reasonable range from 0.13 to 0.32.

Factor Exposures

First-step regression results are displayed in Table 1.6. For ease of interpretation, instead of reporting the raw beta estimates, we report the product of the beta estimates and the standard deviations of the respective factors. Consequently, the table shows weekly realized portfolio excess returns (in bps) in response to a one standard deviation shock to each of the factors.

Default betas are positive and statistically significant throughout portfolios and almost monotonically increasing along both the liquidity and credit quality dimensions. Default risk is economically important with a one standard deviation shock to the default factor having an

1.3. Pricing of Liquidity Risk

Panel A: Bid-Ask-Spread-Sorted Portfolios								
Credit Rating	DEF				LIQ			
	Bid-Ask Spread				Bid-Ask Spread			
	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4
AAA-AA	13.13 [11.78]	19.85 [8.83]	34.10 [7.08]	46.41 [16.00]	4.86 [2.30]	2.91 [1.16]	3.25 [0.95]	5.39 [1.15]
A	17.36 [11.05]	22.20 [18.61]	30.45 [12.50]	59.52 [10.55]	5.30 [1.97]	8.25 [2.26]	8.92 [2.33]	14.21 [1.80]
BBB	22.65 [16.55]	28.45 [17.03]	36.81 [13.03]	60.87 [9.36]	8.19 [3.42]	11.61 [2.67]	13.30 [2.46]	19.03 [2.19]
BB	50.93 [8.02]	63.01 [11.10]	77.22 [8.56]	92.88 [9.57]	15.75 [5.95]	20.89 [3.78]	25.92 [2.89]	32.20 [3.98]
B-CCC	75.70 [6.02]	106.22 [8.38]	155.65 [8.44]	198.60 [4.93]	58.40 [3.24]	26.48 [1.92]	46.38 [3.66]	80.51 [1.89]

Panel B: Price-Impact-Sorted Portfolios								
Credit Rating	DEF				LIQ			
	Price Impact				Price Impact			
	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4
AAA-AA	10.52 [12.99]	19.94 [8.00]	30.35 [9.63]	51.12 [11.99]	4.20 [2.31]	4.61 [1.76]	4.17 [1.25]	2.93 [0.64]
A	13.05 [13.51]	20.02 [11.40]	29.63 [11.37]	68.99 [12.49]	4.96 [2.31]	7.21 [2.26]	11.60 [2.28]	12.50 [1.65]
BBB	17.10 [11.34]	24.98 [14.53]	41.10 [11.92]	66.26 [12.78]	6.17 [2.51]	11.46 [2.76]	13.55 [2.46]	21.59 [2.47]
BB	35.63 [16.55]	66.15 [11.75]	82.18 [10.88]	108.10 [8.99]	10.64 [3.93]	22.22 [3.55]	32.03 [3.54]	26.98 [3.96]
B-CCC	71.54 [8.88]	116.00 [11.08]	142.19 [7.65]	204.54 [4.47]	51.36 [3.22]	45.45 [3.53]	41.84 [3.22]	63.12 [1.50]

Table 1.6: Results of Time-Series Regressions.

The table displays first-step regression results at the level of individual portfolios. Reported are beta estimates times the standard deviation of the respective factors; i.e., the weekly realized portfolio excess returns (in basis points) in response to a one standard deviation shock to the factors. In brackets are t -statistics based on Newey and West (1987) heteroscedasticity and autocorrelation consistent standard errors with 24 lags. Time series consist of 276 weekly observations from October 11, 2006 to February 1, 2012.

impact on portfolio excess returns between 11 bps and 205 bps.

Liquidity betas are positive throughout portfolios and statistically significant at the five-percent level for 29 out of the 40 portfolios. Liquidity betas also tend to increase along the liquidity and credit quality dimensions. However, especially along the liquidity dimension there are exceptions indicating that portfolios with higher bid-ask spreads or price impact measures do not necessarily exhibit higher liquidity risk. Liquidity risk is also economically important with a one standard deviation shock to the liquidity factor having an impact on portfolio excess returns between 3 bps and 81 bps. Unreported adjusted R^2 s of the regressions range from 19% to 78% across portfolios.³⁶

³⁶Unreported results for nested one-factor specifications of regression (1.4) show that, on their own, each of the

Chapter 1. Liquidity Risk in Credit Default Swap Markets

Spec.	1	2	3	4	5	6
c				-0.39 (-0.86) [-0.91]	0.25 (0.61) [0.50]	-0.34 (-1.01) [-1.03]
λ_{DEF}	2.23 (4.70) [4.84]		1.43 (6.26) [5.49]	2.40 (3.59) [3.70]		1.60 (4.20) [4.06]
λ_{LIQ}		2.53 (4.30) [4.47]	0.94 (4.82) [2.77]		2.42 (3.86) [4.00]	0.92 (5.07) [2.71]
R^2	0.94 [0.89,0.99]	0.88 [0.68,1.00]	0.97 [0.95,0.99]	0.95 [0.91,0.98]	0.88 [0.70,1.00]	0.98 [0.96,0.99]

Table 1.7: Results of Cross-Sectional Regressions.

The table displays results of several specifications of the second-step regression. Specifications of $\hat{E}[r_{i,t}^e] = c + \hat{\beta}_i^{DEF} \lambda_{DEF} + \hat{\beta}_i^{LIQ} \lambda_{LIQ} + u_i$ are estimated from expected excess returns and beta estimates inferred from time series that consist of 276 weekly observations from October 11, 2006 to February 1, 2012. Reported are factor price of risk estimates (in basis points), t -statistics based on asymptotic generalized method of moments standard errors that account for error-in-variables problems (in parenthesis), t -statistics based on Kan, Robotti, and Shanken's (2013) asymptotic standard errors that account for error-in-variables problems and potential model misspecification (in brackets), cross-sectional R^2 s, and their 95% confidence intervals. Standard errors are heteroscedasticity and autocorrelation consistent through the use of Newey and West's (1987) method with 24 lags.

Factor Prices of Risk

Second-step regression results for alternative specifications of the cross-sectional regression (1.6) are displayed in Table 1.7. The table shows regression coefficients with t -statistics that account for EIV and heteroscedastic and autocorrelated errors reported in parentheses, and t -statistics that in addition account for potential model misspecification reported in brackets. The table also shows cross-sectional R^2 s with 95% confidence intervals in brackets. These confidence intervals are computed along the lines of Kan et al. (2013), with details deferred to Appendix A.5.

Factors carry significant prices of risk in one- and two-factor models regardless of whether standard errors are adjusted for potential model misspecification or not. For instance, in the two-factor specification with an imposed zero-intercept restriction (specification 3), the most conservative t -statistics are 5.49 and 2.77 for the default and liquidity factors, respectively. Restricting intercepts to zero is inconsequential because they are small in magnitude and not statistically significant (see specifications 4–6). In other words, model specifications cannot be rejected based on the average abnormal expected returns that they produce. The positive sign on the price of liquidity risk implies that sellers of credit protection earn higher expected excess returns on contracts with higher liquidity exposures. Cross-sectional R^2 s are substantial across model specifications, which is, in part, a consequence of using less-noisy and forward-
two factors constitutes a significant explanatory variable of CDS portfolio excess returns.

1.3. Pricing of Liquidity Risk

Spec.	BM	1	2	3	4	5	6	7	8	9
Def. Risk	3.30 [1.08]	2.69 [0.88]	2.72 [0.89]	5.13 [1.68]	3.59 [1.17]	2.92 [0.97]	3.35 [1.11]	3.12 [1.02]	1.75 [0.80]	2.78 [0.88]
Liq. Risk	2.21 [0.59]	2.04 [0.55]	2.20 [0.59]		2.14 [0.59]	1.59 [0.43]	2.22 [0.59]	2.02 [0.54]	1.61 [0.46]	2.21 [0.59]
Charact./ Add. Risk		1.14 [0.21]	0.93 [0.15]	0.00 [0.00]		1.32 [0.24]	-0.13 [-0.02]	0.52 [0.09]	2.49 [0.35]	0.33 [0.18]

Table 1.8: Decompositions of Expected Excess Returns.

The table displays decompositions of expected excess returns for the benchmark model specification and the robustness checks. Annualized expected excess returns are decomposed into contributions of characteristics/additional factor risk premia. Reported are the contributions of these components (in % per year) to the difference in the expected excess return between the bid-ask-spread-sorted B-CCCQ4 and AAA-AAQ1 portfolios and, in brackets, the contributions of these components to the average expected excess return across the 40 portfolios. Specification identifiers are given in the second row of the table.

looking information when estimating expected excess returns. For instance, our preferred model specification (specification 3; henceforth, the benchmark model or BM specification) has a cross-sectional R^2 of 0.97, and the specification test of Kan et al. (2013) cannot reject the null hypothesis $H_0 : R^2 = 1$ (the p -value is 0.21).

To assess the economic importance of the risk factors, we use the benchmark model specification to decompose the annualized expected excess return on each portfolio into default and liquidity risk premia (defined as $52 \times \hat{\beta}_i^F \times \hat{\lambda}_F$, $F \in \{DEF, LIQ\}$). We summarize the results in two ways. First, we consider the contributions of these components to the difference in the expected excess return between the two extreme (bid-ask-spread-sorted) portfolios B-CCCQ4 and AAA-AAQ1. We use the term *expected return differential* to refer to this difference. Second, we consider the contributions to the average expected excess return across all portfolios.

The first column of Table 1.8 shows the contributions to the expected return differential and, in brackets, the contributions to the average expected excess return. Default and liquidity risk contribute 3.30% and 2.21% per year, respectively, to the expected return differential (the remainder is a pricing error). In case of the average expected excess return, default and liquidity risk contribute 1.08% and 0.59% per year, respectively. That liquidity risk is economically important is in contrast to the findings in Bongaerts et al. (2011) and is addressed in more detail below.

Financial Intermediaries and the Pricing of Risk

We now study time-variation in factor prices of risk and the relation to the risk-bearing capacity of financial intermediaries. On each observation date we estimate the following conditional version of Equation (1.6):

$$\hat{E}_t[r_{i,t+1}^e] = \hat{\beta}_i^{DEF} \lambda_{DEF,t} + \hat{\beta}_i^{LIQ} \lambda_{LIQ,t} + u_{i,t}, \quad (1.8)$$

Chapter 1. Liquidity Risk in Credit Default Swap Markets




Figure 1.4: Factor Prices and Intermediary Equity Capital.

The figure displays three-month centered moving averages of time-varying factor price of risk estimates (black lines, left hand scales) and the negative value of intermediary equity capital (gray lines, right hand scales). Panel A displays the factor price of default risk and Panel B displays the factor price of liquidity risk. The factor price of risk estimates are in basis points and obtained by cross-sectional regressions of conditional expected excess returns on full sample beta estimates. Intermediary equity capital is in 100 billion USD and given by the aggregate market capitalization of financial institutions that make up the G14 group of major credit derivatives dealers. The time series consist of 276 weekly observations between October 11, 2006 and February 1, 2012. Dotted vertical lines correspond to (from left to right) the collapse of two Bear Stearns structured-credit hedge funds on June 20, 2007, the Bear Stearns near-bankruptcy on March 17, 2008, the default of Lehman Brothers on September 15, 2008, and July 1, 2011 marking the beginning of a month in which a sharp sell-off in non-core European government bonds intensified the European sovereign debt crisis.

where $\lambda_{i,t}$ s denote conditional factor prices of risk.³⁷ Figure 1.4 displays time series of the resulting factor prices of risk (black lines). Panel A shows the price of default risk, while Panel B shows the price of liquidity risk, and for expositional purposes we display three-month moving averages that smooth out higher frequency fluctuations. The magnitudes of both factor prices are particularly high in the aftermath of the Lehman Brothers default and the AIG bailout. Other periods during which the factor prices of default and liquidity risk are high include the initial phase of the financial crisis as well as the last part of the sample period when the European sovereign debt crisis intensified.

Many recent models of intermediary asset pricing imply that prices of risk (in absolute values) correlate negatively with the risk-bearing capacity of financial intermediaries. Moreover, in most models the risk-bearing capacity of the intermediary sector correlates positively with its equity capital. For instance, in the models of He and Krishnamurthy (2013) and Kondor and Vayanos (2014), intermediaries have wealth-dependent (effective) risk aversion and their risk-bearing capacity is increasing in equity capital. Alternatively, in the models of Gromb and Vayanos (2002) and Brunnermeier and Pedersen (2009), intermediaries face funding

³⁷ Note that because betas in Equation (1.8) are fixed at their full sample estimates, sample means of factor price of risk estimates in the conditional model coincide with those in the unconditional model; that is $\lambda_{DEF} = \frac{1}{T} \sum_{t=1}^T \lambda_{DEF,t}$ and $\lambda_{LIQ} = \frac{1}{T} \sum_{t=1}^T \lambda_{LIQ,t}$.




Figure 1.5: CDS Spread Decomposition of Bid-Ask-Spread-Sorted Portfolios.

The figure displays five-year CDS spreads (in % per year) of the bid-ask-spread-sorted portfolios. CDS spreads are decomposed into expected default losses, factor risk premia, and pricing errors implied by the benchmark model specification. The horizontal axis displays portfolio identifiers.

constraints and their risk-bearing capacity is increasing in equity capital and decreasing in margin requirements.

To investigate the impact of intermediary frictions on the pricing of CDS contracts, we proxy the risk-bearing capacity of the intermediary sector by the market capitalization of the financial institutions that make up the G14 group of major credit derivatives dealers, see Section 1.2.5. In addition to factor prices of risk, Figure 1.4 also displays the negative value of the G14 market capitalization (gray lines). Evidently, periods of low intermediary equity capital coincide with high factor prices of risk. Indeed, the time-series correlations between intermediary equity capital and the prices of default and liquidity risk are -0.71 and -0.32, respectively. This is consistent with intermediary frictions affecting the pricing of CDS contracts and with recent work by Siriwardane (2015) who shows in a post-crisis sample that capital losses for sellers of credit protection lead to wider CDS spreads.

Decomposing CDS Spreads

As an alternative illustration of economic importance, we decompose CDS spreads instead of expected excess returns. The CDS spread of each portfolio is decomposed into components due to default and liquidity risk, as well as an additional component reflecting the expected default loss. The decomposition is based on the conditional model described in the previous section, and computational details are given in Appendix A.3. Figure 1.5 displays the resulting decomposition for the 20 bid-ask-spread-sorted portfolios. The corresponding figure for the

Chapter 1. Liquidity Risk in Credit Default Swap Markets

20 price-impact-sorted portfolios can be found at the end of Appendix A.

We first consider the extreme (bid-ask-spread-sorted) portfolios. For the AAA-AAQ1 portfolio, the average CDS spread is 45 bps of which expected default losses and the default risk premium account for 7 bps (17%) and 25 bps (56%), respectively, while the liquidity risk premium accounts for 15 bps (34%) (the remaining -3 bps is a pricing error). At the other end of the spectrum is the B-CCCQ4 portfolio with an average CDS spread of 1710 bps of which expected default losses and the default risk premium account for 830 bps (49%) and 537 bps (31%), respectively, while the liquidity risk premium accounts for 347 bps (20%) (here the pricing error is -4 bps). Averaging the relative contributions across all portfolios, we find that expected default losses and the default risk premium account for 29% and 47% of model-implied CDS spreads, respectively, while the liquidity risk premium accounts for 24%.

That expected default losses only account for a relatively small fraction of credit spreads, in particular for highly rated firms, is well known; see, e.g., Elton, Gruber, Agrawal, and Mann (2001) and Driessen (2005) for corporate bond yield spreads, and Berndt et al. (2005) for CDS spreads. There is some disagreement concerning the size of the default risk premium with some papers using structural credit risk models to argue that only relatively small default risk premia are consistent with historical default records and equity risk premia (again, in particular for highly rated firms; see, e.g., Huang and Huang (2012)), while other papers, mainly using reduced-form credit risk models, argue that default risk premia are sizable. Our analysis indicates that the default risk premium is the largest component of CDS spreads. Most importantly, however, we find evidence for a sizable liquidity risk premium.

1.3.4 Robustness Checks

We conduct a range of robustness checks; we control for the contract-specific level of illiquidity, consider an alternative notion of liquidity risk suggested by Bongaerts et al. (2011), investigate an alternative construction of our tradable liquidity factor, and include additional risk factors in the asset pricing model. For each robustness check, first-step regression results are summarized in the text, second-step regression results are reported in Table 1.9, and results of the expected excess return decomposition are reported in Table 1.8.³⁸ We only report results for additional factors that are likely to capture similar effects as our liquidity factor. Because our two measures of economic importance typically give similar results, we only comment on the expected return differential.

Contract-Specific Level of Illiquidity

A number of studies have shown that CDS spreads increase with the contract-specific level of illiquidity; see, e.g., Tang and Yan (2007), Bühler and Trapp (2009), Bongaerts et al. (2011),

³⁸First-step regression results are available upon request. In the text, we summarize the number of portfolios that load significantly on each additional/alternative factor as well as the sign of betas. We also remark if betas with respect to the default and liquidity factors change significantly upon inclusion of an additional factor.

Qiu and Yu (2012), Lesplingart et al. (2012), and Pires et al. (2014). Therefore, we control for the contract-specific level of illiquidity when assessing liquidity risk. Specifically, we separately add bid-ask spreads and the price impact measure as portfolio characteristics to the second-step regression. Unreported results show that, on their own, both bid-ask spreads and the price impact measure are significantly and positively related to expected excess returns, corroborating findings in previous papers. However, in conjunction with default and liquidity betas neither bid-ask spreads nor the price impact measure are significantly related to expected excess returns (see specifications 1 and 2 in Tables 1.8 and 1.9, respectively). As such, it appears that default and liquidity risk largely subsume the effect of the contract-specific level of illiquidity. The statistical significance and economic importance of liquidity risk are very similar to the benchmark case.

Alternative Notion of Liquidity Risk

Bongaerts et al. (2011) also investigate the relative importance of contract-specific illiquidity (measured by the level of transaction costs) and liquidity risk in the cross section of expected excess returns on CDS contracts. However, they consider a notion of liquidity risk that is very different from ours, namely covariation between innovations to contract-specific transaction costs and the return on a nontraded default factor. When this covariance is negative—as it is the case empirically—transaction costs rise in states with high aggregate default risk and unwinding hedge positions becomes more expensive. This makes CDS contracts less effective hedges against default risk and should lead to less demand for credit protection and lower expected excess returns for credit protection sellers. However, the empirical analysis in Bongaerts et al. (2011) reveals that the premium associated with this notion of liquidity risk is economically negligible. To investigate if this result also holds true in our more recent and broader sample of CDSs, we replace, in the cross-sectional regression, betas capturing our notion of liquidity risk with betas capturing their notion (see specification 3 in Tables 9 and 10).³⁹ The price associated with this alternative notion of liquidity risk is both statistically insignificant and economically negligible, confirming the results of Bongaerts et al. (2011).⁴⁰ A possible reason for the lack of importance of Bongaerts et al.'s (2011) notion of liquidity risk may be that the majority of CDS contracts are marked to market on a daily basis. This implies that protection buyers realize gains when aggregate default risk increases without having to unwind their positions and incur transaction costs.

Alternative Construction of CDS Market Liquidity Factor

We consider an alternative construction of the tradable liquidity factor in which the excess re-

³⁹ Specifically, we follow Bongaerts et al. (2011) in estimating single-factor betas of bid-ask spread innovations with respect to the default factor. We use Bongaerts et al.'s (2011) time-series model of liquidity to compute bid-ask spread innovations and orthogonalize innovations for stock market returns (i.e., returns on the “nonhedge” asset in their terminology). These liquidity betas have the expected negative sign but are insignificant for most portfolios.

⁴⁰ The same result obtains when adding the alternative liquidity betas to the benchmark model specification in which case the pricing of our notion of liquidity risk is virtually unaffected by the additional betas.

Spec.	1	2	3	4	5	6	7	8	9
λ_{DEF}	1.16 (5.55) [3.48]	1.18 (6.31) [3.35]	2.23 (4.02) [4.37]	1.54 (6.04) [6.44]	1.30 (6.90) [4.57]	1.48 (6.67) [5.15]	1.35 (8.66) [3.74]	1.27 (9.16) [4.13]	1.27 (5.91) [4.54]
λ_{LIQ}	0.87 (5.53) [2.14]	0.94 (4.96) [2.19]		1.13 (5.35) [3.67]	0.70 (3.15) [1.53]	0.95 (5.02) [3.06]	0.87 (5.46) [2.32]	0.79 (4.57) [2.23]	0.94 (5.04) [3.72]
λ_X	207.35 (0.86) [0.85]	2152.72 (0.86) [0.85]	-0.87 (-0.01) [-0.01]		-0.05 (-1.61) [-1.05]	0.56 (0.35) [0.24]	-4.21 (-0.42) [-0.27]	13.95 (2.61) [2.31]	-15.60 (-5.07) [-3.50]
R^2	0.97 [0.95,1.00]	0.98 [0.95,1.00]	0.94 [0.89,0.99]	0.97 [0.95,1.00]	0.98 [0.96,1.00]	0.97 [0.95,0.99]	0.97 [0.95,0.99]	0.98 [0.95,1.00]	0.98 [0.96,1.00]

Table 1.9: Robustness of Cross-Sectional Regressions.

The table displays results of a series of robustness checks. Specifications of $\hat{E}[r_{i,t}^e] = \hat{\beta}_i^{DEF} \lambda_{DEF} + \hat{\beta}_i^{LIQ} \lambda_{LIQ} + \bar{X}_i \lambda_X + u_i$ (specifications 1 and 2) and $\hat{E}[r_{i,t}^e] = \hat{\beta}_i^{DEF} \lambda_{DEF} + \hat{\beta}_i^{LIQ} \lambda_{LIQ} + \hat{\beta}_i^X \lambda_X + u_i$ (specifications 3–9) are estimated from expected excess returns, beta estimates, and sample means of characteristics inferred from time series that consist of 276 weekly observations from October 11, 2006 to February 1, 2012. Specification identifiers are given in the second row of the table. Reported are factor price of risk estimates (in basis points), t -statistics based on asymptotic generalized method of moments standard errors that account for error-in-variables problems (in parenthesis), t -statistics based on Kan, Robotti, and Shanken’s (2013) asymptotic standard errors that account for error-in-variables problems and potential model misspecification (in brackets), cross-sectional R^2 -s, and their 95% confidence intervals. Standard errors are heteroscedasticity and autocorrelation consistent through the use of Newey and West’s (1987) method with 24 lags.

turns on the individual index arbitrage strategies are weighted by the inverse of their conditional volatilities (see specification 4 in Tables 1.8 and 1.9). This is akin to the construction of the time series momentum factor in Moskowitz, Ooi, and Pedersen (2012), see Appendix A.4 for details. The resulting factor is very highly correlated with our original liquidity factor and we obtain results that are similar to those in the benchmark case.

Additional Factors

The Hu, Pan, and Wang (2013) “Noise” measure. Hu et al. (2013) argue that their “Noise” measure is a broad illiquidity measure that captures the availability of arbitrage capital. Therefore, we include innovations to their “Noise” measure as an additional factor in the model (see specification 5 in Tables 1.8 and 1.9).⁴¹ Seventeen portfolios load significantly on this factor, all with a negative sign. However, the factor price of risk is not statistically significant, while the price of CDS market liquidity risk remains statistically significant when ignoring potential model misspecification. This is consistent with the analysis in Section 1.2.5, which identified the “Noise” measure as the variable that is most strongly related to variation in CDS market illiquidity. In terms of economic importance, the “Noise” factor contributes 1.32% per year to the expected return differential, while the contribution of the CDS market liquidity factor is reduced somewhat to 1.59%.

Corporate bond market illiquidity factor. Lin et al. (2011) find that exposure to market-wide corporate bond liquidity is priced in the cross section of corporate bond returns. Given the relation between CDS spreads and corporate bond yields, such exposure may also be priced in the cross section of CDS returns. Therefore, we include innovations to a corporate bond market illiquidity measure as an additional factor in the model (see specification 6 in Tables 1.8 and 1.9). The corporate bond market illiquidity measure is an aggregate of bond-specific Amihud (2002) illiquidity measures, see Appendix A.4 for details.⁴² Seventeen portfolios load significantly on the corporate bond market illiquidity factor, all with a negative sign. However, the factor price of risk is neither statistically significant nor economically important, while results for CDS market liquidity risk are almost identical to the benchmark case.

Stock market illiquidity factor. Both Acharya, Amihud, and Bharath (2013) and Bongaerts, de Jong, and Driessen (2012) find that exposure to aggregate liquidity in the stock market is priced in the cross section of corporate bond returns. To investigate if this also holds true for CDS returns, we include innovations to a stock market illiquidity measure as an additional

⁴¹Because the “Noise” measure is very persistent, we use its innovations as a factor rather than the measure itself. Innovations are the residuals of an AR(2) model for the time series of the “Noise” measure. The reported results are not sensitive to the choice of time series model and using first-differences of the “Noise” measure gives very similar results as well. The same comments apply to the corporate bond and stock market illiquidity factors and the volatility factor that are analyzed in the robustness checks below. We do not take innovations of our liquidity factor because it exhibits virtually no autocorrelation (see Table 1.4).

⁴²At a monthly frequency, the corporate bond market illiquidity measure is highly correlated with the corporate bond market illiquidity measure of Dick-Nielsen, Feldhütter, and Lando (2012). Because Dick-Nielsen et al.’s (2012) measure is only available at a monthly frequency, we construct our own corporate bond market illiquidity measure.

Chapter 1. Liquidity Risk in Credit Default Swap Markets

factor in the model (see specification 7 in Tables 1.8 and 1.9). The stock market illiquidity measure is an aggregate of stock-specific Amihud (2002) illiquidity measures, see Appendix A.4 for details.⁴³ Twenty-one portfolios load significantly on the stock market illiquidity factor, all with a negative sign. However, the factor price of risk is statistically insignificant and of limited economic importance, while results for CDS market liquidity risk are similar to the benchmark case.

Stock market factor. We control for stock market risk by including a factor which is the excess return on an equally weighted portfolio of the S&P 500 and EURO STOXX 50 indices (see specification 8 in Tables 1.8 and 1.9). Fourteen portfolios load significantly on the stock market factor, all but one with a positive sign, and its price of risk is statistically significant and has the expected positive sign. Economically, the factor contributes 2.49% per year to the expected return differential, and reduces the contribution of the default factor substantially to 1.75%.⁴⁴ The price of liquidity risk remains statistically significant, but the contribution to the expected return differential is reduced somewhat to 1.61% per year.

Volatility factor. Several papers find that volatility is an important risk factor in asset markets; see, e.g., Ang, Hodrick, Xing, and Zhang (2006) and Bongaerts et al. (2012) for evidence from the stock and corporate bond markets, respectively. Therefore, we include innovations to the VIX index as a factor in the model (see specification 9 in Tables 1.8 and 1.9). Twelve portfolios load significantly on the volatility factor, mostly with a negative sign, and its price of risk is statistically significant and has the expected negative sign. Economically, however, the volatility factor is of limited importance. Results for liquidity risk are similar to the benchmark case.

1.4 Conclusion

We analyze whether liquidity risk is priced in the cross section of returns on CDS contracts. First, we construct a model-independent measure of CDS market illiquidity by aggregating deviations of credit index levels from their no-arbitrage values implied by the index constituents' CDS spreads. Second, based on index arbitrage strategies, we design a tradable liquidity factor that is highly negatively correlated with innovations to the CDS market illiquidity measure. Third, we define liquidity risk as covariation between CDS returns and the liquidity factor and show that liquidity risk is both statistically significant and economically important for the pricing of CDSs. In particular, liquidity risk increases CDS spreads and the expected excess returns earned by sellers of credit protection. Consistent with recent models of intermediary asset pricing, we find that illiquidity and risk premia correlate negatively with proxies for the

⁴³ Readily available measures of stock market illiquidity such as those of Pástor and Stambaugh (2003) and Sadka (2006) are only available at a monthly frequency, which is why we construct our own stock market illiquidity measure.

⁴⁴ While the price of default risk is largely unaffected, the spread in default betas across portfolios shrinks substantially when the stock market factor is included in the model.

risk-bearing capacity of CDS market intermediaries.

2 Market Structure and Transaction Costs of Index CDSs

This chapter is based on joint work with Pierre Collin-Dufresne and Anders B. Trolle in which we study the two-tiered structure of the index CDS market after the implementation of the Dodd-Frank Act. We identify dealer-to-customer (D2C) trades and interdealer (D2D) trades. Transaction costs and price impacts are larger for D2C trades and increase with trade size, quoted bid-ask spread, and volatility. D2C trades Granger-cause D2D trades consistent with the interdealer market being used for managing inventory risk. Unique order-book data show the important role of mid-market matching and workup for reducing transaction costs and price impacts of D2D trades. D2C trades are competitive relative to executable bids and offers in the interdealer market, suggesting that the market structure delivers favorable prices for customers who value immediacy.

2.1 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 classical two-tiered over-the-counter (OTC) market in which global derivatives dealers provide liquidity to their institutional customers in the dealer-to-customer (D2C) segment of the market, and dealers trade among themselves in the interdealer (D2D) segment of the market. New comprehensive 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). These regulated trading platforms are required to offer trading in order books, thus opening up

Chapter 2. Market Structure and Transaction Costs of Index CDSs

the market to all-to-all trading, in which customers could compete with dealers for liquidity provision. However, SEFs are also allowed to offer trading via request for quote (RFQ), which more closely mimics traditional trading in OTC markets. 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 (mostly via order books) run by interdealer brokers (IDBs).¹

The endurance of this bifurcated market structure 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).² On the other hand, some market participants have accused dealers of resisting a transition to an all-to-all market structure in order to preserve their “monopoly” on liquidity provision; see, e.g., Managed Funds Association (2015).³ In light of this controversy, the purpose of the paper is twofold: first, using transaction data, we provide a detailed characterization of the two-tiered market structure. Second, we analyze transaction costs and price impacts across market segments and different credit indices, and estimate dealer profits from liquidity provision.

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 transaction, 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 amount traded in the D2C segment is USD 9.843 billion and USD 3.705 billion for CDX.IG and CDX.HY, respectively. In the D2D segment, the corresponding numbers are USD 1.354 billion and USD 0.402 billion. Outright

¹ Referring to both the index CDS and the interest rate swap markets, a recent article summarized the current situation as “...dealer banks still trade together privately in one segment of the market and the buy side 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 and Loss*, October 22, 2015.

² Even for the most liquid index CDSs, there are often not more than 100 trades per day and contract notional amounts are frequently in excess of USD 100 million.

³ Dealers are confronted with similar accusations pertaining to the single-name CDS market in which they allegedly conspired to shut down emerging all-to-all trading venues. Recently, global derivatives dealers agreed to settle a civil lawsuit brought by a group of institutional investors for USD 1.87 billion (see “Banks Near Pact on Swaps Suit,” *Wall Street Journal*, September 12, 2015). Investigations by U.S. and European antitrust authorities are ongoing as well.

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 executed on SEFs.

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. 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 trade types are the focus of the paper.

We measure transaction cost as the difference between the transaction price and the contemporaneous value of Markit's intraday mid-quote (the effective half-spread). 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 significantly higher than those of D2D trades. For CDX.IG, average transaction costs are 0.137 basis points (bps) and 0.088 bps for D2C and D2D trades, respectively. The corresponding numbers for CDX.HY are 0.674 bps and 0.402 bps, respectively. The differences in transaction costs are mostly due to D2C trades having larger price impacts than D2D trades. For CDX.IG, average price impacts are 0.106 bps and 0.063 bps for D2C and D2D trades, respectively. The corresponding numbers for CDX.HY are 0.508 bps and 0.246 bps, respectively. The larger price impact of D2C trades likely reflects the institutional nature of the index CDS market in which customers are sophisticated investors who may be better than dealers at interpreting public information regarding aggregate credit risk in the economy.⁵ In contrast, D2D trades mainly serve to manage dealers' inventory risk (see, e.g., Reiss and Werner (1998)). 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 not informationally motivated but rather motivated by investors seeking to maintain a liquid credit exposure with a relatively constant maturity profile. 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 investigate how trade characteristics and market conditions affect transaction costs and price impacts. Transaction costs and price impacts increase with trade size, quoted bid-ask spread, and volatility implied by index swaptions; i.e., options on index CDSs. Our findings regarding differences in transaction costs and price impacts of D2C and D2D trades are robust to controlling for these determinants in trade-by-trade regressions. 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.

We also analyze the dynamics of D2C trades, D2D trades, and quotes using a vector autoregressive (VAR) model in the spirit of Hasbrouck (1991a, 1991b). Order flow is persistent and characterized by one-way Granger causality, with D2C trades Granger-causing D2D trades, which is consistent with inventory management taking place in the interdealer market. In

⁵In 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.

Chapter 2. Market Structure and Transaction Costs of Index CDSs

line with our findings based on the above-mentioned 15-minute price impact measure, D2C trades have larger contemporaneous and cumulative effects on quotes than D2D trades.

Finally, we investigate how the use of trading protocols that are only available in the interdealer market contribute to the differences in transaction costs and price impacts of D2C and D2D trades. 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; i.e., by means of quantity exchange at a fixed price (see, e.g., Duffie and Zhu (2015)).⁶ In contrast to marketable orders that execute against the best bid or offer on the order book, the execution of orders for matching and workup is uncertain because it depends on interests from the other side of the market.

Mid-market matching is the dominant trading protocol and accounts for 52.2% and 58.6%, respectively, of the trading volume in five-year on-the-run CDX.IG and CDX.HY. Workup is also frequently used and accounts for 19.1% and 14.9%, respectively. Mid-market matches have significantly lower transaction costs and price impacts than order-book trades. This is consistent with Zhu's (2014) venue-selection model, in which liquidity traders prefer a mid-point dark pool (essentially equivalent to continuous mid-market matching) that offers price improvement but does not guarantee execution, while execution risk causes informed traders to prefer an exchange that guarantees immediate execution at a market maker's bid or offer. By design, a workup is initiated by an order-book trade and occurs at the same transaction price. However, we find that price impacts of workups are close to those of order-book trades implying that this trading protocol allows to expand the size of an order-book trade with little additional price impact. These results suggest that size-discovery trading protocols attract liquidity-motivated trading and contribute to lowering overall transaction costs and price impacts of D2D trades.

We also use the GFI data to estimate dealer profits from liquidity provision in five-year on-the-run index CDSs. Assuming that dealers immediately close D2C trades by mid-market matches, estimated profits are USD 0.433 million and USD 0.808 million per day in case of CDX.IG and CDX.HY, respectively. However, assuming that dealers instead close positions at the best bid or offer on the order book, estimated profits are negative. Because mid-market matching is only possible when there is interest from the other side of the market, this suggests that dealers only make profits through their willingness to bear inventory risk.

From a regulatory perspective, our results show that the current two-tiered market structure delivers favorable prices for customers who value immediacy. The prices that customers obtain via RFQ are often better than those available on the order books of IDB SEFs. Indeed, 96.0% and 96.6% of the D2C trades in CDX.IG and CDX.HY, respectively, are executed at prices

⁶The two trading protocols differ in how the fixed price is determined and in the time span over which quantity can be exchanged. Mid-market matching is possible at a broker-determined price until the broker resets the price, while workup is possible at the price of an initiating order-book trade for a short period of time following trade execution.

that are strictly more favorable than the best bid or offer on the order book of the GFI Swaps Exchange. This suggests that regulators should not necessarily strive for all-to-all trading in swap markets.⁷ While customers who value immediacy would not be able to save transaction costs by executing their trades on the order books of IDB SEFs, transaction costs could be reduced at the expense of execution certainty either through liquidity supplying order-book trades or through mid-market matching.

The paper is related to a number of studies documenting the impact of the implementation of Dodd-Frank Act provisions on swap market liquidity. Loon and Zhong (2016) show that post-trade transparency has a positive impact on liquidity in the index CDS market. Benos, Payne, and Vasios (2016) show that pre-trade transparency (SEF mandate) 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 implementation of the Dodd-Frank Act and compare liquidity and transaction costs across the two segments of the market. Moreover, we contribute to the literature by showing how some unique features of swap trading such as the packaging of trades, mid-market matching, and workup affect the cost at which a swap can be traded.

Consistent with our results, Biswas, Nikolova, and Stahel (2015) find that, in the single-name CDS market, D2D trades have lower transaction costs than D2C trades. However, their transaction cost estimates are relatively imprecise due to a lack of transaction timestamps. Also, they do not investigate the price impact of trades, a cost component that we show is crucial for the comparison of D2C and D2D transaction costs.

In some respects, our results differ from those of studies that analyze transaction costs in the corporate and municipal bond markets, in which dealers seem to exert market power, and retail-sized trades have significantly higher transaction costs than institutional-sized trades (see, e.g., Harris and Piwowar (2006), Edwards, Harris, and Piwowar (2007), and Green, Hollifield, and Schürhoff (2007)).⁸ Consistent with the institutional nature of the index CDS market, we find D2C transaction costs that increase with trade size.

The paper is organized as follows: Section 2.2 describes the structure of the index CDS market and the regulatory reforms set forth by the Dodd-Frank Act. Section 2.3 discusses the data and the identification algorithms. Section 2.4 compares D2C and D2D transaction costs and investigates how transaction costs vary with trade characteristics and market conditions. Section 2.5 analyzes the dynamics of trades and quotes using VAR methods. Section 2.6 uses GFI data to investigate transaction costs across different interdealer trading protocols and to estimate dealer profits from liquidity provision. Section 2.7 concludes, and data-related details and robustness checks are contained in Appendix B.

⁷ The regulatory implications go beyond the index CDS market. For instance, the interest rate swap market—which has been subject to the same set of regulatory reforms—remains two-tiered as well. Moreover, the Dodd-Frank Act constitutes a template for over-the-counter derivatives market regulations that other jurisdictions are going to implement in the coming years.

⁸In similar vein, Schultz (2001) finds that corporate bond trades of less active institutional investors have higher transaction costs than those of the most active institutional investors.

2.2 The Index CDS Market

This section briefly 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.

2.2.1 Index Credit Default Swaps

An index CDS is a standardized credit derivative contract on a diversified index of creditors. 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 often use the “par spread” which is the fixed spread that makes the upfront amount equal to zero. We use these 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 creditors being revised according to credit rating and liquidity criteria.⁹ Creditors that fail to maintain a credit rating within a specified range, due to either upgrades or downgrades, and creditors whose single-name CDSs have deteriorated significantly in terms of their trading activity are replaced by the most actively traded creditors meeting the credit rating requirements. Liquidity is typically 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 creditors with investment-grade credit ratings, and the latter comprises 100 North American creditors with non-investment-grade credit ratings. These indices are the focus of the paper.

2.2.2 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 traded with their institutional customers. D2C trades were either negotiated over the phone or executed electronically on trading platforms such as MarketAxess or Tradeweb.¹⁰ Electronic trade execution was typically via name-disclosed RFQs that enable querying multiple dealers simultaneously for an executable one-sided market of a given notional amount.

⁹An index's series number uniquely determines the creditors in the index.

¹⁰Electronic 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.

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.

2.2.3 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, a reporting requirement, and a trade execution requirement.¹¹

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 trade in which case dissemination is delayed by at least 15 minutes.¹²

The trade execution requirement mandates that the most liquid index CDSs trade on SEFs and via one of two trading methods: the order book or an RFQ that is transmitted to at least three other market participants on the SEF.¹³ Since the trade execution requirement took effect, trades in five-year on-the-run and immediate off-the-run index CDSs on CDX.IG and CDX.HY have been subject to the requirement.¹⁴ 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 B.1 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.

¹¹See 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).

¹²Block trades have notional amounts that exceed certain minimum block sizes and are exempt from immediate dissemination to protect liquidity providers in large transactions from front running. Minimum block sizes depend on the index CDS 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).

¹³For an interim one-year period, it was sufficient to transmit RFQs to at least two other participants.

¹⁴In addition, trades in five-year on-the-run and immediate off-the-run index CDSs on iTraxx Europe and iTraxx Europe Crossover have been 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.

Chapter 2. Market Structure and Transaction Costs of Index CDSs

However, several years into the new regulatory regime, the index CDS market remains two-tiered and all-to-all trading has yet to materialize.¹⁵

Dealers trade with their institutional customers on SEFs run by incumbent operators of electronic 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. Dealers trade with each other on 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 (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 customers because of the risk of uncontrolled information leakage of proprietary trading strategies.¹⁶

2.3 Data and Identification Algorithms

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

2.3.1 Data

Our empirical analysis is based on trades and quotes over a two-year period from October 2, 2013 (when most SEFs started operating) 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

¹⁵Implicitly, 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)).

¹⁶Trading via RFQ also entails a certain amount of information leakage, but in this case the customer 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 (see, e.g., “How to Game a SEF: Banks Fear Arrival of Arbitrageurs,” *Risk Magazine*, March 19, 2014).

at least USD 100 million,¹⁷ 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 exception that exempts the trade from the clearing and trade execution requirements.¹⁸ 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).¹⁹ 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., e-mails that dealers send to their institutional customers 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.²⁰

Figure 2.1 shows trades and the mid-point of Markit intraday 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 quotes between 7:00 a.m. and 5:30 p.m., New York time, 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 turn out to be delta hedges of index swaption trades, see below. Data processing also shows that the trades are composed of 139 D2C trades executed on D2C SEFs and 26 D2D trades executed on IDB SEFs.

2.3.2 Identification of SEFs

In devising the SEF identification algorithm, we use SEF-reported trading volumes from Clarus FT.²¹ 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.

¹⁷The 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).

¹⁸This 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).

¹⁹There 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.

²⁰Quotes from runs older than 15 minutes are discarded from the computation and a five-minute memory prevents repeated computations of the same composite.

²¹Clarus FT is the standard data source for SEF-reported daily trading volumes. In Appendix B.3, we describe the Clarus FT data in detail.

Chapter 2. Market Structure and Transaction Costs of Index CDSs




Figure 2.1: CDX.IG Trades and Mid-Quotes on May 6, 2015.

The figure shows transaction prices of all dealer-to-customer (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 index CDS spreads and expressed in basis points (bps). Series 24 of CDX.IG was on-the-run on May 6, 2015.

Specifically, we associate with each SEF the format of trade reports whose aggregate trade size corresponds to the SEF-reported trading volume (Appendix B.2 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 focusing on trades executed on SEFs is not restrictive because the most actively traded index CDSs are subject to the trade execution requirement. The majority of trade reports that we do not capture come from the period before the requirement took effect. These are most likely D2C trades because any D2D trade facilitated by an IDB would have been on-SEF.

2.3.3 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 (Appendix B.2 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 transactions 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.²² 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 transactions 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 level 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 as the index CDS spread moves. For CDX.IG, the reference level is usually set in spread multiples of 0.5 bps.²³ We identify index swaption and tranche swap delta hedges as index CDS transactions that have the same underlying index and contract tenor as an index swaption or tranche swap transaction. 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 transaction 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 current level 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 2.1.

2.3.4 SEF Order Flow

Table 2.1 displays descriptive statistics of the enriched transaction data that allows to distin-

²²Typically, the underlying of both index CDSs is the same but there are also curve trades in which the two index CDSs reference different index series.

²³Because CDX.HY is quoted in terms of a price, the reference level is usually set in price multiples of 0.125%.

SEF	USD MM					% of Trds					% of Vlm				
	Trds	Sz	Vlm (ActVlm)	5Y	OTR	Bspk	Clrd	Blck	Cppd	Crv	Rll	Swptn	Trmch		
Panel A1: CDX.IG Dealer-To-Customer															
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-Customer															
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 2.1: Descriptive Statistics of On-SEF Index CDS Trades.

The table shows descriptive statistics of on-SEF dealer-to-customer (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. Sz is median trade size. Vlm (ActVlm) is (actual SEF-reported) daily volume. 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 (Blck) is the percentage of cleared (block) trades. Cppd is the percentage of trades with capped notional amounts. Crv (Rll) is the percentage of daily volume due to curve trades (index rolls). Swptn (Trmch) is the percentage of daily volume due to index swaption (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.

guish between D2C and D2D trades and between outright 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 executed on a given SEF.

In terms of the notional amount traded, D2C trades in CDX.IG account for a daily trading volume of USD 9.843 billion, on average, and those in CDX.HY account for a daily trading volume of USD 3.705 billion, on average. In comparison, D2D trades in the two indices account for average daily trading volumes of USD 1.354 billion and USD 0.402 billion, respectively.²⁴ These averages appear in parenthesis in Table 2.1 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. Table 2.1 shows that 21.2% and 2.3% of D2C trades in CDX.IG and CDX.HY, respectively, are disseminated with capped trade sizes, while the corresponding numbers for D2D trades are 6.8% and 1.4%, respectively.²⁵ As a consequence, transaction-data-based average daily trading volumes are downward biased.²⁶

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.²⁷ Outright trades account for most of the trading volume and, among package transactions, index rolls are most popular, accounting for 5.0% and 8.9% (17.3% and 21.7%) of D2C (D2D) trading volume in CDX.IG and CDX.HY, respectively. The fact that there are

²⁴D2D 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.

²⁵In 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).

²⁶The 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 trade size of D2C trades in CDX.IG is approximately USD 250 million. Similarly, conditional on being capped, the average trade 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 trade sizes of D2C and D2D trades in CDX.HY are approximately USD 225 million and USD 160 million, respectively.

²⁷Loon 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 eventual transaction cost differences between D2C and D2D trades because the vast majority of both D2C and D2D trades are non-bespoke and centrally cleared. We observe an increasing share of end-user exempt transactions prior to February 10, 2014 (around 80% of trades on February 7, 2014 are end-user exempt) and not a single end-user exempt trade afterwards. We are not aware of no-action reliefs issued by the CFTC that expired on February 10, 2014 and could explain the decline.

Chapter 2. Market Structure and Transaction Costs of Index CDSs

Type	CDX.IG		CDX.HY	
	D2C	D2D	D2C	D2D
Outright				
Five-year on-the-run	90.0	79.9	91.2	85.5
Five-year immediate off-the-run	5.3	0.8	5.1	1.4
Other	0.7	1.2	0.2	2.1
Package				
Roll five-year on-the-run immediate off-the-run	3.4	7.6	3.2	6.4
Other roll	0.1	2.9	0.1	1.0
Curve	0.1	3.0	0.0	0.0
Swaption delta hedge	0.3	2.4	0.2	2.2
Tranche swap delta hedge	0.1	2.1	0.1	1.4

Table 2.2: Percentages of On-SEF Index CDS Trades by Trade Type.

The table shows percentages of on-SEF dealer-to-customer (D2C) and dealer-to-dealer (D2D) index CDS trades in CDX.IG and CDX.HY 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.

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 executed on order books do not qualify as block trades.

As explained in Section 2.2.1, liquidity in the index CDS market concentrates in on-the-run index CDSs and, in particular, those with a five-year contract tenor. Therefore, we separately break down total D2C and D2D trades into fractions of transactions in specific contracts and report results in Table 2.2. Specifically, the table shows fractions of outright trades in five-year on-the-run and immediate off-the-run index CDSs, and fractions of transactions that are part of index rolls between these contracts. For completeness, we also report fractions of outright trades in other index CDS contracts, fractions of transactions that are part of other index rolls and curve trades, and fractions of index swaption and tranche swap delta hedges.

In case of both indices, trades in five-year on-the-run and immediate off-the-run index CDSs make up almost the entire D2C on-SEF trading activity. More than 90% of D2C trades are outright trades in five-year on-the-run index CDSs, around 5% of D2C transactions are outright trades in five-year immediate off-the-run index CDSs, and index rolls between these contracts account for another 3%. The residual transactions account for less than 2% of D2C trades. D2D on-SEF trading activity is a little more diverse but outright trades in five-year on-the-run and immediate off-the-run index CDSs and index rolls between these contracts nevertheless account for 88.3% and 93.3% of D2D trades in CDX.IG and CDX.HY, respectively.

From Table 2.1 it follows that outright trades account for 94.3% of D2C trading volume in CDX.IG. Breaking down the volume share along the lines of Table 2.2 shows that outright trades in five-year on-the-run and immediate off-the-run index CDSs account for 88.5% and 5.3% of D2C trading volume, respectively, and D2C index rolls between these contracts account

for another 4.9%. Thus, in aggregate these trades account for a comparable high amount of both trades and volume. For D2D trades in CDX.IG, the aggregate share of outright trades in five-year on-the-run and immediate off-the-run index CDSs and index rolls between these contracts is slightly lower in terms of volume than in terms of trades. Nevertheless, trades in these contracts account for 80.3% of D2D trading volume in CDX.IG. Similar results obtain for trades in CDX.HY.

2.4 Transaction Cost Comparison

In order to analyze what determines transaction costs of D2C and D2D trades in the two-tiered index CDS market, we focus on outright trades in on-the-run index CDSs and index rolls between on-the-run and immediate off-the-run index CDSs, all with a five-year contract tenor.²⁸ As highlighted by the preceding discussion, together these trades account for the majority of both transactions and trading volumes in the index CDS market.

2.4.1 Transaction Cost Decomposition

We measure transaction costs by effective half-spreads with respect to Markit's intraday mid-quote. Recognizing that spreads reflect both dealer revenue and the information content of trade, we further decompose effective half-spreads into realized half-spreads and price impacts. Specifically,

$$\underbrace{q_t(p_t - m_t)}_{=\text{EfficSprd}_t} = \underbrace{q_t(p_t - m_{t+\Delta})}_{=\text{RlzdSprd}_t} + \underbrace{q_t(m_{t+\Delta} - m_t)}_{=\text{PrcImp}_t}, \quad (2.1)$$

where p_t is the transaction price of the t -th trade in index CDS i (we suppress dependence on i because all our analyses separately focus on one type of trade in CDSs on a given index), m_t is the mid-point of the latest composite quote in the 15-minute period prior to trade execution, and $m_{t+\Delta}$ is the mid-point of the first 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,²⁹ and m_t ($m_{t+\Delta}$) is the corresponding difference in quote mid-points.³⁰ Trade direction, q_t , is inferred by the Lee and Ready (1991) algorithm and equals +1 (−1) in case of protection-buyer-initiated (protection-

²⁸ In Appendix B.6, we provide an analysis of outright trades in immediate off-the-run index CDSs. Results are consistent with those of outright trades in on-the-run index CDSs. For the other trade types there are too few transactions to reliably measure transaction costs.

²⁹ Following market convention, p_t is the on-the-run minus the immediate off-the-run index CDS spread.

³⁰ Specifically, m_t is the corresponding difference in mid-points of the latest 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-points of 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.

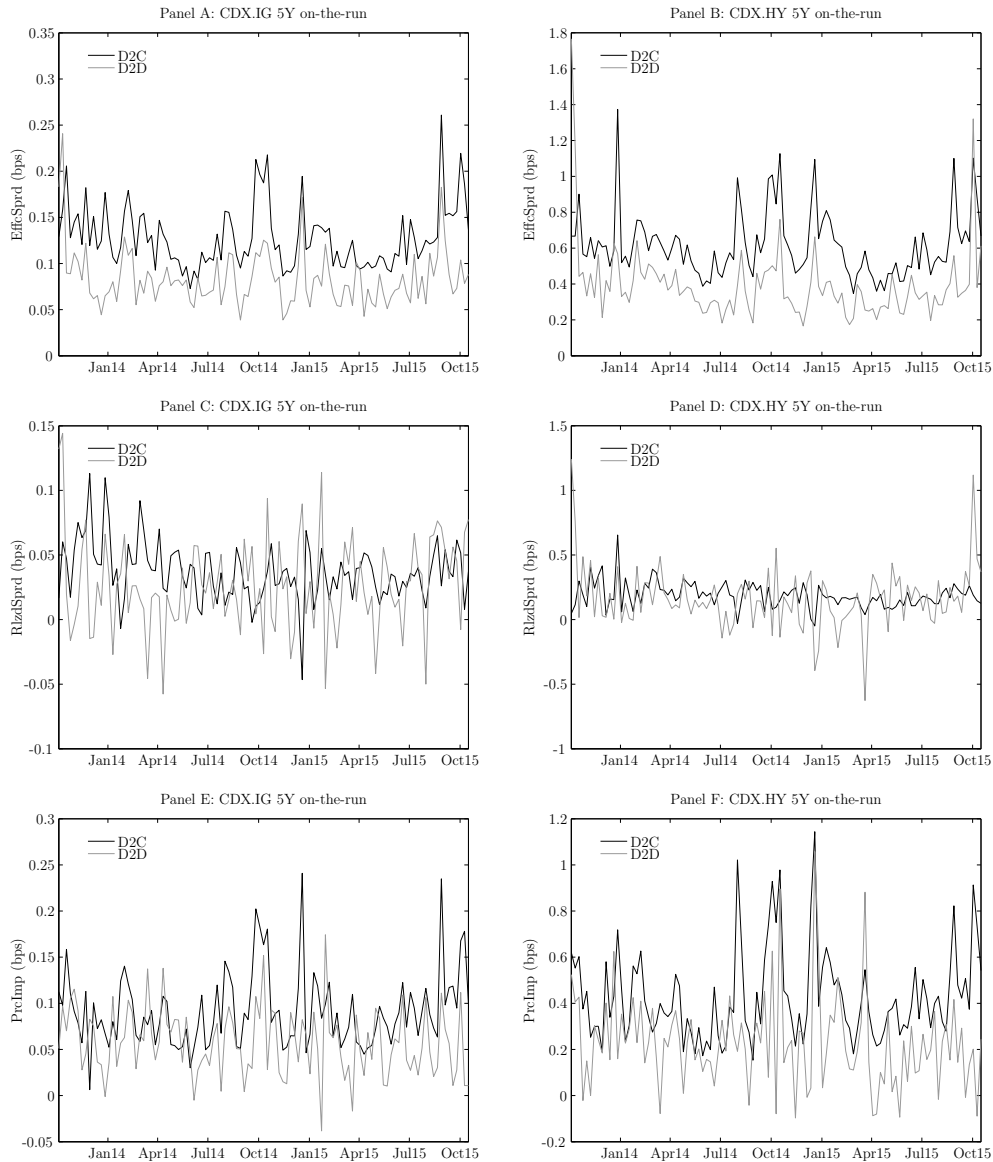


Figure 2.2: Weekly Average Effective Half-Spreads, Realized Half-Spreads, and Price Impacts. Panels A and B, Panels C and D, and Panels E and F, show weekly sample means of effective half-spreads (EfficSprd), realized half-spreads (RlzdSprd), and price impacts (Prclmp) of outright dealer-to-customer (D2C) and dealer-to-dealer (D2D) trades in five-year on-the-run index CDSs on CDX.IG and CDX.HY, respectively. EfficSprd 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. Prclmp is defined as $q_t \times (m_{t+\Delta} - m_t)$. Both transaction prices and quotes are in terms of index CDS 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.

seller-initiated) trades.

Assuming that one counterparty of each trade is a liquidity providing dealer, Equation (2.1) can be interpreted as follows: the effective half-spread measures the liquidity providing dealer's revenue if she were able to immediately close her position at the prevailing mid-quote. If instead it takes the dealer Δ units of time to close her position (and again assuming that she is able to do so at the then prevailing mid-quote), her revenue is the realized half-spread. The revenue is less than the effective half-spread if the price moves against the dealer while she is reversing the trade over time. Price impact captures such trade-induced price moves or adverse selection costs.

2.4.2 Descriptive Statistics

Figure 2.2 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 are consistent with this in that there is no systematic difference between the realized half-spreads of D2C and D2D trades.

Table 2.3 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 2.2. In case of CDX.IG, average effective half-spreads are 0.137 bps and 0.088 bps for D2C and D2D trades, 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 mostly due to D2C trades having larger price impacts than D2D trades. For CDX.IG, average price impacts are 0.106 bps and 0.063 bps for D2C and D2D trades, 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 average per trade revenues (as captured by realized half-spreads) across D2C and D2D trades.

As explained in Section 2.3.3, index rolls are liquidity motivated. Consistent with a non-informational motive for trade, Table 2.3 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 effective half-spreads and price impacts across D2C and D2D trades.

Table 2.4 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

Chapter 2. Market Structure and Transaction Costs of Index CDSs

Type	Dealer-To-Customer			Dealer-To-Dealer			D2C-D2D		
	Effc Sprd	Rlzd Sprd	Prc Imp	Effc Sprd	Rlzd Sprd	Prc Imp	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 2.3: Effective Half-Spreads, Realized Half-Spreads, and Price Impacts by Trade Type.

Panels A and B show sample means of effective half-spreads (EfficSprd), realized half-spreads (RlzdSprd), and price impacts (PrcImp) of dealer-to-customer (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. EfficSprd 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 index CDS 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.

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, Hotchkiss, and Sirri (2007), Harris and Piwowar (2006), and Green et al. (2007). This reflects structural differences between these markets: the index CDS market is purely institutional with sophisticated market participants trading in large sizes; in contrast, bond markets have retail segments with unsophisticated market participants trading in small sizes and with dealers who seem to exert market power.

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 large transactions.

2.4. Transaction Cost Comparison

Trade Size	Dealer-To-Customer			Dealer-To-Dealer			D2C-D2D		
	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 2.4: 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-customer (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 index CDS 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.

2.4.3 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 during which these 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 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

$$y_t = \alpha_{D2C} D2C_t + \alpha_{D2D} D2D_t + \beta' X_t + \epsilon_t, \quad (2.2)$$

Chapter 2. Market Structure and Transaction Costs of Index CDSs

where y_t is either the effective half-spread, the realized half-spread, or the price impact of the t -th trade in index CDS i (as before, we suppress dependence on i), $D2C_t$ and $D2D_t$ are dummy variables for D2C and D2D trades, respectively, and X_t is a vector of control variables.³¹ 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 for the five-year on-the-run index CDS prior to trade execution, the corresponding mid-quote, and end-of-day three-month at-the-money implied index swaption volatility (the swaption's underlying is the five-year on-the-run 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 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 2.4 shows that D2C transaction costs and price impacts tend to increase with trade size, and we include a reference level dummy to account for potentially unidentified index swaption and tranche swap delta hedges.³²

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 2.3 because the latter correspond to coefficient estimates of a restricted version of Equation (2.2) which excludes control variables.

Table 2.5 displays regression results. Accounting for trade characteristics and market conditions does not materially change the conclusions from Table 2.3. 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 2.3) and statistically significant. For CDX.HY, the difference is 0.219 bps (in comparison to 0.273 bps in Table 2.3) and statistically significant. For both indices, the estimated regression coefficients show that transaction costs increase with trade size, when bid-ask spreads widen (i.e., when liquidity deteriorates), and when implied volatility increases. In addition, trades with reference level transaction prices are more expensive.

³¹Comparing regression-based methods of 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" (Bessembinder 2003, p. 8).

³²One 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 unless they are traded with a delta hedge in a made available to trade index CDS. This requirement was temporarily overruled by a no-action relief (see CFTC Letter No. 14-12 (Feb. 10, 2014) and its extensions CFTC Letter No. 14-62 (May. 1, 2014) and CFTC Letter No. 14-137 (Nov. 10, 2014)). Typically, the delta hedge of an off-SEF index swaption or tranche swap trade would nevertheless be executed on-SEF in order to satisfy other regulatory requirements.

2.4. Transaction Cost Comparison

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

Table 2.5: 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-customer (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 index CDS 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; index CDS 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.

Chapter 2. Market Structure and Transaction Costs of Index CDSs

Again, transaction cost differentials are mostly due to differences in price impact. For CDX.IG, the difference in price impacts of D2C and D2D trades is a statistically significant 0.027 bps that accounts 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 that accounts 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 when bid-ask spreads widen and when implied volatility increases. 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.

Finally, consistent with a non-informational motive for trade, regression results for index rolls (displayed in Table 2.6) do not reveal significant differences in effective half-spreads and price impacts of D2C and D2D index rolls.³³ The estimated regression coefficients show that index roll transaction costs are insensitive to the size rolled and increase when bid-ask spreads widen.

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 in the same index CDS and 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 2.7 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 the same 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 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.

³³Because trade sizes of index rolls are relatively large (e.g., more than 50% of CDX.IG index rolls have capped on-the-run leg trade sizes), the regressions for index rolls only include a dummy variable for block-sized index rolls which is based on the trade size of the on-the-run leg.

³⁴Similar matching methods have, e.g., 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.

2.4. Transaction Cost Comparison

	CDX.IG			CDX.HY		
	EffcSprd	RlzdSprd	Prclmp	EffcSprd	RlzdSprd	Prclmp
D2C	0.047** (13.11)	0.024** (5.97)	0.023** (6.31)	0.373** (16.52)	0.224** (8.12)	0.149** (5.52)
D2D	0.049** (9.70)	0.028** (3.82)	0.021** (3.17)	0.360** (10.98)	0.134** (3.09)	0.226** (5.77)
BLCK	0.002 (0.47)	-0.007 (-1.27)	0.009 (1.89)	0.022 (0.84)	0.017 (0.52)	0.004 (0.15)
BAS	0.356** (3.96)	0.175* (2.18)	0.181* (2.30)	0.224** (4.94)	0.223* (2.37)	0.001 (0.01)
SPRD/100	-0.062 (-1.69)	-0.019 (-0.58)	-0.043 (-1.17)	-0.053 (-1.33)	0.002 (0.04)	-0.055 (-0.95)
VLTLTY	0.041 (0.68)	-0.061 (-1.33)	0.102 (1.82)	0.632** (2.58)	-0.127 (-0.29)	0.759 (1.47)
<i>N</i>	1,281	1,281	1,281	1,423	1,423	1,423
D2C – D2D	-0.002	-0.004	0.002	0.013	0.090	-0.077
<i>p</i> -value	0.72	0.61	0.79	0.65	0.02	0.06

Table 2.6: Regressions Controlling for Index Roll 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 (Prclmp) of dealer-to-customer (D2C) and dealer-to-dealer (D2D) index rolls between five-year on-the-run and immediate off-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 difference between on-the-run and immediate off-the-run transaction prices and m_t is the difference between the latest on-the-run and immediate off-the-run mid-quotes 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. Prclmp is defined as $q_t \times (m_{t+\Delta} - m_t)$. Both transaction prices and quotes are in terms of index CDS 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), and for block-sized trades (BLCK; +USD 100 MM for CDX.IG and +USD 25 MM 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 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.

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.

In Appendix B.7, we repeat this section's analyses using an alternative intraday mid-quote

Chapter 2. Market Structure and Transaction Costs of Index CDSs

Matching Trade Size	Dealer-To-Customer			Dealer-To-Dealer			D2C-D2D		
	Effc Sprd	Rlzd Sprd	Prc Imp	Effc Sprd	Rlzd Sprd	Prc Imp	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 2.7: Effective Half-Spreads, Realized Half-Spreads, and Price Impacts of Matched Pairs. Panels A and B show sample means of effective half-spreads (EfficSprd), realized half-spreads (RlzdSprd), and price impacts (Prclmp) of matched pairs of outright dealer-to-customer (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. EfficSprd 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. Prclmp is defined as $q_t \times (m_{t+\Delta} - m_t)$. Both transaction prices and quotes are in terms of index CDS 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 identical trade size) that occur within a 15-minute window bracketing the D2D trade. In case of more than one matching D2C trade, the EfficSprd, RlzdSprd, and Prclmp 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.

from Credit Market Analysis. The results we obtain are consistent with the ones reported here. We also show that the 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 trade characteristics.

2.5 The Dynamics of Trades and Quotes

The evidence thus far has revealed that D2C trades have both higher transaction costs and larger price impacts than D2D trades. In order to investigate the price discovery process that gives rise to the differential price impact, we analyze a VAR model that accounts for the

two-tiered structure of the index CDS market and the market-specific quote provision in form of dealer runs. In comparison to the relatively simple ad hoc decomposition of the effective half-spread that we used in the previous section, the VAR model accounts for persistence in order flow and dynamic interactions between quote revisions and trades.

2.5.1 VAR Framework and Model 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,

$$\Delta m_t = \sum_{j=1}^{10} \alpha_j \Delta m_{t-j} + \sum_{j=0}^{10} \beta_j x_{t-j}^{\text{D2C}} + \sum_{j=0}^{10} \gamma_j x_{t-j}^{\text{D2D}} + \epsilon_t^{\Delta m}, \quad (2.3a)$$

$$x_t^{\text{D2C}} = \sum_{j=1}^{10} \delta_j \Delta m_{t-j} + \sum_{j=1}^{10} \zeta_j x_{t-j}^{\text{D2C}} + \sum_{j=1}^{10} \eta_j x_{t-j}^{\text{D2D}} + \epsilon_t^{x, \text{D2C}}, \quad (2.3b)$$

$$x_t^{\text{D2D}} = \sum_{j=1}^{10} \kappa_j \Delta m_{t-j} + \sum_{j=0}^{10} \lambda_j x_{t-j}^{\text{D2C}} + \sum_{j=1}^{10} \rho_j x_{t-j}^{\text{D2D}} + \epsilon_t^{x, \text{D2D}}, \quad (2.3c)$$

where t indexes the t -th quote revision (i.e., computation of a composite quote) 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, $\epsilon_t^{\Delta m}$, $\epsilon_t^{x, \text{D2C}}$, and $\epsilon_t^{x, \text{D2D}}$, are uncorrelated because we resolve contemporaneous effects by including contemporaneous trade-related variables in Equations (2.3a) and (2.3c).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 adjust quotes in response to trades.

Hasbrouck (1991a, 1991b) argues that VAR systems like the one in Equations (2.3a) to (2.3c) provide a flexible and robust framework in which permanent (information-driven) and transitory (microstructure-driven) quote changes can be separated. Specifically, because microstructure effects fade away in the long-run, the system-implied long-run cumulative mid-quote change in response to a shock of trade-related variables measures the information content of trade (or the price impact). The latter can be conveniently estimated from the vector moving average (VMA) representation of the VAR model; that is,

$$\Delta m_t = a_0 \epsilon_t^{\Delta m} + a_1 \epsilon_{t-1}^{\Delta m} + \dots + b_0 \epsilon_t^{x, \text{D2C}} + b_1 \epsilon_{t-1}^{x, \text{D2C}} + \dots + c_0 \epsilon_t^{x, \text{D2D}} + c_1 \epsilon_{t-1}^{x, \text{D2D}} + \dots, \quad (2.4a)$$

$$x_t^{\text{D2C}} = d_0 \epsilon_t^{\Delta m} + d_1 \epsilon_{t-1}^{\Delta m} + \dots + z_0 \epsilon_t^{x, \text{D2C}} + z_1 \epsilon_{t-1}^{x, \text{D2C}} + \dots + h_0 \epsilon_t^{x, \text{D2D}} + h_1 \epsilon_{t-1}^{x, \text{D2D}} + \dots, \quad (2.4b)$$

$$x_t^{\text{D2D}} = k_0 \epsilon_t^{\Delta m} + k_1 \epsilon_{t-1}^{\Delta m} + \dots + l_0 \epsilon_t^{x, \text{D2C}} + l_1 \epsilon_{t-1}^{x, \text{D2C}} + \dots + r_0 \epsilon_t^{x, \text{D2D}} + r_1 \epsilon_{t-1}^{x, \text{D2D}} + \dots. \quad (2.4c)$$

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

It immediately follows from Equation (2.4a), that the price impact of a single protection-buyer-initiated D2C trade is

$$\Lambda_{x,D2C} = \lim_{n \rightarrow \infty} \sum_{j=0}^n \mathbb{E} \left[\Delta m_{t+j} \middle| \Omega_t^{x,D2C} \right] = \lim_{n \rightarrow \infty} \sum_{j=0}^n b_j = \sum_{j=0}^{\infty} b_j, \quad (2.5)$$

where $\Omega_t^{x,D2C} = \{e_t^{x,D2C} = 1, e_t^{\Delta m} = e_t^{x,D2D} = 0, e_s^{\Delta m} = e_s^{x,D2C} = e_s^{x,D2D} = 0, s < t\}$ denotes the event of an isolated unit-sized shock of the D2C-trade-related variable.³⁶ Similarly, the price impact of a single protection-buyer-initiated D2D trade is $\Lambda_{x,D2D} = \sum_{j=0}^{\infty} c_j$.

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, e.g., 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, e.g., 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, $\sigma_{\Delta \bar{p}}^2$, can be explicitly expressed in terms of error term variances and VMA-representation parameters; that is,

$$\sigma_{\Delta \bar{p}}^2 = \left(\sum_{j=0}^{\infty} a_j \right)^2 \sigma_{\Delta m}^2 + \left(\sum_{j=0}^{\infty} b_j \right)^2 \sigma_{x,D2C}^2 + \left(\sum_{j=0}^{\infty} c_j \right)^2 \sigma_{x,D2D}^2, \quad (2.6)$$

where $\sigma_{\Delta m}^2 = \mathbb{V}(e_t^{\Delta m})$, $\sigma_{x,D2C}^2 = \mathbb{V}(e_t^{x,D2C})$, and $\sigma_{x,D2D}^2 = \mathbb{V}(e_t^{x,D2D})$. Equation (2.6) 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 (2.6), 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 (2.6) is the basis of our price discovery metric, Hasbrouck's (1991b) R^2 , that expresses each component's variance as a fraction of $\sigma_{\Delta \bar{p}}^2$.

We estimate the VAR model using all quote changes between 7:00 a.m. and 5:30 p.m., New York time, during our sample period.³⁷ We splice together intraday quote changes of five-year on-the-run index CDSs and the corresponding numbers of signed outright trades in order to create a continuous vector-valued time series from which we can estimate the VAR model. We exclude quote changes that span long periods of time presumably because of technical issues with the composite computation.³⁸ Finally, we winsorize quote changes at the 0.1% and 99.9%

³⁶A single protection-buyer-initiated D2C trade is an event in $\Omega_t^{x,D2C}$ 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.

³⁷When estimating the VAR model, we assume that the system is in steady-state at the beginning of each trading day.

³⁸The fact that, over these time spans, there are typically neither quotes for CDX.IG nor for CDX.HY suggests technical disruptions.

quantile of their distribution.

2.5.2 Results

Panels A1 and A2 of Table 2.8 display VAR coefficient estimates for CDX.IG and CDX.HY, respectively. The results for both indices are similar and, therefore, the discussion focuses on CDX.IG. The significant coefficients of contemporaneous trade-related variables in Equation (2.3a) 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 generally positive and significant coefficients of lagged D2C-(D2D-)trade-related variables in Equation (2.3b) (Equation (2.3c)) indicate positively autocorrelated D2C (D2D) trades. This reflects persistence in order flow, a pervasive feature of trade in financial markets.³⁹ Consistent with dealers hedging customer trades in the inter-dealer market, coefficients of contemporaneous and lagged D2C-trade-related variables in Equation (2.3c) are generally positive and significant.

The generally positive coefficients of lagged mid-quote changes in Equation (2.3b) suggest that quote changes are positively related to D2C trades. This is in contrast to the negative relation implied by inventory control considerations of an individual dealer who sets quotes so as to elicit customer trades in the direction of inventory (i.e., who reduces quotes to elicit protection-buyer-initiated customer trades when being a net protection buyer and, vice versa, when being a net protection seller). Instead, the positive relation may reflect momentum-driven trading by customers.

Granger causality tests reveal that the dynamic interaction between D2C- and D2D-trade-related variables is characterized by one-way Granger causality with D2C trades Granger-causing D2D trades. This is consistent with inventory management taking place in the interdealer market. Many market participants, in fact, view D2D trades as primarily hedging motivated. In support of this view, price discovery fractions of D2D trades in Panel C of Table 2.8 are virtually zero.

Figure 2.3 shows trade-induced cumulative quote revisions implied by the estimated VAR models. Specifically, the figure tracks the cumulative quote revision following single protection-buyer-initiated D2C and D2D trades, respectively. Consistent with the evidence based on the simple price impact measure of Section 2.4, we find that a single protection-buyer-initiated D2C trade has a larger cumulative effect on quotes than a single protection-buyer-initiated D2D trade. A formal statistical test regarding the ultimate price impact of a trade—i.e., the long-run limit of the cumulative quote revisions exhibited in Figure 2.3—is provided in Panel B of Table 2.8 and rejects the hypothesis of identical price impacts of D2C and D2D trades.

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

Chapter 2. Market Structure and Transaction Costs of Index CDSs

	Coefficient Estimates					Granger Causality Tests		
	$\sum_{j=1}^{10} \Delta m_{t-j}$	x_t^{D2C}	$\sum_{j=1}^{10} x_{t-j}^{\text{D2C}}$	x_t^{D2D}	$\sum_{j=1}^{10} x_{t-j}^{\text{D2D}}$	Δm	x^{D2C}	x^{D2D}
Panel A1: CDX.IG								
Δm_t	0.344 (67.08)	0.009 (77.19)	0.017 (48.00)	0.003 (14.79)	0.005 (8.67)		4793.6 [<0.01]	141.7 [<0.01]
x_t^{D2C}	2.069 (22.07)		0.239 (37.40)		0.022 (1.91)	1096.3 [<0.01]		15.0 [0.13]
x_t^{D2D}	-0.172 (-3.37)	0.023 (19.87)	0.030 (8.67)		0.137 (21.96)	27.2 [<0.01]	102.4 [<0.01]	
Panel A2: CDX.HY								
Δm_t	0.254 (42.61)	0.042 (81.68)	0.074 (49.49)	0.011 (9.48)	0.020 (6.07)		5351.7 [<0.01]	84.5 [<0.01]
x_t^{D2C}	0.410 (15.28)		0.322 (48.10)		-0.007 (-0.49)	851.6 [<0.01]		11.7 [0.31]
x_t^{D2D}	-0.071 (-5.70)	0.023 (21.21)	0.027 (8.59)		0.131 (19.49)	49.3 [<0.01]	119.7 [<0.01]	
Panel B: Price Impact								
	CDX.IG			CDX.HY				
	D2C	D2D	D2C – D2D	D2C	D2D	D2C – D2D		
Λ	0.060 (65.51)	0.019 (12.24)	0.041 (23.53)	0.256 (68.47)	0.049 (6.87)	0.207 (25.80)		
Panel C: Price Discovery								
	CDX.IG			CDX.HY				
	D2C	D2D	Trade-Unrelated	D2C	D2D	Trade-Unrelated		
R^2	28.61	0.83	70.56	37.73	0.30	61.97		

Table 2.8: 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-customer (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 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 index CDS 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.

We relegate estimation of VAR model specifications that take trade size into account to Appendix B.9. The results we obtain are consistent with the ones we report here. When estimating

2.6. Why is Trade in the Interdealer Market Cheaper?

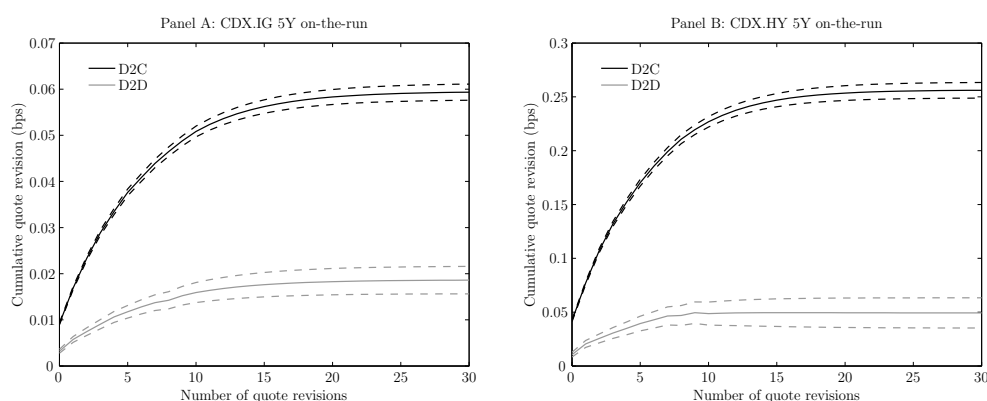


Figure 2.3: VAR-Model-Implied Price Impact.

The panels show cumulative quote revisions in response to either a single protection-buyer-initiated dealer-to-customer (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 index CDS 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.

alternative VAR model specifications, we observe that the use of trade size does not add to the explanatory power of VAR models. This is reminiscent of Jones, Kaul, and Lipson (1994) who find that in the equity market trade size has little incremental explanatory power above that contained in the number of transactions.

2.6 Why is Trade in the Interdealer Market Cheaper?

In intermediating D2D trades, IDBs have developed a variety of trading protocols that are not available in the D2C segment. These include mid-market matching and workup. The distinctive feature of these two trading protocols is that trade occurs through size discovery; i.e., through quantity exchange at a fixed price (see, e.g., Duffie and Zhu (2015)). Because the price at which the exchange takes place is fixed, the price is insensitive to price pressure and allows for exchange of potentially large quantities with little price impact.

We use unique order-book data from the main IDB SEF, the GFI Swaps Exchange, to investigate how the use of these trading protocols contributes to the low transaction costs and small price impacts of D2D trades.⁴⁰ We also use this data to estimate dealer profits from liquidity provision.

⁴⁰Focusing on trades executed on the GFI Swaps Exchange is not restrictive because it is the IDB SEF facilitating the majority of D2D trading volume (see Table 2.1). Other IDB SEFs also offer matching and workup, but order-book data are unavailable.

2.6.1 Size-Discovery Trading Mechanisms

In addition to conventional order book and RFQ trading protocols (both of which are price-discovery trading mechanisms), the GFI Swaps Exchange offers both matching and workup trading protocols. For five-year on-the-run index CDSs, the primarily used size-discovery trading mechanism is continuous mid-market matching. The mid-market level is set by a GFI broker and is usually somewhere between the best bid and offer currently resting on the 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 fed onto the trading screen that displays the 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 potential interests but they know that interests must be at least of a minimum size.⁴¹ Any opposing interests for matching at the mid-market level immediately result in a trade.

Workup sessions on the GFI Swaps Exchange are initiated by trades on the order book. During these sessions, 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 a trade 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 that can work up trade size. Subsequent to the exclusivity period, other market participants can join the trade for another 30 seconds before the workup session terminates. In contrast to continuous mid-market matching, market participants are informed about which side of the market has unfilled interests.

2.6.2 Data and Identification of Mid-Market Matches and Workups

The GFI data consist of the best bid and offer quotes that rest on the order book of the GFI Swaps Exchange. In addition, the data include the mid-market levels that GFI brokers set for mid-market matching. From this data we identify order-book trades, mid-market matches, and workups (Appendix B.4 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 likely voice-brokered RFQs.

Table 2.9 shows volume shares of the different trade categories. We separately report volume shares for outright trades in five-year on-the-run and other index CDSs, for index rolls between five-year on-the-run and immediate off-the-run index CDSs, and for 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 index CDSs, mid-market matches account for 52.2% and 58.6% of trading volume in case of CDX.IG and CDX.HY, respectively. 20.0% and 16.5% of trading volume in these contracts gets executed at the best bid or offer and about the same share of volume is traded in ensuing workup

⁴¹Current minimum sizes are USD 25 million for CDX.IG and USD 10 million for CDX.HY.

2.6. Why is Trade in the Interdealer Market Cheaper?

	Mid-Market Matching	Order Book	Trade Workup	Other Method
Panel A: CDX.IG				
Outright				
Five-year on-the-run	52.2	20.0	19.1	8.8
Other	13.6	8.8	18.4	59.1
Package				
Roll five-year on-the-run immediate off-the-run	20.1	18.8	26.4	34.8
Other	7.9	25.3	25.6	41.1
Total	42.6	20.4	20.6	16.4
Panel B: CDX.HY				
Outright				
Five-year on-the-run	58.6	16.5	14.9	10.0
Other	6.5	8.2	14.0	71.3
Package				
Roll five-year on-the-run immediate off-the-run	31.2	20.3	14.9	33.6
Other	5.1	25.3	27.2	42.4
Total	49.3	17.3	15.4	18.0

Table 2.9: 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 of volume shares. 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.

sessions.⁴² Together size-discovery trading mechanisms account for the majority of trading volume in five-year on-the-run index CDSs, with aggregate volume shares of 71.3% and 73.5% in case of CDX.IG and CDX.HY, respectively.

2.6.3 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 (2.2). 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

$$y_t = \alpha + \beta_{\text{MTCH}}\text{MTCH}_t + \beta_{\text{WRKUP}}\text{WRKUP}_t + \beta_{\text{OTHER}}\text{OTHER}_t + \gamma' X_t + \epsilon_t, \quad (2.7)$$

where y_t and X_t are defined as before and MTCH_t , WRKUP_t , and OTHER_t are dummy variables for mid-market matches, workups, and trades with unidentified trading protocol. Thus,

⁴²About half of the five-year on-the-run trades that occur on the order book are subsequently worked up.

Chapter 2. Market Structure and Transaction Costs of Index CDSs

α estimates the average effective half-spread (or, depending on the dependent variable used, realized half-spread or price impact) of an outright order-book trade with trade size in the first quartile of the trade size distribution that is executed when average market conditions prevail, and β s estimate effective half-spread differences with respect to order-book trades.

Table 2.10 displays regression results. First, compare order-book trades and 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 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 informed trades occurring on the order book. 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 an exchange or a mid-point dark pool (essentially equivalent to continuous mid-market matching). Sending an order to a mid-point dark pool involves a trade-off between potential price improvement and the risk of no execution. Because informed traders face higher execution risk than liquidity traders, they concentrate on the exchange that guarantees immediate execution at a market maker's bid or offer.

Next, compare order-book trades and workups. There are no significant differences in effective half-spreads. This is by design as workups are executed at the transaction prices of the initiating order-book trades. There are also no significant differences in price impacts. Because a workup follows the initiating order-book trade very closely in time, and because of 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 order-book trade. The result, therefore, indicates that the additional price impact of a workup is close to zero.

Overall, our results show that size-discovery trading protocols attract liquidity-motivated trading and contribute to lowering overall transaction costs and price impacts of D2D trades.

2.6.4 Estimates of Profits from Liquidity Provision

We use the GFI data to estimate dealer profits from liquidity provision 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

2.6. Why is Trade in the Interdealer Market Cheaper?

	CDX.IG			CDX.HY		
	EffcSprd	RlzdSprd	PrcImp	EffcSprd	RlzdSprd	PrcImp
CONST	0.102** (30.43)	-0.013 (-1.93)	0.115** (15.75)	0.476** (12.92)	-0.043 (-0.72)	0.519** (9.50)
MTCH	-0.040** (-9.95)	0.003 (0.47)	-0.043** (-5.72)	-0.152** (-5.10)	0.008 (0.15)	-0.160** (-3.24)
WRKUP	0.005 (1.29)	0.001 (0.12)	0.005 (0.53)	0.016 (0.84)	-0.032 (-0.85)	0.048 (1.27)
OTHER	0.040* (2.04)	0.098** (3.34)	-0.058** (-3.47)	0.354* (2.11)	0.575** (3.22)	-0.222** (-3.24)
MDM	0.010** (3.22)	0.008 (1.31)	0.002 (0.28)	0.002 (0.08)	0.048 (1.08)	-0.046 (-1.17)
LRG	0.007 (1.08)	0.019* (1.97)	-0.013 (-1.29)	0.002 (0.10)	0.033 (0.66)	-0.030 (-0.61)
BLCK	-0.002 (-0.09)	0.063* (2.24)	-0.065** (-2.75)	-0.025 (-0.34)	0.345 (1.94)	-0.370* (-2.43)
RFRNC	0.015** (2.90)	0.003 (0.22)	0.012 (1.02)	0.060 (1.33)	0.092 (1.29)	-0.032 (-0.54)
BAS	0.312** (4.24)	0.150 (1.50)	0.161 (1.22)	0.216** (4.96)	0.079 (0.87)	0.137 (1.50)
SPRD/100	0.119 (1.95)	-0.020 (-0.19)	0.140 (1.95)	0.306 (1.64)	0.377 (1.73)	-0.071 (-0.77)
VLTITY	0.138** (3.22)	-0.017 (-0.20)	0.156 (1.58)	0.486 (1.44)	-1.347* (-1.98)	1.833** (2.99)
<i>N</i>	6,623	6,623	6,623	6,844	6,844	6,844

Table 2.10: 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 outright order-book trades, mid-market matches, workups, and other 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 index CDS 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 trade workups (WRKUP), for trades with an other method of trade execution (OTHER), 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), for trades with transaction prices at typical reference levels (RFRNC; index CDS 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 (VLTITY). 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 1,336 (3,640) [1,187] {460} and 1,102 (4,261) [1,032] {449} outright order-book trades (mid-market matches) [trade workups] {other trades} in five-year on-the-run index CDSs on CDX.IG and CDX.HY, respectively.

Chapter 2. Market Structure and Transaction Costs of Index CDSs

at the best offer (bid) that prevails at trade execution on the order book.⁴³

In the first scenario estimated profits are USD 0.433 million and USD 0.808 million per day in case of CDX.IG and CDX.HY, respectively, or USD 1.241 million per day 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.⁴⁴ This suggests that liquidity providers only make profits through their willingness to bear inventory risk (see, e.g., Grossman and Miller (1988)).

The results show that liquidity provision in the D2C segment is very competitive and that institutional investors who value immediacy would not be able to save transaction costs by executing their trades on the order books of IDB SEFs. 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 best bid or offer.⁴⁵ Transaction costs can only be reduced at the expense of execution certainty either through liquidity supplying order-book trades or through mid-market matching.

2.7 Conclusion

Using transaction data, we study the market structure and transaction costs of index CDSs after the implementation of the Dodd-Frank Act. We identify D2C trades and D2D trades. Transaction costs and price impacts are larger for D2C trades and increase with trade size, quoted bid-ask spread, and volatility. D2C trades Granger-cause D2D trades consistent with the interdealer market being used for managing inventory risk. Unique order-book data show the important role of mid-market matching and workup for reducing transaction costs and price impacts of D2D trades. D2C trades are competitive relative to executable bids and offers in the interdealer market, suggesting that the current two-tiered market structure delivers favorable prices for customers who value immediacy. While these customers would not be able to save transaction costs by executing their trades on interdealer order books, transaction costs could be reduced at the expense of execution certainty either through liquidity supplying order-book trades or through mid-market matching.

⁴³We 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.

⁴⁴Trades 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 the GFI data does not include the depth available at the best bid and offer.

⁴⁵In 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.

3 Index CDS Trading Costs around the Introduction of SEFs

I document a decline of trading costs and profits from liquidity provision in the index credit default swap market over a two-and-a-half-year period during which Dodd-Frank Act provisions were implemented. Consistent with better comparison shopping and stronger price competition on regulated trading platforms, I find lower trading costs and profits from liquidity provision for trades executed on swap execution facilities (SEFs) in comparison to bilaterally negotiated trades. The results suggest that Dodd-Frank rules introducing SEFs had a diminishing effect on trading costs and profits from liquidity provision.

3.1 Introduction

In this paper, I document a reduction of trading costs in the billion dollar index credit default swap (CDS) market over the course of a two-and-a-half-year period during which the U.S. Commodity Futures Trading Commission (CFTC) implemented Dodd-Frank Act provisions that regulate trade in this formerly unregulated over-the-counter (OTC) market. I provide evidence in support of lower profits from liquidity provision driving the decline in the cost of trading. These results are consistent with a statutory goal of the Dodd-Frank Act to promote pre- and post-trade transparency in that improving comparison shopping and increasing competition among liquidity providers constitute key functions of pre- and post-trade transparency in OTC markets.¹

In this respect, the CFTC enacted three important requirements: the trade reporting requirement, the minimum trading functionality requirement, and the trade execution requirement. The trade reporting requirement mandates real-time trade reporting for regulatory purposes and public dissemination of trade details, the minimum trading functionality requirement mandates that regulated trading platforms (so-called swap execution facilities (SEFs)) operate order books for all the swaps in which they offer trading, and the trade execution requirement mandates on-SEF trade execution of so-called required transactions in made available to trade

¹Other key functions of pre- and post-trade transparency in OTC markets are price discovery and monitoring agent-delegated trade execution.

Chapter 3. Index CDS Trading Costs around the Introduction of SEFs

(MAT) swaps by prescribed methods of trade execution that are limited to the order book and the request-for-quote (RFQ) trading protocol. The latter two requirements are the starting point of this paper.

First, I find a significant decline of trading costs and profits from liquidity provision around the effective date of the minimum trading functionality requirement when SEFs were introduced. Trading costs of both required and non-required (or permitted) transactions are larger in the 60-calendar-day period before the effective date of the minimum trading functionality requirement than in the period between the effective dates of the minimum trading functionality and trade execution requirements. For instance, the trading costs of required transactions in CDX.IG contracts (i.e., contracts on broad-based credit indices that are composed of North American investment-grade companies) decrease from 0.287 bps to 0.180 bps and those of permitted transactions decrease from 0.295 bps to 0.269 bps.

Second, using a difference-in-differences approach, I do not find an effect of mandatory on-SEF trade execution on trading costs. Trading costs of both required and permitted transactions in CDX.IG decrease further to 0.156 bps and 0.215 bps in the 60-calendar-day period after the effective date of the trade execution requirement. But trading costs of transactions that are subject to the trade execution requirement do not decrease significantly more than those of transactions that are not subject to the requirement.

Similarly, profits from providing liquidity in both required and permitted transactions are higher in the period before the effective date of the minimum trading functionality requirement than in the period between the effective dates of the minimum trading functionality and trade execution requirements. Per trade profits from liquidity provision decrease further in the period after the effective date of the trade execution requirement but, in comparison to permitted transactions, not significantly more for required transactions. Thus, a difference-in-differences approach does not reveal an effect of mandatory on-SEF trade execution on profits from liquidity provision.

These results suggest that the minimum trading functionality requirement had an effect on trading costs and per trade profits from liquidity provision. But anecdotal evidence is that SEF order books failed to attract liquidity.² This suggests that, rather than the specific minimum trading functionality, it is the pre-trade transparency provided for by regulated trading platforms in general that drives the observed declines in trading costs and per trade profits from liquidity provision. In this respect, the results suggest that, once pre-trade transparency has been provided for, there is no incremental effect on trading costs and profits from liquidity provision associated with mandatory pre-trade transparency in the form of a trade execution requirement. The results are robust to controlling for trade characteristics as well as market liquidity and volatility and also obtain for transactions in CDX.HY contracts (i.e., contracts on broad-based credit indices that are composed of North American high-yield companies).

²See “RFQ vs. CLOB—The Battle So Far,” Blog post, *TabbFORUM*, May 27, 2015.

SEFs provide for pre-trade transparency because they offer methods of trade execution that, in comparison to bilateral negotiations, facilitate comparison shopping and create direct price competition among liquidity providers. Therefore, I compare trading costs and profits from liquidity provision of trades that are executed on SEFs (on-SEF) with those that are not (off-SEF). I focus on the period after the effective date of the minimum trading functionality requirement but before the effective date of the trade execution requirement when trading on SEFs was voluntary for both required and permitted transactions. The differences are striking. For the most actively traded five-year on-the-run (5Y OTR) contracts, on-SEF trading costs are 40%–50% lower than off-SEF trading costs and the per trade profits from on-SEF liquidity provision are less than 25% of those from off-SEF liquidity provision. Similarly, for the less actively traded five-year immediate off-the-run (5Y OFF) contracts, on-SEF trading costs are 37% lower than off-SEF trading costs and the per trade profits from on-SEF liquidity provision are 29%–54% of those from off-SEF liquidity provision. For other contracts, there are no statistically discernable differences in on-SEF and off-SEF trading costs and profits from liquidity provision.

In contrast to many other OTC markets, the index CDS market is not characterized by a complete absence of pre-trade transparency. Instead, credit derivatives dealers provide their institutional clients via electronic pricing messages (usually e-mails) with indicative bid and offer quotes on a variety of index CDSs. While quotes are indicative and for no particular notional amount, there is an implicit understanding that an instrument's standard notional amount can be executed at or near these quotes without additional bargaining.³ The frequency with which quotes are updated heavily depends on the particular contract. Quotes on 5Y OTR and 5Y OFF contracts are updated several hundred times a day (on aggregate across dealers) while other contracts are only quoted a few times per day, presumably upon a client's request. Consequently, in case of being directly contacted by a client for trade in a 5Y OTR or 5Y OFF contract, the dealer can be relatively sure that her indicative quotes were aggressive in comparison to the rest of the market.⁴ Thus, the dealer is likely to win the trade and, therefore, has no incentive to improve upon her quotes. In contrast, when queried for quotes via a SEF, the dealer is simultaneously competing with other dealers and uncertain about her odds of winning the trade and, therefore, may improve upon her quotes.⁵ Consistently, I find that on-SEF trades in 5Y OTR and 5Y OFF contracts are significantly more likely to get executed within the quoted composite bid-ask spread that prevails at trade execution than off-SEF trades.

The results concerning differential trading costs and profits from liquidity provision of on-SEF

³Some dealers even maintain so-called “dealer-run” platforms with two-sided-market streams that allow for trade initiation at the streamed quotes. Trade execution is, however, at the dealer's discretion and access to platforms is usually restricted to a dealer's most important clients.

⁴Many institutional investors have access to the quotes of several credit derivatives dealers.

⁵While during the period under consideration there has not been a requirement for querying a minimum number of dealers in case of required transactions, best execution practice among market participants seems to be seeking quotes from two to three dealers (see “The SEF RFQ Minimum is Moving to 3. Does it matter? Nope,” Blog post, *Greenwich Associates*, July 17, 2014, for survey-based evidence from the interest rate swap market).

Chapter 3. Index CDS Trading Costs around the Introduction of SEFs

and off-SEF trades are robust to controlling for trade characteristics and the endogenous choice of trading venue. The choice model that I estimate reveals that the likelihood of on-SEF trade execution decreases with trade size and that the more actively traded 5Y OTR and 5Y OFF contracts are more likely to be executed on SEFs than other contracts. Despite the fact that illiquidity and volatility tend to be highly correlated, they have opposite effects on the likelihood of on-SEF trade execution: the likelihood decreases with the quoted bid-ask spread of the 5Y OTR contract but increases with the at-the-money implied volatility of short-term index options on the same contract. The rationale is that fast trade execution on SEFs is important when volatility is high, while low information leakage in bilateral negotiations is important when liquidity is low.

The paper is related to a number of studies analyzing the impact of Dodd-Frank Act provisions on swap market liquidity and trading costs. Loon and Zhong (2016) document a positive impact of various aspects of the Dodd-Frank reform package (such as, mandatory trade reporting, central clearing, and trade on SEFs) on index CDS liquidity in the first year following the trade reporting requirement. Relative to what has been documented by Loon and Zhong (2016), I provide evidence of a longer term trend of declining trading costs in the Dodd-Frank regulatory regime that seems primarily due to lower profits from liquidity provision. I also consider an aspect of the Dodd-Frank Act that Loon and Zhong (2016) do not, namely, mandatory pre-trade transparency due to the trade execution requirement. Benos et al. (2016) find positive impacts of the minimum trading functionality requirement and the trade execution requirement on trading costs of interest rate swaps. In contrast, I do not find an effect of mandatory on-SEF trade execution on trading costs of index CDSs. My analysis also differs from Benos et al. (2016) in how trading costs are measured. I directly measure trading costs at the transaction level by effective half-spreads, thereby, exploiting trade report data in a more comprehensive manner and being able to include per trade profits from liquidity provision in my analysis. Finally, Collin-Dufresne, Junge, and Trolle (2016) focus on the structure of the index CDS market that, in spite of the implementation of Dodd-Frank Act provisions, continues to be two-tiered. They show that differences in on-SEF trading costs of client and interdealer trades are due to low price impact of interdealer trades that serve to manage inventory risk. In comparison, I focus on differences in trading costs of on-SEF and off-SEF trades and provide evidence in support of relatively stronger price competition on SEFs, which is consistent with the competitive pricing of on-SEF client trades documented by Collin-Dufresne et al. (2016).

The paper is also related to studies analyzing the effect of regulations that enforce transparency upon OTC market. Bessembinder et al. (2006), Edwards et al. (2007), and Goldstein et al. (2007) document positive effects of mandatory post-trade transparency on the costs of trading corporate bonds. Asquith, Covert, and Pathak (2013), on the other hand, find a differential effect of mandatory post-trade transparency on the liquidity of thinly traded high-yield corporate bonds in that both price dispersions and trading volumes drop upon trade reporting becoming mandatory. In comparison, my results suggest that providing for pre-trade transparency has an effect on trading costs, while there is no incremental effect of mandatory pre-trade

transparency.

The results concerning the choice between on-SEF and off-SEF trade execution are consistent with other studies of mechanism choice and venue selection in that (off-SEF) bilateral negotiation tends to be chosen for off-the-run instruments (see, e.g., Barclay, Hendershott, and Kotz (2006)), for larger and less standard trades (see, e.g., Hendershott and Madhavan (2015)), and when spreads are wide and volatility is low (see, e.g., Bessembinder and Venkataraman (2004)). However, in contrast to Bessembinder and Venkataraman (2004) and Hendershott and Madhavan (2015), I do not find evidence for strategic selection of the lower cost trading venue. This could be due to significant costs associated with SEF onboarding and compliance that are not captured by the simple effective half-spread measure of trading costs used in the analysis.

The paper is organized as follows. Section 3.2 summarizes the relevant parts of the Dodd-Frank Act and discusses their potential impact on profits from liquidity provision. Section 3.3 provides institutional details about the index CDS market and describes the trade and quote data. Section 3.4 presents results of the analyses of trading costs and profits from liquidity provision and, where necessary, supplements methodological details. Section 3.5 concludes the paper.

3.2 The Dodd-Frank Act

Title VII of the Dodd-Frank Act provides for a regulatory reform of U.S. OTC swap markets with the objective to promote financial stability, pre- and post-trade transparency, and the trading on SEFs. For the majority of swaps, the regulatory agency charged with the implementation of the Act was the CFTC. In order to promote post-trade transparency, the CFTC enacted a trade reporting requirement providing for real-time trade reporting and public dissemination of transaction data. Effective December 31, 2012, swap dealers were required to report their trades to so-called swap data repositories that collect and publicly disseminate the transaction data. Other market participants were required to report their swap trades in subsequent months.⁶ In order to promote pre-trade transparency, the CFTC enacted a minimum trading functionality and a trade execution requirement both of which are tightly linked to a new type of regulated trading platform, the SEF, that was introduced in order to regulate trading in swap markets.

3.2.1 Minimum Trading Functionality and Trade Execution Requirements

A SEF is “a trading system or platform in which multiple participants have the ability to execute or trade swaps by accepting bids and offers made by multiple participants in the

⁶Specifically, so-called major swap market participants were required to report their swap trades from February 28, 2013 onwards, and all other market participants were required to report their trades from April 10, 2013 onwards. The stated dates apply to interest rate and index credit default swaps only.

Chapter 3. Index CDS Trading Costs around the Introduction of SEFs

facility or system” (see Section 1(a)(50) of the Commodity Exchange Act (CEA)). Essentially, the definition ensures that any person or trading platform that facilitates the execution of swaps and is subject to CFTC oversight has to comply with SEF regulations. Compliance with SEF regulations amongst others requires that SEFs operate order books as minimum trading functionalities for the swaps that they list for trading. But trades on SEFs do not necessarily have to be executed on the order book because SEFs are allowed to offer other methods of trade execution in addition to the order book. In fact, trades, in general, do not have to be executed on SEFs at all. Only some trades in MAT swaps, so-called required transactions, have to be executed on SEFs due to the trade execution requirement. Specifically, the requirement applies to all non-block trades in MAT swaps that are not packaged with a non-MAT swap and where none of the counterparties is an end-user hedging commercial risks (see below). In case that the requirement applies, it requires that trades are executed either against an order resting on the order book or against a response to a RFQ which was transmitted to at least three other market participants. Trade execution requirements come into effect whenever a swap is made available to trade by means of a SEF-initiated MAT determination that is consistent with the CEA and CFTC regulations. Once being in effect, the requirements apply to all SEFs and not only the one that filed the MAT determination. The minimum trading functionality requirement came into effect on the compliance date of SEF regulations, October 2, 2013, and trade execution requirements for MAT swaps came into effect in February 2014.

As mentioned above, the trade execution requirement contains some exceptions. First, it allows for off-SEF execution of so-called block trades.⁷ Block trades are large-sized trades distinguished by notional amounts that exceed pre-defined thresholds.⁸ A block-sized trade executed on an order book would typically have a large price impact and, therefore, the requirement allows for off-SEF trade execution. Similarly, a block-sized trade executed via an RFQ that is disseminated to at least three other market participants would typically have a larger price impact than if it were bilaterally negotiated because of information leakage associated with requesting quotes from multiple dealers.

Second, end-users that trade swaps for hedging commercial risks are exempt from the trade execution requirement and, therefore, do not have to execute their trades on SEFs. This is because costly compliance with regulatory rules would discourage such end-users from hedging via swaps.

Third, MAT swaps that are packaged with non-MAT swaps are temporarily exempt from the trade execution requirement because SEFs and central clearing counterparties were unable to handle the processing necessary to guarantee clearing of all legs of a package upon trade execution when the trade execution requirements for MAT swaps came into effect. This would

⁷In fact, off-SEF trade execution is part of the definition of a block trade. This, however, does not mean that on-SEF trade execution of block-sized trades is not allowed, it only means that such trades would not be disseminated with a delay. Moreover, that part of the definition of a block trade was temporarily overruled by a no-action relief that the CFTC granted on September 19, 2014.

⁸Notional thresholds are set by the CFTC with the intention that around 50% of the aggregate notional amount traded is disseminated in real-time.

leave the parties to the trade with the risk that one leg of their trade would not be accepted for clearing and, as a consequence, be declared invalid from the beginning during the post-trade processing stage. Because many packages involve legs that partially offset each other's risk, the parties' risk profile would ultimately be more risky than intended upon trade execution.

3.2.2 Potential Impact of Requirements on Profits from Liquidity Provision

In order to understand how and why the two requirements might affect profits from liquidity provision, it is worth contrasting swap trading in the presence of SEF regulations with how trading used to be when swap markets were OTC and, to a large extent, exempt from regulatory oversight. The most obvious difference that trade in OTC swap markets was characterized by a complete absence of post-trade transparency will not be discussed because real-time trade reporting was already in effect when SEFs were introduced.

First, trade in OTC swap markets was dealer intermediated and, as a consequence, there was at least one dealer counterparty on each and every trade. The minimum trading functionality requirement in principle allows for liquidity provision by non-dealer market participants, with an adverse effect on profits from liquidity provision due to increased competition.

Second, the pre-trade information available to traders in OTC swap markets comprised of indicative dealer quotes (alike the electronically distributed ones described in the introduction) as well as quotes collected by directly contacting a dealer over the phone. The latter were in general only firm "as long as the breath is warm," i.e., they could be subject to change upon a repeat contact. Because a repeat contact signals less favorable valuations of competing dealers, it gives rise to strategic pricing behavior by dealers that generally leads to less favorable terms offered upon a repeat contact (see, e.g., Zhu (2012)). In comparison, the trade execution requirement ensures that a significant share of trades takes place on electronic trading platforms and by means of trading protocols ensuring that quotes are executable and come from multiple liquidity providers simultaneously competing for the trade. This eliminates the possibility for strategic dealer pricing upon a repeat contact and increases dealer competition relative to traditional bilateral negotiations in OTC markets. Both effects are likely to reduce profits from liquidity provision.

3.3 Institutional Details and Data

Focus of the empirical analysis are index CDSs. This section provides the necessary institutional background about the market in which these contracts trade. Moreover, it describes the data that are used in the analysis.

3.3.1 The Index CDS Market

An index CDS is a standardized credit derivative on an index of creditors (i.e., a credit index)

Chapter 3. Index CDS Trading Costs around the Introduction of SEFs

that generates the same payoff as a diversified portfolio of the creditors' single-name CDSs. Specifically, over the life of the index CDS contract, the credit protection seller agrees to compensate the credit protection buyer for the loss associated with each credit event that pertains to one of the creditors in the index and with that fraction of the contract's notional amount which is proportional to the creditor's weight in the index.⁹ In return, the credit protection buyer agrees to make fixed-rate quarterly premium payments on that fraction of the contract's notional amount which remains after all preceding credit events have been written-off completely.¹⁰

Index composition follows a rules-based approach according to which creditors in the index are revised every six months. The revisions ensure that the index is composed of those creditors of a specific credit rating grade whose single-name CDSs are most liquidly traded in the market. Whenever there is a revision, a new index—or, more precisely, a new series of an index—is launched.¹¹ The most recently launched index is referred to as on-the-run and all previously launched indices are referred to as off-the-run. Liquidity and trading activity concentrates in on-the-run index CDSs and, in particular, those with a five-year maturity.

Index CDSs used to be traded in a relatively opaque OTC market. With the implementation of Dodd-Frank Act provisions and, in particular, with the approval of a MAT determination for 5Y OTR and 5Y OFF contracts on the most popular credit indices, the index CDS market has undergone significant changes. Since the effective date of the trade execution requirement, February 26, 2014, the majority of trades are executed on SEFs. Nevertheless, the structure of the index CDS market has been largely unaffected and trades continue to be dealer intermediated due to almost exclusive use of RFQ methods for trade execution (see, e.g., Collin-Dufresne et al. (2016)).

The two most popular credit indices of North American creditors with investment-grade and high-yield ratings are the CDX North American Investment Grade (CDX.IG) and the CDX North American High Yield (CDX.HY) index. The former comprises 125 creditors and the latter comprises 100 creditors. Index CDS contracts on these indices have maturities between one and ten years and are traded with fixed spreads and points upfront. That is, the premium payments of all contracts on a particular index accrue at the same fixed rate (the fixed spread specified in the contract terms) and counterparties exchange an upfront payment (the points upfront per dollar of notional amount) at trade initiation in order to compensate each other for the potentially non-zero present value of the contract. Nevertheless, contracts are rarely quoted in terms of points upfront. Instead, CDX.IG contracts are usually quoted in terms of par spreads (i.e., in terms of the fixed rates that give zero present values of contracts; henceforth referred to as index CDS spreads) and CDX.HY contracts are usually quoted in terms of bond-

⁹For instance, a 60% default loss of a creditor in an equally-weighted index of 100 creditors corresponds to an 0.6% loss on the notional of the index CDS contract.

¹⁰Continuing with the example of one creditor's default in an equally-weighted index of 100 creditors, following the default, premium payments are made on 99% of the notional whereas they were made on the full notional of the index CDS contract before the default.

¹¹An index's series number uniquely determines the creditors in the index.

equivalent prices. Contracts on CDX.IG and CDX.HY are the focus of my analysis.

3.3.2 Data and Descriptive Statistics

For the analysis, I use trade and quote data over the 611-trading-day period from January 2, 2013 to June 30, 2015. The trade data come from publicly disseminated trade reports by the Bloomberg Swap Data Repository, the Depository Trust & Clearing Corporation Data Repository, and the Intercontinental Exchange Trade Vault (Appendix C contains details concerning the sample construction). The trade reports comprise, amongst others, information to identify the underlying of the index CDS contract, the contract term, the execution timestamp, the transaction price, the (capped) trade size, and an on-SEF trade execution indicator. The quote data come from Markit and comprise time-stamped dealer composite bid and ask quotes. Composite quotes are updated on an intraday basis whenever a dealer sends out an electronic pricing message with indicative quotes to one of the subscribers to Markit's quote parsing and streaming service. In order to preserve pricing message anonymity, Markit computes composites using individual-dealer quotes from each dealer's latest pricing message within the past 15 minutes. Quote timestamps correspond to the time when the updating pricing message was made.

As can be seen from Table 3.1, the data exhibit significant differences in the trading and quoting activity among contracts.¹² Trading activity, both in terms of the notional amount traded and in terms of the number of trades, concentrates in 5Y OTR index CDS contracts. For instance, 5Y OTR CDX.IG accounts for 85% of total volume in CDX.IG contracts and for 87% of total trades, and 5Y OTR CDX.HY accounts for 86% of total volume in CDX.HY contracts and for 91% of total trades. Similarly, dealer quotes on 5Y OTR contracts alone account for around 60% of all quotes in case of both indices.

Daily trading volumes of 5Y OTR contracts are large. On an average trading day, the aggregate notional amount of 5Y OTR CDX.IG contracts is USD 7.8 billion and that of 5Y OTR CDX.HY contracts is USD 3.2 billion. Particularly noteworthy is the fact that daily trading volumes are due to a relatively small number of trades, with only 138 5Y OTR CDX.IG and 169 5Y OTR CDX.HY trades per day, on average. As a consequence of high daily trading volumes and relatively few trades per day, it follows that trade sizes have to be large. It should also be noted that actual trading volumes are even larger because reported volumes are downward biased due to capped trade sizes and due to trading of non-U.S. institutions that are not subject to the CFTC's trade reporting requirement.

Although being low-high frequency in comparison to other markets (such as, the equity or foreign exchange markets where trading and quoting takes place at a millisecond frequency),

¹²The number of contracts, N , in the Total row does not add up to those of the 5Y OTR, 5Y OFF, and Other rows because five of the twelve contracts that are either 5Y OTR or 5Y OFF become further off-the-run during the sample period (i.e., they become what I categorize as Other). Netting those out of the 83 and 67 Other CDX.IG and CDX.HY contracts and adding the remaining seven 5Y OTR or 5Y OFF contracts gives the Total row value.

Contract	N	Total			Daily Volume			Trades Per Day			Quotes Per Day					
		Volume	Trades	Quotes	\bar{N}_V	Mean	10 th	90 th	Mean	10 th	90 th	Mean	10 th	90 th		
Panel A: CDX.IG																
5Y OTR	6	4,764	84,447	299,041	1.0	7.8	3.8	12.8	1.0	138	74	221	1.0	489	330	674
5Y OFF	6	413	6,330	189,953	1.0	0.7	0.0	1.4	0.9	10	0	22	1.0	311	191	448
Other	83	458	6,087	16,445	57.1	0.7	0.1	1.5	4.3	10	2	20	10.1	27	2	56
Total	85	5,635	96,864	505,439	59.1	9.2	4.5	14.7	6.2	159	84	246	12.1	827	558	1151
Panel B: CDX.HY																
5Y OTR	6	1,977	103,262	246,954	1.0	3.2	1.3	5.9	1.0	169	74	295	1.0	404	288	543
5Y OFF	6	227	7,455	142,444	1.0	0.4	0.0	0.8	0.9	12	1	24	1.0	233	106	352
Other	67	99	2,937	28,722	49.8	0.2	0.0	0.4	2.2	5	0	11	15.8	47	11	90
Total	69	2,303	113,654	418,120	51.8	3.8	1.5	6.6	4.1	186	84	323	17.7	684	476	925

Table 3.1: Descriptive Statistics Trade and Quote Data.

Panels A and B show descriptive statistics of the trade and quote data for CDX.IG and CDX.HY, respectively. Descriptive statistics are computed separately for five-year on-the-run (5Y OTR), five-year immediate off-the-run (5Y OFF), and all other (Other) index CDS contracts. N is the aggregate number of index CDS contracts, Volume is the aggregate trading volume (in USD billion), Trades is the aggregate number of trades, and Quotes is the aggregate number of quotes. \bar{N}_V is the time series mean of the aggregate number of non-matured index CDS contracts, \bar{N}_T is the time series mean of the aggregate number of traded index CDS contracts, and \bar{N}_Q is the time series mean of the aggregate number of quoted index CDS contracts. The sample period is January 2, 2013 to June 30, 2015.

quote revisions for 5Y OTR contracts occur relatively frequently. For instance, the average numbers of 489 and 404 quotes per day on 5Y OTR CDX.IG and CDX.HY contracts, respectively, suggest that during a ten-hour trading day quotes are updated every one to one and a half minutes.

The 5Y OFF contract, i.e., the five-year contract referencing the previously launched index series, also accounts for a significant share of trading volume (7% for CDX.IG and 10% for CDX.HY) and trades (7% in case of both indices). On an average trading day, there are ten 5Y OFF CDX.IG and twelve 5Y OFF CDX.HY trades but around 10% of trading days see no trading at all. Nevertheless, dealer quote activity remains relatively high with 200 to 300 quote updates per day.

The residual trading volume is due to trading in non-five-year contracts on the on-the-run and immediate off-the-run series and contracts on further off-the-run series (collectively referred to as Other contracts in what follows).¹³ On an average trading day, there are around 57 such CDX.IG contracts and 50 such CDX.HY contracts but less than a handful of those actually trade. However, the few ones that trade seem to trade more than once which could, for instance, be due to hedging activity by the dealer that facilitated the trade.¹⁴ Similarly, relatively few contracts are quoted but those that are quoted have more than one quote per day, on average, which is consistent with OTC market practice of getting quotes from a few dealers prior to trade execution.

The differences in trading and quoting activity lead to disproportional large shares of trades in Other contracts that are being dropped when combining the trade and quote data. The combined data consists of all trades with an available intraday (i.e., from the same trading day) quote prior to trade execution and another intraday quote from an update that follows trade execution by at least 15 minutes. The 15-minute period is chosen for two reasons: first, it takes time for information to get incorporated into quotes; and second, block trades are disseminated with a delay of 15 minutes and the 15-minute period ensures that the quotes are updated when trade occurrence is public knowledge.¹⁵ Moreover, the 15-minute period ensures that Markit composite quotes before and after trade execution are based on distinct sets of dealer quotes.

The combined data set consists of 90,983 CDX.IG and 108,986 CDX.HY trades. The trades comprise 83,550 5Y OTR, 6,002 5Y OFF, and 1,431 Other CDX.IG trades and 100,667 5Y OTR,

¹³Even within these inactively traded contracts, trading activity is relatively concentrated. For instance, in case of CDX.IG two of the 83 other contracts, namely, seven- and ten-year CDX.IG series 9, account for 35% of trades and 43% of trading volume of the remaining contracts. Series 9 is the last series that was launched before the financial crisis when tranche swaps were particularly popular and, anecdotally, trading activity is due to hedging of the considerable amount of outstanding tranche swaps.

¹⁴On a per-contract basis, the sample means reported in Table 3.1 suggest that the average Other CDX.IG contract trades $10/57.1 = 0.2$ times per day or once a week and the average Other CDX.HY contract trades $5/49.8 = 0.1$ times per day or once every two weeks.

¹⁵This is only the case for some part of the sample period because prior to July 30, 2013 all trade were disseminated with a delay of 30 minutes. Moreover, for the first year following the compliance date of the CFTC's real-time trade reporting requirement (i.e., December 31, 2012) block trades were disseminated with a delay of 30 minutes.

7,070 5Y OFF, and 1,249 Other CDX.HY trades. These correspond to 99%, 95%, and 24% of the CDX.IG trades and 97%, 95%, and 43% of the CDX.HY trades, respectively, that are exhibited in Table 3.1. Given prevailing quotes, I classify trades at index CDS spreads greater than the prevailing mid-quote as protection buyer initiated and trades at spreads less than the prevailing mid-quote as protection seller initiated. Following Lee and Ready (1991), trades at the mid-quote are classified using the tick rule with trades on an up-tick being classified as protection buyer initiated and trades on a down-tick being classified as protection seller initiated.

3.4 Dodd-Frank Regime Trading Costs

Trading costs are measured by effective half-spreads with respect to the prevailing mid-quote. Spreads reflect both the liquidity provider's revenue and the trade's information content and can be decomposed accordingly:

$$\underbrace{q_t(p_t - m_t)}_{\text{=Effective Half-Spread}} = \underbrace{q_t(p_t - m_{t+\Delta})}_{\text{=Realized Half-Spread}} + \underbrace{q_t(m_{t+\Delta} - m_t)}_{\text{=Price Impact of Trade}}, \quad (3.1)$$

where p_t is the transaction price, m_t is the mid-quote prevailing at trade execution t , and $m_{t+\Delta}$ is the mid-quote of the first quote update at least 15 minutes after trade execution. Both transaction prices and quotes are in terms of index CDS spreads and expressed in basis points. Trade direction, q_t , equals +1 (−1) for protection-buyer-initiated (protection-seller-initiated) trades.

The intuition for decomposing effective half-spreads according to Equation (3.1) is as follows: the effective half-spread measures the liquidity provider's revenue if he were able to immediately close his position at the prevailing mid-quote. If instead it takes the liquidity provider at least 15 minutes to close his position (and again assuming that he is able to do so at the then prevailing mid-quote), his revenue is the realized half-spread. The liquidity provider's revenue is less than the effective half-spread if the price moves against him while he is reversing the trade over time. Such trade-induced price moves or adverse selection costs are captured by the trade's price impact.

Panels A and B of Figure 3.1 shows monthly averages of daily trade-size weighted effective and realized half-spreads of 5Y OTR CDX.IG and CDX.HY, respectively, over the two-and-a-half-year period following the effective date of the CFTC's trade reporting requirement. Focusing on the effective half-spread, the panels show a large and almost steady decline in the trading costs of 5Y OTR CDX.IG and CDX.HY, respectively. Half-spreads of both indices compressed by about 46% from 0.237 bps in January 2013 to 0.131 bps in June 2015 for 5Y OTR CDX.IG and, similarly, from 1.210 bps to 0.647 bps for 5Y OTR CDX.HY. The spread compression is consistent with the generally positive impact that various aspects of the Dodd-Frank reform package (such as mandatory trade reporting, central clearing, and trade on SEFs) had on a

3.4. Dodd-Frank Regime Trading Costs

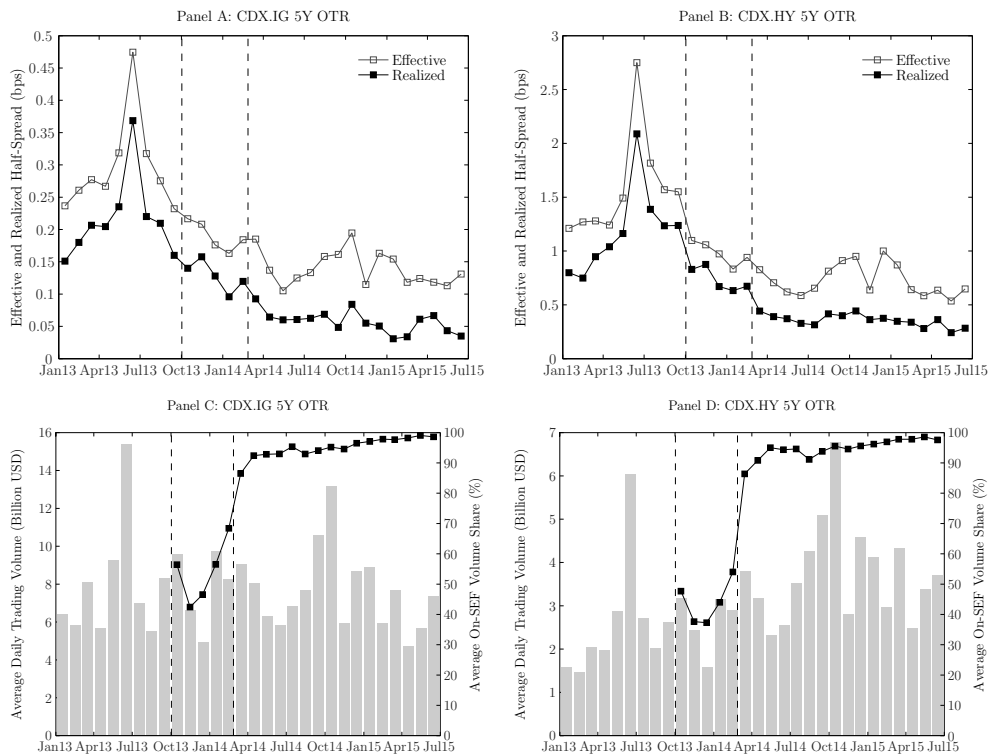


Figure 3.1: Effective and Realized Half-Spreads and Fractions of On-SEF Volume.

Panels A and B show monthly averages of daily trade-size-weighted effective and realized half-spreads of transactions in five-year on-the-run (5Y OTR) index CDSs on CDX.IG and CDX.HY, respectively. Panels C and D show monthly averages of daily trading volumes (gray bars, left hand scales) of transactions in 5Y OTR index CDSs on CDX.IG and CDX.HY, respectively, as well as monthly averages of the daily volume share of transactions executed on SEFs (black squares, right hand scales). The sample period is January 2, 2013 to June 30, 2015. The dashed vertical lines correspond to (from left to right) the effective date of the minimum trading functionality requirement, October 2, 2013, and the effective date of the trade execution requirement, February 26, 2014.

variety of liquidity measures in the first year of real-time trade reporting, as documented by Loon and Zhong (2016). Relative to what has been documented by Loon and Zhong (2016), Figure 3.1 provides evidence of a longer term trend of declining transaction costs in the Dodd-Frank regulatory regime that seems primarily due to lower profits from liquidity provision. This is because realized half-spreads decline almost in lockstep with effective half-spreads from 0.151 bps in January 2013 to 0.035 bps in June 2015 for 5Y OTR CDX.IG and, similarly, from 0.799 bps to 0.283 bps for 5Y OTR CDX.HY.¹⁶ Apart from lower profits from liquidity provision, the correlated decline suggests that price impact, which is the difference between effective and realized half-spreads, remains relatively constant over time.

It should be noted that, in the above, I do not measure half-spread declines from their peak in June 2013 but from the first month of the sample period, January 2013. The peak in June

¹⁶The time series correlation between daily trade-size weighted effective and realized half-spreads is 0.86.

Chapter 3. Index CDS Trading Costs around the Introduction of SEFs

2013 coincides with a bond market sell-off that lasted from May 2013 to July 2013 and was accompanied by widening credit spreads and sharply rising Treasury yields.¹⁷ The increase in credit risk also led to wider index CDS spreads, with 5Y OTR CDX.IG (CDX.HY) rising from 74.3 bps (358.9 bps) on May 2, 2013 to 97.6 bps (477.9 bps) on June 24, 2013 before reverting back to 74.8 bps (369.5 bps) on July 31, 2013.

Panels C and D of Figure 3.1 show monthly averages of daily trading volume for 5Y OTR CDX.IG and CDX.HY, respectively. Trading volumes seem to be unaffected by the minimum trading functionality and trade execution requirements. However, the decomposition of trading volume changed due to the trade execution requirement. As can be read off the panels' right hand scales, the monthly average of the fraction of daily trading volume that is executed on SEFs increased to about 90% in the month following the effective date of the trade execution requirement. The latter was announced on January 28, 2014, about one month before it came into effect. But even before the announcement, on-SEF volume shares among 5Y OTR contracts were relatively high, averaging 47% and 40%, respectively, for CDX.IG and CDX.HY during the third quarter of 2013. Abstracting from the unusually high half-spreads in between May 2013 and July 2013, a comparison of upper and lower panels suggests that the spread compression coincides with an increased use of on-SEF trade execution.

3.4.1 Difference-in-Differences Analysis

In order to quantify the impact of the minimum trading functionality and trade execution requirements, I focus on the period from 60 calendar days before the effective date of the minimum trading functionality requirement to 60 calendar days after the effective date of the trade execution requirement, i.e., August 3, 2013 to April 26, 2014. The pre-event window of 60 calendar days is chosen so as to not overlap with the effective date of the CFTC's block trade rules on July 30, 2013. Prior to that date all trades were disseminated with a delay of 30 minutes while after that date only block trades are disseminated with a delay.¹⁸ Because August 3, 2013 is only three trading days after the effective date of block trade rules, I also experimented with shorter 30-calendar-day pre- and post-event windows in order to rule out confounding effects of block trade rules and found similar albeit weaker results.

As described above, block trades are exempt from the trade execution requirement and so are end-user exempt trades and trades that are packaged with non-MAT index CDSs. In what follows, I will abstract from the fact that the last two types of trades are not subject to the trade execution requirement and consider all non-block trades in 5Y OTR and 5Y OFF index CDSs on CDX.IG and CDX.HY (i.e., those CDX.IG and CDX.HY contracts that are MAT) as required transactions, i.e., as transactions that are subject to the on-SEF trade execution requirement.¹⁹

¹⁷See "The Recent Bond Market Selloff in Historical Perspective," Blog post, *Liberty Street Economics*, August 5, 2013.

¹⁸In addition, a uniform cap of USD 100 million was applied to the trade size of all trades in the period prior to the effective date of block trade rules.

¹⁹In the period after the effective date of the trade execution requirement, 93.2% (5,535 out of 5,942) and 92.5%

To the extent that the requirement did not have an impact on trades that are not subject to the requirement, so-called permitted transactions, misclassifying the two types of trades will make detection of the impact more difficult (i.e., it will bias results against finding an impact).²⁰ In principle, end-user exempt trades could be taken into account because they are flagged in the transaction data. However, the flag exhibits a behavior that cannot be reconciled with primarily institutional participants in the index CDS market. Accordingly, the fraction of end-user exempt trades increases from 36.9% of trades on August 5, 2013 to 73.0% of trades on February 7, 2014. Then it drops to zero on February 10, 2014 and never exceeds 0.7% of trades for the remainder of the sample period. Taken at face value, this suggests either a significant change in the composition of market participants or their trading behavior, neither one of which appears plausible. Therefore, I do not account for the fact that end-user exempt trades are not subject to the trade execution requirement. Because package transactions are not flagged in the transaction data, I cannot account for the fact that these trades are not subject to the trade execution requirement. Block trades in MAT index CDSs and all trades in non-MAT index CDSs form the control group of permitted transactions.

In order to address the question whether the minimum trading functionality and trade execution requirements had an impact on trading costs and profits from liquidity provision, I use a difference-in-differences approach that essentially assesses whether costs and profits of required transactions decline relatively more than those of permitted transactions over the pre- and post-event windows associated with a given event. Specifically, I estimate the following regression

$$y = \beta_0 + \beta_1 \text{MAT} + \beta_2 \text{CMP} + \beta_3 \text{CMP} \times \text{MAT} + \beta_4 \text{EXC} + \beta_5 \text{EXC} \times \text{MAT} + X' \beta + \epsilon \quad (3.2)$$

where y is either the effective or realized half-spread of the t -th trade in index CDS i on date d (dependence on t , i , and d is suppressed for notational convenience), MAT is a dummy variable taking the value one if the t -th trade is a non-block trade in an index CDS i that is MAT (i.e., a required transaction according to the above definition), CMP is a dummy variable taking the value one if the date d on which the t -th trade was executed is on or after the compliance date of SEF rules (October 2, 2013) which include the requirement for a minimum trading functionality, and EXC is a dummy variable taking the value one if the date d on which the t -th trade was executed is on or after the effective date of the trade execution requirement (February 26, 2014), X is a set of control variables, and ϵ is an error term. The coefficient estimate of β_5 is the difference-in-differences estimator of the causal effect of the

(6,049 out of 6,543) of the CDX.IG and CDX.HY trades that I classified as required transactions are actually executed on a SEF.

²⁰Spillover effects might lead to an impact on trading costs and profits from liquidity provision of permitted transactions. As long as this impact is less than the one on required transactions, the statement remains correct. For an example of a regulatory intervention (the full repeal of the uptick rule for short sales by the U.S. Securities and Exchange Commission) associated with significant spillover effects, see Boehmer, Jones, and Zhang (2015).

Chapter 3. Index CDS Trading Costs around the Introduction of SEFs

trade execution requirement on trading costs and profits from liquidity provision in that

$$\begin{aligned} \beta_5 = & \mathbb{E}[y|\text{CMP} = 1, \text{EXC} = 1, \text{MAT} = 1, X] - \mathbb{E}[y|\text{CMP} = 1, \text{EXC} = 0, \text{MAT} = 1, X] \\ & - (\mathbb{E}[y|\text{CMP} = 1, \text{EXC} = 1, \text{MAT} = 0, X] - \mathbb{E}[y|\text{CMP} = 1, \text{EXC} = 0, \text{MAT} = 0, X]). \end{aligned} \quad (3.3)$$

That is, β_5 is the difference of mean differences of effective or realized half-spreads of required and permitted transactions, respectively, before and after the effective date of the trade execution requirement for MAT index CDSs, where before means before the effective date of the trade execution requirement but after the SEF compliance date. Thus, a significantly negative β_5 indicates a spread-compressing effect of the trade execution requirement. Similarly, β_3 is the difference of mean differences of effective or realized half-spreads of required and permitted transactions, respectively, before and after the compliance date of SEF rules, where after means after the compliance date of SEF rules but before the effective date of the trade execution requirement; that is,

$$\begin{aligned} \beta_3 = & \mathbb{E}[y|\text{CMP} = 1, \text{EXC} = 0, \text{MAT} = 1, X] - \mathbb{E}[y|\text{CMP} = 0, \text{EXC} = 0, \text{MAT} = 1, X] \\ & - (\mathbb{E}[y|\text{CMP} = 1, \text{EXC} = 0, \text{MAT} = 0, X] - \mathbb{E}[y|\text{CMP} = 0, \text{EXC} = 0, \text{MAT} = 0, X]). \end{aligned} \quad (3.4)$$

But β_3 does not capture a causal effect because, in contrast to the trade execution requirement that only applies to required transactions in MAT index CDSs, the minimum trading functionality requirement affects all trades in all index CDS contracts (provided that contracts are listed on a SEF which is the case for all non-matured CDX.IG and CDX.HY contracts). Nevertheless, β_3 in itself is an insightful coefficient because it captures potentially differential changes in effective and realized half-spreads of required and permitted transactions over time.

I estimate Equation (3.2) with and without control variables. The control variables that I consider can be grouped into variables that are trade specific and variables that control for market conditions. The trade-specific control variables include a dummy variable for block trades (BLCK) which accounts for the fact that there are no block trades that are subject to the trade execution requirement, and dummy variables for trades in 5Y OFF and Other index CDS contracts (5YOFF and OTHER, respectively). The variables that control for market conditions are the bid-ask spread of the 5Y OTR contract prevailing at trade execution (BAS) and the end-of-day at-the-money implied volatility of a three-month index option on the 5Y OTR contract (VLTLTY).

As can be seen from specifications (1) and (4) of Table 3.2, trading costs of permitted transactions in the 60-calendar-day period prior to the compliance date of SEF rules are statistically indistinguishable from those of required transactions (β_1 is insignificantly different from zero and non-uniformly signed across indices). Trading cost of permitted transactions in the period prior to the compliance date of SEF rules are higher than those in the period between effective dates ($\beta_2 < 0$ albeit insignificantly so for CDX.IG). In comparison to permitted transactions, trading costs of required transactions decline significantly more ($\beta_3 < 0$) over the two periods. For CDX.IG, the magnitude of the decline is 0.080 bps or 28.0% of the level of trading

3.4. Dodd-Frank Regime Trading Costs

	CDX.IG			CDX.HY		
	(1)	(2)	(3)	(4)	(5)	(6)
CNSTNT	0.295** (17.999)	0.193* (2.188)	0.130 (1.509)	1.532** (20.183)	1.283** (3.052)	0.906* (2.202)
MAT	-0.009 (-0.433)	0.090 (1.022)	0.072 (0.831)	0.024 (0.244)	0.247 (0.594)	0.203 (0.495)
CMP	-0.026 (-1.295)	-0.036 (-1.682)	0.029 (1.614)	-0.213* (-2.208)	-0.297** (-3.192)	0.132 (1.511)
CMP×MAT	-0.080** (-3.439)	-0.069** (-3.124)	-0.048* (-2.472)	-0.406** (-3.614)	-0.319** (-2.942)	-0.277** (-2.729)
EXC	-0.055** (-2.662)	-0.047* (-2.263)	-0.011 (-0.506)	-0.309** (-4.423)	-0.249** (-3.973)	-0.121 (-1.921)
EXC×MAT	0.030 (1.912)	0.021 (1.349)	0.022 (1.412)	0.094 (1.408)	0.020 (0.359)	0.008 (0.142)
BLCK		0.081 (0.951)	0.083 (0.975)		0.075 (0.184)	0.060 (0.149)
5YOFF		0.051** (4.076)	0.068** (5.111)		0.383** (4.565)	0.454** (6.676)
OTHER		0.406** (5.504)	0.428** (5.918)		2.282** (5.825)	2.316** (5.952)
BAS			0.967** (7.119)			0.422** (4.390)
VLTLTY			0.314* (2.211)			1.118 (1.040)
<i>N</i>	26,708	26,708	26,708	27,934	27,934	27,934
<i>R</i> ²	0.027	0.042	0.069	0.041	0.071	0.086

Table 3.2: Difference-in-Differences Regression Specifications for Effective Half-Spreads.

The table shows OLS estimates of regression specifications for the difference-in-differences estimator of the causal effect of the minimum trading functionality and trade execution requirements (t -statistics based on standard errors clustered by date are shown in parenthesis). The dependent variable is the effective half-spread defined as $q_t \times (p_t - m_t)$, where p_t is the transaction price and m_t is the mid-quote prevailing at trade execution t . Both transaction prices and quotes are in terms of index CDS spreads and expressed in basis points. Trade direction, q_t , equals +1 (-1) for protection-buyer-initiated (protection-seller-initiated) trades and is inferred by the Lee and Ready (1991) algorithm. The explanatory variables include a dummy variable for non-block trades in made available to trade index CDS contracts (MAT), a dummy variable for trades executed on or after the SEF compliance date (CMP), a dummy variable for trades executed on or after the effective date of the trade execution requirement (EXC), a dummy variable for block trades (BLCK), a dummy variable for trades in the five-year immediate off-the-run contract (5YOFF), a dummy variable for trades in contracts other than the five-year on-the-run and immediate off-the-run contract (OTHER), the bid-ask spread of the five-year on-the-run contract prevailing at trade execution (BAS), and the end-of-day at-the-money implied volatility of a three-month index option on the five-year on-the-run contract (VLTLTY). Continuous explanatory variables are demeaned. N is the number of trades, R^2 is the coefficient of determination, and ** and * denote statistical significance at the 1% and 5% level, respectively. Regression specifications are estimated from all trades between August 3, 2013 and April 26, 2014.

costs of required transactions prior to the compliance date of SEF rules. For CDX.HY, the magnitude of the effect is more substantial in absolute terms, 0.406 bps, but of the same order

Chapter 3. Index CDS Trading Costs around the Introduction of SEFs

of magnitude in relative terms, 26.1%. Moreover, trading costs of permitted transactions in the period between effective dates are significantly higher than those in the 60-calendar-day period after the effective date of the trade execution requirement ($\beta_4 < 0$). However, there is no evidence of a causal effect of the trade execution requirement on trading costs as the declines in trading costs of permitted and required transactions are statistically indistinguishable (β_5 is insignificantly different from zero and incorrectly signed).

Two aspects of the evidence in Table 3.2 deserve more attention. First, the large decline in the cost of required transactions following the SEF compliance date which could, for instance, be due to an increased use of SEFs in anticipation of the trade execution requirement.²¹ Second, the large decline in the cost of permitted transactions following the trade execution requirement which is suggestive of a significant spillover effect of the requirement. This could, for instance, be due to the fact that traders who become SEF participants in order to comply with the trade execution requirement for MAT index CDSs also execute their trades in non-MAT index CDSs on SEFs. However, as observed by Boehmer et al. (2015), any changes in market conditions will confound estimates of the spillover effect and controlling for market conditions becomes important. In line with their observation, trade-specific control variables do not affect spillover effect estimates (see specifications (2) and (5)), while controlling for market conditions leaves spillover effects that are insignificantly different from zero (see specifications (3) and (6)).

However, the decline in the cost of required transactions following the SEF compliance date does not seem to occur in anticipation of the MAT determination because most of the decline accrues prior to the date on which the determination was filed (October 28, 2013). Estimating a variant of regression (3.2) that focuses on the SEF compliance date only (i.e., EXC terms in Equation (3.2) are omitted) gives β_3 estimates of -0.076 bps (t -statistic -2.554), -0.068 bps (t -statistic -2.382), and -0.063 bps (t -statistic -2.203) in case that specifications are estimated from all CDX.IG trades between September 9, 2013 and October 25, 2013 and control variables coincide with those of specifications (1), (2), and (3), respectively. Similarly, estimating this variant of regression (3.2) from all CDX.HY trades during the above period gives β_3 estimates of -0.412 bps (t -statistic -2.147), -0.410 bps (t -statistic -2.067), and -0.309 bps (t -statistic -1.873) in case that control variables coincide with those of specifications (4), (5), and (6), respectively. Moreover, placebo tests that use the date on which the MAT determination was filed instead of the SEF compliance date but that are otherwise identical to the aforementioned variant of regression (3.2) (both in terms of the specification and in terms of the number of trading days in the pre- and post-event windows) give β_3 estimates that are insignificantly different from zero.

Table 3.3 shows the results of the difference-in-differences regression for realized half-spreads. Consistent with the overall decline in profits from liquidity provision in Figure 3.1, specifications (1) and (4) show that realized half-spreads of permitted transactions decrease over

²¹The MAT determination that led to the trade execution requirement was filed by Tradeweb SEF on October 28, 2013 and its certification was announced on January 28, 2014.

3.4. Dodd-Frank Regime Trading Costs

	CDX.IG			CDX.HY		
	(1)	(2)	(3)	(4)	(5)	(6)
CNSTNT	0.216** (12.327)	0.074 (0.857)	0.037 (0.431)	1.260** (16.241)	0.931 (1.715)	0.785 (1.449)
MAT	0.012 (0.550)	0.147 (1.698)	0.136 (1.586)	-0.046 (-0.430)	0.251 (0.465)	0.232 (0.431)
CMP	-0.023 (-1.076)	-0.029 (-1.299)	0.009 (0.423)	-0.254* (-2.553)	-0.322** (-3.174)	-0.161 (-1.429)
CMP×MAT	-0.088** (-3.587)	-0.079** (-3.323)	-0.067** (-3.073)	-0.304* (-2.485)	-0.233 (-1.823)	-0.216 (-1.723)
EXC	-0.052** (-2.934)	-0.049** (-2.866)	-0.027 (-1.556)	-0.387** (-4.946)	-0.344** (-4.848)	-0.280** (-3.894)
EXC×MAT	0.008 (0.477)	0.001 (0.078)	0.002 (0.110)	0.105 (1.373)	0.045 (0.659)	0.042 (0.613)
BLCK		0.117 (1.360)	0.117 (1.370)		0.163 (0.308)	0.159 (0.299)
5YOFF		0.077** (5.981)	0.086** (6.677)		0.464** (5.682)	0.497** (6.544)
OTHER		0.420** (5.368)	0.433** (5.575)		1.984** (4.220)	1.997** (4.252)
BAS			0.628** (3.908)			0.223** (3.139)
VLTLTY			0.094 (0.595)			-0.269 (-0.300)
<i>N</i>	26,708	26,708	26,708	27,934	27,934	27,934
<i>R</i> ²	0.021	0.031	0.038	0.027	0.043	0.045

Table 3.3: Difference-in-Differences Regression Specifications for Realized Half-Spreads.

The table shows OLS estimates of regression specifications for the difference-in-differences estimator of the causal effect of the minimum trading functionality and trade execution requirements (t -statistics based on standard errors clustered by date are shown in parenthesis). The dependent variable is the realized half-spread defined as $q_t \times (p_t - m_{t+\Delta})$, where p_t is the transaction price and $m_{t+\Delta}$ is the first mid-quote that follows trade execution t by at least 15 minutes. Both transaction prices and quotes are in terms of index CDS spreads and expressed in basis points. Trade direction, q_t , equals +1 (-1) for protection-buyer-initiated (protection-seller-initiated) trades and is inferred by the Lee and Ready (1991) algorithm. The explanatory variables include a dummy variable for non-block trades in made available to trade index CDS contracts (MAT), a dummy variable for trades executed on or after the SEF compliance date (CMP), a dummy variable for trades executed on or after the effective date of the trade execution requirement (EXC), a dummy variable for block trades (BLCK), a dummy variable for trades in the five-year immediate off-the-run contract (5YOFF), a dummy variable for trades in contracts other than the five-year on-the-run and immediate off-the-run contract (OTHER), the bid-ask spread of the five-year on-the-run contract prevailing at trade execution (BAS), and the end-of-day at-the-money implied volatility of a three-month index option on the five-year on-the-run contract (VLTLTY). Continuous explanatory variables are demeaned. N is the number of trades, R^2 is the coefficient of determination, and ** and * denote statistical significance at the 1% and 5% level, respectively. Regression specifications are estimated from all trades between August 3, 2013 and April 26, 2014.

time ($\beta_2 < 0$) and significantly more so in case of required transactions ($\beta_3 < 0$). Again, there is evidence of a significant spillover effect in that realized half-spreads of permitted transactions

Chapter 3. Index CDS Trading Costs around the Introduction of SEFs

are significantly lower in the period after the effective date of the trade execution requirement ($\beta_4 < 0$) but there is no evidence of a causal effect of the requirement itself (β_5 insignificantly different from zero and incorrectly signed). Adding control variables affects results of the two indices in a different manner. For CDX.IG, the spillover effect becomes insignificant whereas it remains significant for CDX.HY. In case of both indices and regardless of the control variables used, realized half-spreads of required transactions decline significantly more upon the SEF compliance date than those of permitted transactions although the evidence is relatively weak in case of CDX.HY (one-sided t -tests of the null hypothesis $H_0 : \beta_3 \geq 0$ are nevertheless rejected at the 5% level).

The results for trading costs are consistent with pre-trade transparency having a positive effect on trading costs and, more generally, liquidity. This is, e.g., a theoretical prediction of Pagano and Röell (1996), who study the impact of transparency on liquidity by comparing trading costs in auction and dealer markets under asymmetric information due to a single insider. They find that average (across trade sizes) expected trading costs for uninformed investors in the more transparent auction market are lower than those in the less transparent dealer market even if the insider optimally selects his trading strategy in each of the two markets. However, in the setup of Pagano and Röell (1996) transparency refers to the liquidity provider's ability to infer whether or not order flow is informed, while arguments for SEF rules and the trade execution requirement typically refer to the liquidity demander's access to a menu of executable prices (either in response to a RFQ or available on an order book).

The results for profits from liquidity provision seem to suggest that the decline in trading costs is due to lower profits from liquidity provision, for instance, because of increased competition for liquidity provision on the order books that SEFs are required to operate because of the minimum trading functionality requirement. In light of anecdotal evidence that the order books of SEFs fail to attract liquidity this explanation seems implausible. However, trading on SEFs differs from trading off SEFs along other dimensions of pre-trade transparency such as comparison shopping that increase competition for liquidity provision on SEFs. The next section investigates this hypothesis by comparing trading costs and profits from liquidity provision of on-SEF trades with those of off-SEF trades. Because prior to the SEF compliance date there were no SEFs and because after the effective date of the trade execution requirement on-SEF trade execution is mandatory for required transactions, I focus on the 95-trading-day period between the two effective dates (October 2, 2013 to February 25, 2014) during which on-SEF trade execution of both required and permitted transactions was voluntary.

3.4.2 On-SEF and Off-SEF Trading Cost Comparison

Table 3.4 shows trade-size-weighted average effective and realized half-spreads of trades executed during the above-mentioned period separately for on-SEF and off-SEF trades. The table shows that in case of both indices on-SEF trades have significantly lower trading costs than off-SEF trades. A breakdown into 5Y OTR, 5Y OFF, and Other contracts shows that this is

3.4. Dodd-Frank Regime Trading Costs

Contract	Trades		Effective Half-Spread		Realized Half-Spread	
	On-SEF	Off-SEF	On-SEF	Off-SEF	On-SEF	Off-SEF
Panel A: CDX.IG						
5Y OTR	7,906	5,287	0.153**	0.257	0.054**	0.221
5Y OFF	170	415	0.196*	0.309	0.085**	0.293
Other	33	163	0.722	0.603	0.541	0.571
Total	8,109	5,865	0.158**	0.276	0.058**	0.241
Panel B: CDX.HY						
5Y OTR	7,386	5,659	0.666**	1.373	0.276**	1.222
5Y OFF	294	703	0.943**	1.493	0.713**	1.320
Other	29	160	2.953	3.027	2.765	2.324
Total	7,709	6,522	0.702**	1.453	0.324**	1.277

Table 3.4: Effective and Realized Half-Spreads.

Panels A and B show trade-size-weighted effective and realized half-spreads of on-SEF and off-SEF trades in index CDSs on CDX.IG and CDX.HY, respectively. Averages are separately computed for transactions in five-year on-the-run (5Y OTR) index CDSs, five-year immediate off-the-run (5Y OFF) index CDSs, and transaction in all other (Other) index CDSs. The effective half-spread is defined as $q_t \times (p_t - m_t)$, where p_t is the transaction price and m_t is the mid-quote prevailing at trade execution t . The realized half-spread is defined as $q_t \times (p_t - m_{t+\Delta})$, where $m_{t+\Delta}$ is the first mid-quote that follows trade execution by at least 15 minutes. Both transaction prices and quotes are in terms of index CDS spreads and expressed in basis points. Trade direction, q_t , equals +1 (−1) for protection-buyer-initiated (protection-seller-initiated) trades and is inferred by the Lee and Ready (1991) algorithm. Trades is the number of trades and ** and * denote trade-size-weighted averages of on-SEF trades that significantly differ from those of off-SEF trades at the 1% and 5% level, respectively. The sample period is October 2, 2013 to February 25, 2014.

due to trades in those contracts that become MAT at the end of the period, namely, 5Y OTR and 5Y OFF. For instance, effective half-spreads of 5Y OTR CDX.IG are 0.153 bps on-SEF and 0.257 bps off-SEF or, stated otherwise, on-SEF trading costs for 5Y OTR CDX.IG are 40% lower than off-SEF trading costs. For 5Y OTR CDX.HY, on-SEF trading costs are even more than 50% lower than off-SEF trading costs, with effective half-spreads being 0.666 bps on-SEF and 1.373 bps off-SEF.

Similarly, profits from on-SEF liquidity provision are significantly lower than those from off-SEF liquidity provision in case of both indices and for both 5Y OTR and 5Y OFF contracts. Differences in profits from on-SEF and off-SEF liquidity provision are even more dramatic than those in trading costs of on-SEF and off-SEF trades, with realized half-spreads of on-SEF trades being less than 25% of those of off-SEF trades in case of 5Y OTR contracts on both CDX.IG and CDX.HY. Moreover, profits from on-SEF liquidity provision are lower than those from off-SEF liquidity provision despite price impacts of on-SEF trades being larger than those of off-SEF trades, as can be seen from the difference between effective and realized half-spreads of on-SEF and off-SEF trades. This rules out that larger profits from off-SEF liquidity provision reflect remuneration for adverse selection.

An issue with the above comparison is the fact that trading costs are a likely determinant

of the endogenous choice whether to trade on a SEF or not. Moreover, the above averages may conceal differences in the characteristics of on-SEF and off-SEF trades. In order to control for both, I estimate a latent variable binary choice model in spirit of Bessembinder and Venkataraman (2004) and Hendershott and Madhavan (2015), in which transaction costs, Y^i , are given by

$$Y^i = X' \beta^i + U^i, \quad (3.5)$$

where X is a set of regressors affecting the cost of trade, U^i is a mean-zero error term with variance σ_i^2 , and the subscript $i = 1$ ($i = 0$) denotes on-SEF (off-SEF) trade execution. The observed choice of trade execution, $D(Z)$, with $D(Z) = 1$ in case of on-SEF trade execution and $D(Z) = 0$ in case of off-SEF trade execution, is due to a latent variable $Z'\theta + U^D$ (U^D is a mean-zero error term with variance 1) that is linked to the choice $D(Z)$ in that $D(Z) = 1$ when $Z'\theta + U^D \geq 0$ and $D(Z) = 0$ otherwise. Hendershott and Madhavan (2015) motivate such a model for the choice between electronic and voice-based trade execution of corporate bond transactions in terms of a trade-off between lower search cost (partly due to more fierce competition in case of simultaneously bidding dealers) and higher information leakage of electronic relative to voice-based trade execution. It is the same trade-off that will determine the choice of on-SEF (electronic) and off-SEF (voice-based) trade execution.

Accounting for the fact that Y^1 is only observed when $D(Z) = 1$ and that Y^0 is only observed when $D(Z) = 0$, and assuming joint normality of U s gives the following conditional means

$$\mathbb{E}[Y^1 | D(z) = 1, x, z] = x' \beta^1 + \mathbb{E}[U^1 | D(z) = 1, z] = x' \beta^1 + \rho_1 \sigma_1 \frac{\phi(z'\theta)}{\Phi(z'\theta)}, \quad (3.6)$$

$$\mathbb{E}[Y^0 | D(z) = 0, x, z] = x' \beta^0 + \mathbb{E}[U^0 | D(z) = 0, z] = x' \beta^0 - \rho_0 \sigma_0 \frac{\phi(z'\theta)}{1 - \Phi(z'\theta)}, \quad (3.7)$$

where x and z are realizations of X and Z , respectively, ρ_i , $i = 0, 1$, is the correlation between U^i and U^D , $\Phi(x)$ denotes the cumulative density function of the standard normal distribution, $\phi(x) = \Phi'(x)$, and $\Phi(z'\theta) = \mathbb{P}(D(z) = 1 | z)$ is the probability of on-SEF trade execution conditional on the realization of Z and parameterized by the coefficient vector θ . If traders strategically select lower cost trade executions, then the conditional means in Equations (3.6) and (3.7) should be smaller than the unconditional means $x' \beta^1$ and $x' \beta^0$ in Equation (3.5) or, in other words, $\mathbb{E}[U^1 | D(z) = 1, z] = \rho_1 \sigma_1 \phi(z'\theta) / \Phi(z'\theta) < 0$ and $\mathbb{E}[U^0 | D(z) = 0, z] = -\rho_0 \sigma_0 \phi(z'\theta) / (1 - \Phi(z'\theta)) < 0$.

The trade characteristics that I consider as explanatory variables of the first-stage probit model are dummy variables for trades with trade sizes in the second, third, and fourth quartile of the trade size distribution (SMLL, MDM, and BLCK, respectively),²² dummy variables for trades in 5Y OFF and Other contracts (5YOFF and OTHER, respectively), a dummy variable for trades

²²The 25%, 50%, and 75% quantiles of the trade size distribution for CDX.IG are USD 25MM, USD 50MM, and USD 100MM, respectively, and those of the trade size distribution for CDX.HY are USD 5MM, USD 10MM, and USD 25MM, respectively.

3.4. Dodd-Frank Regime Trading Costs

	CDX.IG	CDX.HY
CNSTNT	0.744** (20.458)	0.420** (11.112)
SMLL	0.121** (3.395)	0.423** (10.744)
MDM	-0.127** (-3.204)	0.056 (1.529)
BLCK	-0.661** (-13.819)	-0.598** (-12.493)
5YOFF	-0.772** (-7.632)	-0.621** (-9.841)
OTHER	-0.960** (-3.792)	-0.761** (-3.376)
RFRNC	-0.247** (-4.803)	-0.157** (-4.366)
UNCLRD	-1.806** (-29.219)	-1.639** (-23.366)
BAS	-1.888** (-4.911)	-0.077* (-1.972)
VLTLTY	2.798** (5.325)	2.094** (4.248)
<i>N</i>	13,974	14,231

Table 3.5: Probit Regressions for the Choice of On-SEF and Off-SEF Trade Execution.

The table shows coefficient estimates of probit regression specifications for the binary choice between on-SEF and off-SEF trade execution (*t*-statistics based on standard errors clustered by date are shown in parenthesis). The dependent variable equals one for trades that are executed on a SEF. The explanatory variables include dummy variables for small-sized trades (SMLL; USD 25–50MM trade size for CDX.IG and USD 5–10MM trade size for CDX.HY), for medium-sized trades (MDM; USD 50–100MM trade size for CDX.IG and USD 10–25MM trade size for CDX.HY), for block-sized trades (BLCK; trade size > USD 100MM for CDX.IG and trade size > USD 25MM for CDX.HY), for trades in the five-year immediate off-the-run contract (5YOFF), for trades in contracts other than the five-year on-the-run or immediate off-the-run contract (OTHER), for trades with reference-level transaction prices (RFRNC), and for trades that are not centrally cleared (UNCLRD), the bid-ask spread of the five-year on-the-run contract prevailing at trade execution (BAS), and the end-of-day at-the-money implied volatility of a three-month index option on the five-year on-the-run contract (VLTLTY). Continuous explanatory variables are demeaned. *N* is the number of trades and ** and * denote statistical significance at the 1% and 5% level, respectively. Probit regression specifications are estimated from all trades between October 2, 2013 to February 25, 2014.

with transaction prices that coincide with a reference level of an index option or tranche swap quote from the same trading day (RFRNC),²³ and a dummy variable for trades that are not centrally cleared (UNCLRD). 5YOFF and OTHER capture differences in the unconditional likelihood of on-SEF trade execution between trades in the 5Y OTR contract and trades in 5Y

²³Both index options and tranche swaps are conventionally traded “with delta.” That is, along with the index option or tranche swap a delta neutralizing notional amount in the corresponding index CDS is traded in the opposite direction. These trades are executed as packages and the index CDS leg is distinguished by the fact that it is executed at a reference level that does not necessarily have to reflect the current index level because it tends to be set at the beginning of the trading day.

Chapter 3. Index CDS Trading Costs around the Introduction of SEFs

OFF and Other contracts (see Table 3.4). RFRNC and UNCLRD account for the fact that only few SEFs facilitate index option and tranche swap packages and uncleared trades. In addition to trade characteristics, I also consider the bid-ask spread of the 5Y OTR contract that prevails at trade execution (BAS) and the end-of-day at-the-money implied volatility of three-month index options written on the 5Y OTR contract (VLTLTY). The latter explanatory variables are demeaned for the ease of interpretation and capture market liquidity and volatility at trade execution.

Table 3.5 shows coefficient estimates of first-stage probit models. In case of both indices the likelihood of on-SEF trade execution decreases with trade size.²⁴ Trades in 5Y OFF and Other contracts are less likely to be executed on SEFs than trades in the 5Y OTR contract. Similarly, trades that are potentially part of index option or tranche swap packages and uncleared trades are less likely to be executed on SEFs than non-packaged cleared trades. Surprisingly, market liquidity and volatility have opposite effects on the likelihood of on-SEF trade execution. When liquidity deteriorates the likelihood of on-SEF trade execution decreases, while the likelihood of on-SEF trade execution increases when volatility is high. The latter is consistent with fast trade execution being important when volatility is high, whereas the former seems to suggest that there are benefits (such as less information leakage) associated with bilateral negotiations when liquidity is low.

Table 3.6 shows second-stage coefficient estimates of the conditional mean specifications in Equations (3.6) and (3.7), respectively. The explanatory variables of unconditional mean trading costs include trade size dummies (SMLL, MDM, and BLCK), contract dummies (5YOFF and OTHER), the reference level dummy (RFRNC), and the bid-ask spread and implied volatility of the 5Y OTR contract (BAS and VLTLTY). Trade size dummies capture sensitivity of trading costs to trade size (both asymmetric information and inventory considerations suggest that transaction costs increase with trade size) and contract dummies account for the fact the 5Y OFF and Other contracts tend to have higher trading costs than the 5Y OTR contract (see Table 3.4). RFRNC accounts for reference levels that do not necessarily reflect current index levels, potentially, giving rise to higher trading costs. Finally, BAS and VLTLTY capture sensitivity of trading costs to market liquidity and volatility.²⁵

Conditional means are estimated by separate regressions for on-SEF and off-SEF trades which in addition to the above explanatory variables include the respective inverse Mill's ratio (MLLSRT), i.e., $\phi(z'\theta)/\Phi(z'\theta)$ in Equation (3.6) and $\phi(z'\theta)/(1 - \Phi(z'\theta))$ in Equation (3.7). Thus, strategic selection of lower cost trade execution is reflected by negative coefficient estimates on MLLSRT. Inference is based on Heckman et al. (2003) and, by casting the two-step procedure into a generalized method of moments (GMM) framework, takes into account that inverse Mill's ratios are generated regressors. Moreover, I allow for a general correlation structure

²⁴The exception being trades in the first quartile of the trade size distribution for which on-SEF trade execution tends to be less likely than for trades in the second quartile of the trade size distribution.

²⁵As before, BAS and VLTLTY are demeaned for the ease of interpretation. Because conditional means are estimated in separate regressions for on-SEF and off-SEF trades (see next paragraph), BAS and VLTLTY are separately demeaned for on-SEF and off-SEF trades.

3.4. Dodd-Frank Regime Trading Costs

	CDX.IG		CDX.HY	
	On-SEF	Off-SEF	On-SEF	Off-SEF
CNSTNT	0.100** (8.515)	0.233** (18.246)	0.623** (9.944)	1.479** (10.462)
SMLL	0.006 (1.101)	-0.020 (-1.537)	-0.089* (-2.171)	-0.147* (-2.187)
MDM	0.026** (3.255)	-0.010 (-0.604)	-0.049 (-1.449)	-0.104 (-1.211)
BLCK	0.054** (4.977)	0.028 (1.588)	-0.001 (-0.021)	-0.029 (-0.334)
5YOFF	0.018 (1.103)	0.068* (2.258)	0.361** (5.231)	0.265** (2.674)
OTHER	0.649** (3.908)	0.322** (7.172)	2.037** (5.469)	2.034** (7.474)
RFRNC	0.084** (5.162)	0.065** (3.916)	0.096 (1.799)	-0.015 (-0.184)
BAS	0.894** (5.033)	0.785** (4.190)	0.261** (7.218)	0.419** (3.596)
VLTLY	0.411** (3.518)	0.600** (3.229)	1.635** (3.818)	4.363** (4.216)
MLLSRT	0.044* (2.013)	0.009 (0.620)	0.090 (1.037)	-0.160 (-1.633)
<i>N</i>	8,109	5,865	7,709	6,522
<i>R</i> ²	0.127	0.044	0.068	0.055

Table 3.6: Effective Half-Spreads in Choice Model for On-SEF and Off-SEF Trade Execution.

The table shows OLS estimates of linear specifications in a latent variable binary choice model (t -statistics based on standard errors clustered by date are shown in parenthesis; standard error computation follows Heckman, Tobias, and Vytalil (2003)). The dependent variable is the effective half-spread defined as $q_t \times (p_t - m_t)$, where p_t is the transaction price and m_t is the mid-quote prevailing at trade execution t . Both transaction prices and quotes are in terms of index CDS spreads and expressed in basis points. Trade direction, q_t , equals +1 (-1) for protection-buyer-initiated (protection-seller-initiated) trades and is inferred by the Lee and Ready (1991) algorithm. The explanatory variables include dummy variables for small-sized trades (SMLL; USD 25–50MM trade size for CDX.IG and USD 5–10MM trade size for CDX.HY), for medium-sized trades (MDM; USD 50–100MM trade size for CDX.IG and USD 10–25MM trade size for CDX.HY), for block-sized trades (BLCK; trade size > USD 100MM for CDX.IG and trade size > USD 25MM for CDX.HY), for trades in the five-year immediate off-the-run contract (5YOFF), for trades in contracts other than the five-year on-the-run or immediate off-the-run contract (OTHER), and for trades with reference-level transaction prices (RFRNC), the bid-ask spread of the five-year on-the-run contract prevailing at trade execution (BAS), the end-of-day at-the-money implied volatility of a three-month index option on the five-year on-the-run contract (VLTLY), and inverse Mill's ratios based on the choice model estimated in Table 3.5 (MLLSRT; $\phi(z'\theta)/\Phi(z'\theta)$ for on-SEF trades and $\phi(z'\theta)/(1 - \Phi(z'\theta))$ for off-SEF trades). Continuous explanatory variables other than MLLSRT are demeaned. N is the number of trades, R^2 is the coefficient of determination, and ** and * denote statistical significance at the 1% and 5% level, respectively. Regression specifications are estimated from the indicated number of on-SEF and off-SEF trades between October 2, 2013 to February 25, 2014.

among error terms of trades that are executed on the same trading day by using cluster-robust inference for GMM estimators (see, e.g., Cameron, Gelbach, and Miller (2011)).

Chapter 3. Index CDS Trading Costs around the Introduction of SEFs

The results show no evidence for strategic selection. This most likely reflects factors that the model is unable to capture such as costly onboarding and compliance processes associated with joining a SEF. In fact, many market participants did not embrace the CFTC's trade execution requirement because of the associated costs and up to date there seems to be a perception among some market participants to avoid on-SEF trade execution whenever possible. On-SEF trading costs of CDX.IG increase with trade size whereas off-SEF trading costs are insensitive to trade size. For CDX.HY, trades in the second quartile of the trade size distribution have significantly lower trading costs than those in the other quartiles. Consistent with Table 3.4, on-SEF and off-SEF trades in 5Y OFF and Other contracts tend to have higher trading costs than those in the 5Y OTR contract. Moreover, trading costs are high when bid-ask spreads are wide and when volatility is high.

As an illustration, Panels A and B of Figure 3.2 show unconditional mean trading costs implied by the latent variable binary choice model, i.e., $x' \beta^i$, $i = 0, 1$, for 5Y OTR CDX.IG and CDX.HY, respectively, by quartiles of the trade size distribution.²⁶ The panels show that trading costs of on-SEF trades in 5Y OTR contracts are significantly lower than those of off-SEF trades regardless of the trade size and after controlling for the endogenous choice of whether to trade on SEFs or not.

Because there is no evidence for strategic selection and because profits from liquidity provision are a less likely determinant of the liquidity demander's choice whether to execute a trade on a SEF or not, I resort to a simpler regression specification when controlling for the fact that average realized half-spreads reported in Table 3.4 may conceal differences in trade characteristics. As before, the explanatory variables include trade size dummies (SMALL, MED, and BLACK), contract dummies (5YOFF and OTHER), the reference level dummy (RFRNC), and the bid-ask spread and implied volatility of the 5Y OTR contract (BAS and VLTLTY). In spirit of the above, I estimate realized half-spread regression specifications separately for on-SEF and off-SEF trades and I use cluster-robust inference.²⁷

Table 3.7 shows the results. Realized half-spreads of on-SEF and off-SEF trades are not systematically related to trade size. Profits from liquidity provision on trades in 5Y OFF and Other contracts seem to be significantly higher than those on trades in 5Y OTR contracts. For CDX.IG, realized half-spreads tend to be high when market liquidity is low and realized half-spreads of off-SEF trades in both indices are sensitive to market volatility. As an illustration, Panels C and D of Figure 3.2 show regression-specification-implied profits from on-SEF and off-SEF liquidity provision in 5Y OTR contracts (as a function of trade size), confirming that even after accounting for trade characteristics, profits from off-SEF liquidity provision are significantly higher than those from on-SEF liquidity provision irrespective of trade size. Similar results are obtained if the regression in addition includes inverse Mill's ratios based on the above first-stage probit model.

²⁶More precisely, the panels show the unconditional mean trading costs of 5Y OTR trades with non-reference level transaction prices and BAS and VLTLTY evaluated at their unconditional means of 0.

²⁷As before, BAS and VLTLTY are separately demeaned for on-SEF and off-SEF trades for ease of interpretation.

3.4. Dodd-Frank Regime Trading Costs

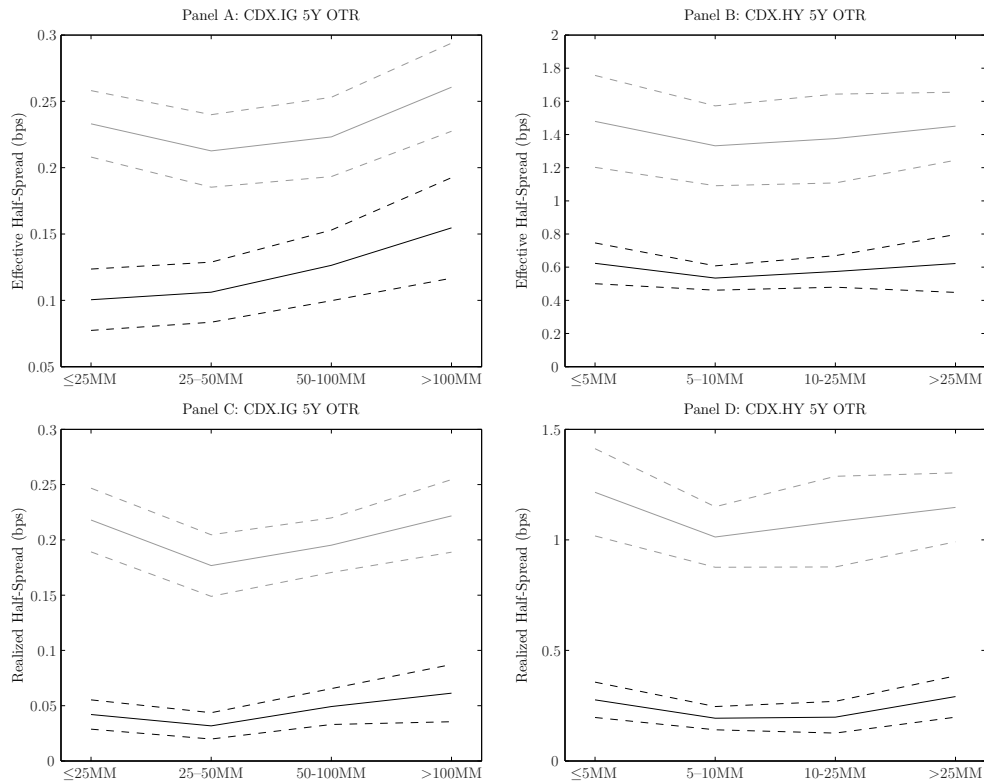


Figure 3.2: Effective and Realized Half-Spreads by Quartiles of the Trade Size Distribution. Panels A and B show effective half-spreads by quartiles of the trade size distribution for on-SEF (black lines) and off-SEF (gray lines) trades in five-year on-the-run (5Y OTR) CDX.IG and CDX.HY, respectively. Effective half-spread estimates are based on the linear specifications in the latent variable binary choice model estimated in Tables 3.5 and 3.6. Panels C and D show realized half-spreads by quartiles of the trade size distribution for on-SEF (black lines) and off-SEF (gray lines) trades in 5Y OTR CDX.IG and CDX.HY, respectively. Realized half-spread estimates are based on the regression specifications estimated in Table 3.7. Dashed lines mark 95% confidence intervals. The sample period is October 2, 2013 to February 25, 2014.

The evidence thus far shows that trades executed on SEFs have lower transaction costs and are less profitable from a liquidity provider's perspective. The CFTC's trade execution requirement that came into effect on February 26, 2014 mandates on-SEF trade execution for a significant share of trades. Thus, one reason for the decline in trading costs and profits from liquidity provision exhibited in Figure 3.1 is the higher share of trades executed on SEFs. As argued above, pre-trade price competition on SEFs is higher than in bilateral negotiations which is a likely explanation for the lower profits from on-SEF liquidity provision. In order to provide some evidence in support of stronger price competition on SEFs, I next compare the fractions of on-SEF and off-SEF trades with transaction prices outside the quoted bid-ask spread that prevails at trade execution. For comparability with the results of this section, I again focus on the period between the effective dates of the minimum trading functionality requirement and the trade execution requirement.

Chapter 3. Index CDS Trading Costs around the Introduction of SEFs

	CDX.IG		CDX.HY	
	On-SEF	Off-SEF	On-SEF	Off-SEF
CNSTNT	0.042** (6.205)	0.218** (14.822)	0.277** (6.785)	1.215** (12.065)
SMLL	-0.010 (-1.330)	-0.041** (-2.697)	-0.083 (-1.869)	-0.202* (-2.208)
MDM	0.007 (0.713)	-0.023 (-1.386)	-0.079 (-1.646)	-0.132 (-1.299)
BLCK	0.019 (1.469)	0.004 (0.224)	0.015 (0.252)	-0.068 (-0.687)
5YOFF	0.046* (2.115)	0.079* (2.277)	0.479** (4.254)	0.272* (2.479)
OTHER	0.510** (7.080)	0.278** (4.204)	1.906 (1.850)	1.532** (3.910)
RFRNC	0.112** (5.970)	0.096** (4.213)	0.230** (3.259)	0.115 (1.125)
BAS	0.556** (3.275)	0.586** (2.581)	0.058 (0.949)	0.193 (1.552)
VLTLTY	-0.001 (-0.007)	0.557* (2.453)	0.109 (0.189)	3.575** (3.145)
<i>N</i>	8,109	5,865	7,709	6,522
<i>R</i> ²	0.039	0.029	0.019	0.021

Table 3.7: Regression Specifications for Realized Half-Spreads of On-SEF and Off-SEF Trades. The table shows OLS estimates of regression specifications for realized half-spreads of on-SEF and off-SEF trades (*t*-statistics based on standard errors clustered by date are shown in parenthesis). The dependent variable is the realized half-spread defined as $q_t \times (p_t - m_{t+\Delta})$, where p_t is the transaction price and $m_{t+\Delta}$ is the first mid-quote that follows trade execution t by at least 15 minutes. Both transaction prices and quotes are in terms of index CDS spreads and expressed in basis points. Trade direction, q_t , equals +1 (−1) for protection-buyer-initiated (protection-seller-initiated) trades and is inferred by the Lee and Ready (1991) algorithm. The explanatory variables include dummy variables for small-sized trades (SMLL; USD 25–50MM trade size for CDX.IG and USD 5–10MM trade size for CDX.HY), for medium-sized trades (MDM; USD 50–100MM trade size for CDX.IG and USD 10–25MM trade size for CDX.HY), for block-sized trades (BLCK; trade size > USD 100MM for CDX.IG and trade size > USD 25MM for CDX.HY), for trades in the five-year immediate off-the-run contract (5YOFF), for trades in contracts other than the five-year on-the-run or immediate off-the-run contract (OTHER), and for trades with reference-level transaction prices (RFRNC), the bid-ask spread of the five-year on-the-run contract prevailing at trade execution (BAS), and the end-of-day at-the-money implied volatility of a three-month index option on the five-year on-the-run contract (VLTLTY). Continuous explanatory variables are demeaned. *N* is the number of trades, *R*² is the coefficient of determination, and ** and * denote statistical significance at the 1% and 5% level, respectively. Regression specifications are estimated from the indicated number of on-SEF and off-SEF trades between October 2, 2013 to February 25, 2014.

3.4.3 Trades at Prices Outside the Quoted Bid-Ask Spread

Markit bid and ask quotes are dealer composites and, therefore, a transaction price outside the quoted bid-ask spread does not necessarily reflect trade execution outside what was quoted by the dealer that facilitated the trade. However, a transaction price outside Markit's quoted

3.4. Dodd-Frank Regime Trading Costs

Contract	Trades		Outside Quoted Spread		Outside Quoted Spread (%)	
	On-SEF	Off-SEF	On-SEF	Off-SEF	On-SEF	Off-SEF
Panel A: CDX.IG						
5Y OTR	7,906	5,287	941	1,514	11.9**	28.6
5Y OFF	170	415	31	132	18.2**	31.8
Other	33	163	18	52	54.5	31.9
Total	8,109	5,865	990	1,698	12.2**	29.0
Panel B: CDX.HY						
5Y OTR	7,386	5,659	594	1,635	8.0**	28.9
5Y OFF	294	703	33	199	11.2**	28.3
Other	29	160	7	57	24.1	35.6
Total	7,709	6,522	634	1,891	8.2**	29.0

Table 3.8: Trades at Prices Outside the Quoted Bid-Ask Spread.

Panels A and B show the number and percentage of on-SEF and off-SEF trades with transaction prices outside the quoted bid-ask spread for CDX.IG and CDX.HY, respectively. Trades is the number of trades and ** and * denote fractions of on-SEF trades with transaction prices outside the quoted bid-ask spread that significantly differ from those of off-SEF trades at the 1% and 5% level, respectively. The sample period is October 2, 2013 to February 25, 2014.

composite bid-ask spread constitutes a valid metric for a comparison of the competitiveness of on-SEF and off-SEF liquidity provision: if order flow with the same characteristics would be executed on and off SEFs and pre-trade price competition would be identical on SEFs and in bilateral negotiations, then one would not expect to see differences in the fractions of on-SEF and off-SEF trades with transaction prices outside the quoted bid-ask spread. At this point, it should also be noted that looking at trades with transaction prices strictly inside the quoted spread (i.e., at trades that look as if they have received price improvement from the quoted bid or ask spread) gives opposite but otherwise almost identical results because only few trades are executed at the quoted composite bid or ask spread.^{28,29}

Table 3.8 shows the number and fraction of trades with transaction prices outside the quoted bid-ask spread. The fraction of trades with transaction prices outside the quoted bid-ask spread is significantly higher for off-SEF trades in case of both CDX.IG and CDX.HY. But there are some differences among contracts. For the contracts that are relatively actively quoted, 5Y OTR and 5Y OFF, the fraction of trades with transaction prices outside the quoted bid-ask spread is significantly lower for on-SEF trades, while there is no statistically discernable difference for Other contracts. For instance, 11.9% of on-SEF trades in 5Y OTR CDX.IG have transaction prices outside the quoted bid-ask spread and so do 8.0% of on-SEF trades in 5Y OTR CDX.HY. In comparison, 28.6% and 28.9% of off-SEF trades in these two contracts have

²⁸Again, due to the fact that quotes are dealer composites, a trade strictly inside the quoted spread does not necessarily reflect trade execution inside what was quoted by the dealer that facilitated the trade.

²⁹For CDX.IG, 82 trades are executed at either the quoted bid spread or the quoted ask spread. For CDX.HY, 170 trades are executed at either the quoted bid price or the quoted ask price (the comparison for CDX.HY is based on prices instead of spreads because the index is quoted in terms of a price and converted spreads mechanically mismatch quoted composite spreads because the latter are rounded).

Chapter 3. Index CDS Trading Costs around the Introduction of SEFs

transaction prices outside the quoted bid-ask spread. As discussed above, on-SEF and off-SEF trades differ in the degree of pre-trade price competition for liquidity provision. Quotes on SEFs are executable and usually come from multiple liquidity providers that simultaneously compete for a trade. In contrast, when bilaterally negotiating trades, quotes are executable only “as long as the breath is warm” and, as a consequence, quotes collected from multiple liquidity providers are subject to strategic price deterioration upon a repeat contact for trade execution. Thus, SEFs allow for better comparison shopping which should ultimately lead to increased competition among liquidity providers. For 5Y OTR and 5Y OFF contracts, the observed differences in the fractions of on-SEF and off-SEF trades with transaction prices outside the quoted bid-ask spread are consistent with stronger competition for liquidity provision of trades that are executed on SEFs in comparison to trades that are bilaterally negotiated off SEFs.

Off-SEF, around 30% of trades have transaction prices outside the quoted bid-ask spread irrespective of the contract.³⁰ In contrast, on-SEF, the fractions of actively quoted contracts (5Y OTR and 5Y OFF) are lower than those of other contracts (Other). One reason that cross-contract differences prevail for on-SEF trades but not for off-SEF trades may be the number of dealers queried for on-SEF trade execution via RFQ.³¹ Because the cost of information leakage is higher for less actively traded contracts, it is plausible that traders query less dealers when they want to trade an Other contract in comparison to the case when they want to trade 5Y OTR or 5Y OFF contracts. But querying more dealers increases pre-trade price competition which leads to better transaction prices and, consequently, less instances where the transaction price is outside the quote bid-ask spread. For on-SEF trades, the cross-contract pattern of the fraction of trades with transaction prices outside the quoted bid-ask spread is consistent with this explanation. Moreover, the absence of a cross-contract pattern for off-SEF trades is consistent with the explanation as well because there is no simultaneous pre-trade price competition when trades are bilaterally negotiated off SEFs.

In order to control for differences in the characteristics of on-SEF and off-SEF trades and the market conditions during which the trades are executed, I estimate trade-by-trade probit regressions where the dependent variable equals one for trades with transaction prices outside the quoted bid-ask spread. The explanatory variables include a dummy variable for trades that are executed on SEFs (SEF) and some of the above control variables. Specifically, I include trade size dummies (SMALL, MDM, and BLCK) in order to control for the fact that quotes reflect prices near which only an instrument’s standard notional amount can be expected to get executed without additional bargaining. I include contract dummies (5YOFF and OTHER) in order to control for differences in the unconditional likelihood with which trades in 5Y OTR, 5Y OFF, and Other contracts are executed at prices outside the quoted bid-ask spread

³⁰ *t*-tests for the null hypothesis that the fraction of off-SEF trades with transaction prices outside the quoted bid-ask spread is the same for 5Y OTR (5Y OFF) and Other contracts fail to reject the null hypothesis at conventional significance levels in case of both indices.

³¹ As mentioned in Sections 3.3, on-SEF trades tend to be executed via RFQ. Also note that during the period under consideration there was no requirement to transmit requests to a minimum number of dealers.

(see Table 3.8). I also include the reference level dummy (RFRNC) in order to control for the fact that reference levels do not necessarily reflect current index levels which increases the likelihood of trade execution at a price outside the quoted bid-ask spread. Finally, I include the bid-ask spread and the implied volatility of the 5Y OTR contract (BAS and VLTLTY) in order to control for liquidity and volatility that prevails in the market at trade execution.³²

Table 3.9 shows coefficient estimates of the probit regressions. The most important result is the strongly significant and negative estimate of the coefficient on the SEF dummy which shows that on-SEF trades are significantly less likely to be executed at prices outside the quoted bid-ask spread. This is strong evidence in support of relatively more pre-trade price competition on SEFs. Moreover, regression results reveal that larger-sized trades are significantly more likely to be executed at prices outside the quoted bid-ask spread than smaller-sized trades.³³ Consistent with quotes being for an instrument's standard notional amount, at- or below-median-sized trades in the second quartile of the trade size distribution are most likely to be executed within the quoted bid-ask spread. In line with Table 3.8, 5Y OFF and Other contracts are more likely to be executed at prices outside the quoted bid-ask spread but in terms of statistical significance results are non-uniform across indices. Finally, trades are more likely to get executed at prices outside the quoted bid-ask spread when liquidity is low and volatility is high.

For 5Y OTR CDX.IG, probit regression estimates imply that 90.7% of on-SEF trades with trade size in the second quartile of the trade size distribution get executed at or within Markit's quoted composite bid-ask spread. In comparison, the estimates imply that 77.8% of off-SEF trades in this contract get executed at or within the quoted spread. This suggests that for bilaterally negotiated off-SEF trades the probability of trade execution at a price outside the quoted bid-ask spread is more than twice that of trades executed on SEFs. For 5Y OTR CDX.HY, the discrepancy is even larger because regression estimates imply that 93.5% of on-SEF trades with trade size in the second quartile of the trade size distribution get executed at or within Markit's quoted composite bid-ask spread while only 76.1% of off-SEF trades do. Overall, this section provides strong evidence in support of higher pre-trade price competition for trades that are executed on SEFs in comparison to bilaterally negotiated trades off SEFs. This supports increased pre-trade price competition as an explanation for the lower trading costs of on-SEF trades and the smaller profits from on-SEF liquidity provision.

3.4.4 Robustness

While allowing for the endogenous choice of whether to trade on a SEF or not, the above comparison of on-SEF and off-SEF trading costs ignores market-structure- and regulation-

³²As before, BAS and VLTLTY are demeaned for ease of interpretation and cluster-robust inference allows for a general correlation structure among error terms of trades that are executed on the same trading day.

³³The probability of trade execution at prices outside the quoted bid-ask spread is an increasing function of trade size only in the region beyond the 25% quantile of the trade size distribution. In contrast, the probability decreases when trade size increases from the first to the second quartile of the trade size distribution.

Chapter 3. Index CDS Trading Costs around the Introduction of SEFs

	CDX.IG	CDX.HY
CNSTNT	-0.744** (-17.643)	-0.521** (-12.569)
SMLL	-0.023 (-0.672)	-0.190** (-3.963)
MDM	0.091* (2.273)	-0.106* (-2.472)
BLCK	0.283** (6.167)	0.037 (0.778)
5YOFF	0.203** (3.067)	0.013 (0.239)
OTHER	0.158 (0.675)	0.224* (1.980)
RFRNC	0.456** (7.834)	0.085 (1.574)
SEF	-0.555** (-13.369)	-0.805** (-20.905)
BAS	2.247** (4.628)	0.058 (1.033)
VLTLTY	3.005** (5.512)	2.977** (4.389)
<i>N</i>	13,974	14,231

Table 3.9: Probit Regressions for Trades at Prices Outside the Quoted Bid-Ask Spread.

The table shows coefficient estimates of probit regression specifications for trades with transaction prices outside the quoted bid-ask spread (t -statistics based on standard errors clustered by date are shown in parenthesis). The dependent variable equals one for trades that are executed at a prices outside the quoted bid-ask spread that prevails at trade execution. The explanatory variables include dummy variables for small-sized trades (SMLL; USD 25–50MM trade size for CDX.IG and USD 5–10MM trade size for CDX.HY), for medium-sized trades (MDM; USD 50–100MM trade size for CDX.IG and USD 10–25MM trade size for CDX.HY), for block-sized trades (BLCK; trade size > USD 100MM for CDX.IG and trade size > USD 25MM for CDX.HY), for trades in the five-year immediate off-the-run contract (5YOFF), for trades in contracts other than the five-year on-the-run or immediate off-the-run contract (OTHER), for trades with reference-level transaction prices (RFRNC), and for trades that are executed on a SEF (SEF), the bid-ask spread of the five-year on-the-run contract prevailing at trade execution (BAS), and the end-of-day at-the-money implied volatility of a three-month index option on the five-year on-the-run contract (VLTLTY). Continuous explanatory variables are demeaned. N is the number of trades and ** and * denote statistical significance at the 1% and 5% level, respectively. Probit regression specifications are estimated from all trades between October 2, 2013 to February 25, 2014.

implied specifics of on-SEF trading. These arise from the bifurcated structure of the index CDS market into dealer-operated client markets and broker-operated interdealer markets. Credit derivatives dealers trade with their institutional clients in the former and manage their inventories in the latter. Because interdealer brokerage falls under the activities specified in the definition of a SEF, interdealer brokers (IDBs) have to comply with SEF rules and, as a consequence, the interdealer market migrated on IDB SEFs when SEF rules became effective on October 2, 2013. On SEFs, client (dealer-to-customer) trades have higher trading costs than interdealer (dealer-to-dealer) trades because the latter serve to manage inventories

3.4. Dodd-Frank Regime Trading Costs

Contract	Trades		Effective Half-Spread		Realized Half-Spread	
	On-SEF	Off-SEF	On-SEF	Off-SEF	On-SEF	Off-SEF
Panel A: CDX.IG						
5Y OTR	4,943	5,287	0.156**	0.257	0.043**	0.221
5Y OFF	90	415	0.225	0.309	0.090**	0.293
Total	5,033	5,702	0.158**	0.261	0.044**	0.226
Panel B: CDX.HY						
5Y OTR	4,808	5,659	0.676**	1.373	0.179**	1.222
5Y OFF	165	703	1.221	1.493	0.702**	1.320
Total	4,973	6,362	0.698**	1.387	0.201**	1.233

Table 3.10: Effective and Realized Half-Spreads when Excluding Interdealer Trades.

Panels A and B show trade-size-weighted effective and realized half-spreads of on-SEF and off-SEF trades in index CDSs on CDX.IG and CDX.HY, respectively. On-SEF trades are limited to dealer-to-customer trades occurring on non-interdealer-broker SEFs. Averages are separately computed for transactions in five-year on-the-run (5Y OTR) index CDSs and five-year immediate off-the-run (5Y OFF) index CDSs. The effective half-spread is defined as $q_t \times (p_t - m_t)$, where p_t is the transaction price and m_t is the mid-quote prevailing at trade execution t . The realized half-spread is defined as $q_t \times (p_t - m_{t+\Delta})$, where $m_{t+\Delta}$ is the first mid-quote that follows trade execution by at least 15 minutes. Both transaction prices and quotes are in terms of index CDS spreads and expressed in basis points. Trade direction, q_t , equals +1 (-1) for protection-buyer-initiated (protection-seller-initiated) trades and is inferred by the Lee and Ready (1991) algorithm. Trades is the number of trades and ** and * denote trade-size-weighted averages of on-SEF trades that significantly differ from those of off-SEF trades at the 1% and 5% level, respectively. The sample period is October 2, 2013 to February 25, 2014.

(see, e.g., Collin-Dufresne et al. (2016)). Thus, in order to rule out that the on-SEF and off-SEF trading cost comparison only reflects differences in trading costs of client (off-SEF) and interdealer (on-SEF) trades, I remove all on-SEF interdealer trades from the sample.³⁴ As in Collin-Dufresne et al. (2016), on-SEF interdealer trades are identified as trades that are executed on an IDB SEF.

Table 3.10 shows trade-size-weighted average effective and realized half-spreads of the remaining trades. In fact, the table shows trade-size-weighted averages for 5Y OTR and 5Y OFF contracts only. This is because, after removing on-SEF interdealer trades, so few (a total of ten in both indices) on-SEF trades in Other contracts remain that trading costs cannot be estimated accurately. The small number of on-SEF client trades in Other contracts likely reflects both high leakage costs associated with requesting quotes on an inactively traded contract and low response rates. Moreover, it suggests that IDB intermediation mitigates such obstacles in the interdealer market. Removing on-SEF interdealer trades has a minimal effect on on-SEF trading costs of 5Y OTR contracts but, in comparison to Table 3.4, increases on-SEF trading costs of 5Y OFF contracts to a degree that they become statistically indistinguishable from off-SEF trading costs. Profits from on-SEF liquidity provision tend to decrease further,

³⁴This assumes that all off-SEF trades are client trades. The assumption is not implausible because interdealer trade typically involves some sort of IDB service (either an IDB-operated order book or voice-brokerage). Moreover, low trading costs of non-removed off-SEF interdealer trades bias the comparison against finding larger off-SEF trading costs.

Chapter 3. Index CDS Trading Costs around the Introduction of SEFs

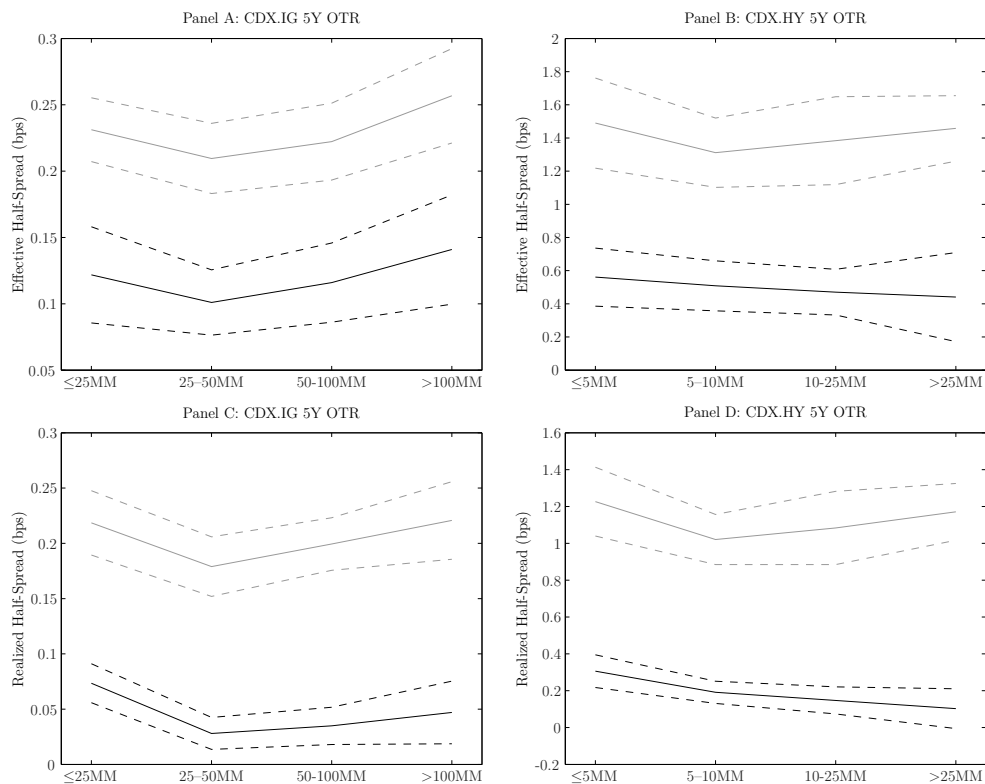


Figure 3.3: Effective and Realized Half-Spreads by Quartiles of the Trade Size Distribution when Excluding Interdealer Trades.

Panels A and B show effective half-spreads by quartiles of the trade size distribution for on-SEF (black lines) and off-SEF (gray lines) trades in five-year on-the-run (5Y OTR) CDX.IG and CDX.HY, respectively. Effective half-spread estimates are based on linear specifications in a latent variable binary choice model that excludes the dummy variable for trades in contracts other than the five-year on-the-run and immediate off-the-run contract (OTHER) but is otherwise identical to the latent variable binary choice model estimated in Tables 3.5 and 3.6. Panels C and D show realized half-spreads by quartiles of the trade size distribution for on-SEF (black lines) and off-SEF (gray lines) trades in 5Y OTR CDX.IG and CDX.HY, respectively. Realized half-spread estimates are based on regression specifications that exclude the dummy variable for trades in contracts other than the five-year on-the-run and immediate off-the-run contract (OTHER) but are otherwise identical to the regression specifications estimated in Table 3.7. The samples from which the latent variable binary choice models and regression specifications are estimated exclude on-SEF dealer-to-dealer trades occurring on interdealer-broker SEFs and all trades in contracts other than the five-year on-the-run and immediate off-the-run contract. Dashed lines mark 95% confidence intervals. The sample period is October 2, 2013 to February 25, 2014.

reinforcing earlier results of significantly lower profits from on-SEF liquidity provision.

Estimating choice models from the samples that exclude on-SEF interdealer trades and all trades in Other contracts gives results that are consistent with those reported in Tables 3.5 and 3.6. Panels A and B of Figure 3.3 show unconditional mean trading costs implied by the latent variable binary choice models estimated from these samples. The panels show

that differences in trading costs of on-SEF and off-SEF trades in 5Y OTR contracts are not due to lower trading costs of on-SEF interdealer trades. Moreover, estimating regression specifications alike the ones in Table 3.7 gives estimates of profits from off-SEF liquidity provision that are significantly higher than those from on-SEF liquidity provision as can be seen from Panels C and D of Figure 3.3. Thus, potentially low profits on on-SEF interdealer trades are not the reason for the difference in profits from on-SEF and off-SEF liquidity provision.

3.5 Conclusion

I document a reduction of trading costs in the index CDS market over the course of a two-and-a-half-year period during which the CFTC implemented Dodd-Frank Act provisions. I provide evidence in support of lower profits from liquidity provision driving the decline in the cost of trading. I find that trading costs and profits from liquidity provision are lower for trades that are executed on SEFs than for bilaterally negotiated off-SEF trades. Trading on SEFs is regulated so as to ensure a minimum degree of pre-trade transparency in OTC markets and, in comparison to bilaterally negotiated trades, facilitates comparison shopping and creates direct price competition among liquidity providers. Consistently, I find that on-SEF trades are significantly more likely to get executed within the quoted bid-ask spread than off-SEF trades. The results suggest that CFTC rules introducing SEFs had a compressing effect on trading costs and profits from liquidity provision.

A Appendix to Chapter 1

A.1 Explanatory Variables of CDS Market Illiquidity

Bid-Ask. Bid and ask quotes for EUR or USD denominated senior five-year CDS contracts come from Bloomberg. Contract-specific bid-ask spreads are monthly averages of daily bid-ask spreads, which are calculated whenever more than five nonnegative daily bid-ask spread observations are available within the month. For each month, *Bid-Ask* is the average of contract-specific bid-ask spreads.

ILLIQ^{CDS}. Single-name CDS data for the construction of *ILLIQ^{CDS}* come from Markit. For each reference name i , the *ILLIQ^{CDS}* measure is the monthly average of absolute spread changes divided by the number of contributors to the spread quotation on date t , $Depth_{i,t}$. That is,

$$ILLIQ_{i,m}^{CDS} = \frac{1}{n_{i,m}} \sum_{t=1}^{n_{i,m}} \frac{|C_{i,t} - C_{i,t-1}|}{Depth_{i,t}}, \quad (\text{A.1})$$

where $n_{i,m}$ is the number of consecutive spread changes in month m and $C_{i,t}$ is the five-year par spread. For each month, *ILLIQ^{CDS}* is the average of *ILLIQ^{CDS}* across those reference names with $n_{i,m} > 5$.

ILLIQ^{IDX}. Data for the construction of *ILLIQ^{IDX}* are those described in Section 1.2.3. For each credit index i , the *ILLIQ^{IDX}* measure is the monthly average of absolute changes in the level of the five-year on-the-run series divided by the number of contributors to the quotation of the index level on date t , $Depth_{i,t}^{IDX}$. That is,

$$ILLIQ_{i,m}^{IDX} = \frac{1}{n_{i,m}} \sum_{t=1}^{n_{i,m}} \frac{|C_{i,t}^{IDX} - C_{i,t-1}^{IDX}|}{Depth_{i,t}^{IDX}}, \quad (\text{A.2})$$

where $n_{i,m}$ is the number of consecutive index level changes in month m and $C_{i,t}^{IDX}$ is defined as in Section 1.2.2. For each month, *ILLIQ^{IDX}* is the average of *ILLIQ^{IDX}* across credit indices.

LIB-OIS. USD LIBOR and OIS rates come from Bloomberg. *LIB-OIS* is the monthly average of daily observations of the spread between three-month LIBOR and OIS rates.

Repo. Repo rates come from Bloomberg. *Repo* is the monthly average of daily observations of the spread between three-month Agency MBS and Treasury general collateral repo rates.

Capital. Data for the construction of *Capital* come from Bloomberg. The market capitalization of each financial intermediary that belongs to the G14 group of major credit derivatives dealers is given by the product of the intermediary's share price (shares denominated in currencies other than USD are converted to USD using spot exchange rates) and the number of shares outstanding.¹ *Capital* is the monthly average of the daily aggregate market capitalization of G14 members.

VIX. VIX index levels come from Bloomberg. *VIX* is the monthly average of daily index levels.

Noise. The Hu, Pan, and Wang (2013) "Noise" measure comes from Jun Pan's website <http://www.mit.edu/~junpan>. *Noise* is the monthly average of daily observations.

CDS-Bond. The average CDS-bond basis across U.S. investment-grade bonds comes from J.P. Morgan. *CDS-Bond* is the monthly average of daily observations.

A.2 Excess Return Computation

This Appendix describes the computation of expected and realized excess returns on a CDS trading at par as well as a portfolio composed of such CDSs. It also describes the computation of realized excess returns on a credit index and its replicating portfolio.

¹In 2005, the G14 included Bank of America, Barclays, Bear Stearns, Citigroup, Credit Suisse, Deutsche Bank, Goldman Sachs, HSBC, J.P. Morgan, Lehman Brothers, Merrill Lynch, Morgan Stanley, UBS, and Wachovia (see "Statement regarding developments in the credit derivatives markets," Press release, Federal Reserve Bank of New York, October 5, 2005). When Nomura was added as a 15th member at the end of August 2011 (see "G14 dealer group adds two members," *Risk.net*, December 1, 2011) the group consisted of Bank of America Merrill Lynch, Barclays, BNP Paribas, Citigroup, Credit Suisse, Deutsche Bank, Goldman Sachs, HSBC, J.P. Morgan, Morgan Stanley, Royal Bank of Scotland, Société Générale, UBS, and Wells Fargo. Bear Stearns and Merrill Lynch left the group when they were acquired by J.P. Morgan and Bank of America on June 2, 2008 and January 2, 2009, respectively. Lehman Brothers dropped out on September 15, 2008 when it defaulted and Wells Fargo replaced Wachovia upon acquisition on January 2, 2009. Because we are unable to determine in which order BNP Paribas, Royal Bank of Scotland, and Société Générale joined the group of G14 dealers, we treat them as group members throughout the entire sample period.

A.2.1 Realized CDS Excess Return

When computing realized excess returns, we assume that contracts are marked to market using the ISDA CDS Standard Model, which is the market standard for determining mark-to-market payments in credit derivatives transactions.² Consider a CDS contract referencing entity i with a notional amount of one dollar and fixed spread C .³ On date t , the present value of the contract from the perspective of the protection seller is

$$PV_t(C; C_{i,t}, R_i^*) = Prem_t(C; C_{i,t}, R_i^*) - Prot_t(C_{i,t}, R_i^*), \quad (\text{A.3})$$

where $C_{i,t}$ denotes the date- t par spread and R_i^* denotes the expected recovery rate on (senior unsecured) debt issued by entity i .⁴ The first term on the right hand side of Equation (A.3) is the date- t present value of the premium leg

$$Prem_t(C; C_{i,t}, R_i^*) = C \times PVBP_t(C_{i,t}, R_i^*), \quad (\text{A.4})$$

where

$$PVBP_t(C_{i,t}, R_i^*) = \sum_{j=1}^J \left(\frac{(t_j - t_{j-1})}{360} D(t, t_j) S_i(t, t_j) - \int_{t \vee t_{j-1}}^{t_j} \frac{(u - t_{j-1})}{360} D(t, u) dS_i(t, u) \right) \quad (\text{A.5})$$

is the date- t present value of a risky annuity with payment dates $t < t_1 < \dots < t_J$ ($t_0 \leq t$ being the start date of the CDS contract and t_J being its maturity), $D(t, t_j)$ is the date- t discount factor applicable to a risk free cash flow on date t_j , and $S_i(t, t_j)$ is the date- t risk neutral survival probability of entity i up to date t_j .⁵ The second term of Equation (A.3) is the date- t present value of the protection leg⁶

$$Prot_t(C_{i,t}, R_i^*) = - \int_t^{t_J} (1 - R_i^*) D(t, u) dS_i(t, u). \quad (\text{A.6})$$

The present value of the contract can be decomposed into an accrual amount, $C \times (t - t_0)/360$, and a residual upfront amount. The par spread is defined such that the upfront amount is

² The ISDA CDS Standard Model is a reduced form model which assumes that (i) credit events occur randomly and independently across reference names at the first jump times of homogeneous Poisson processes with constant intensities; (ii) interest rates evolve independent of the occurrence of credit events; and (iii) in case that a credit event occurs, creditors recover a constant fraction of the reference obligation's par value.

³ We follow market standard in assuming that CDS contracts are standardized with respect to their spread payment dates and maturities. Payment dates of standardized contracts fall on the 20th of March, June, September, and December of each year, and the maturity date is the first payment date that follows the trade date by the term of the contract. CDS contracts trade with standardized payment dates and maturities since 2003, see O'Kane (2008).

⁴ Note that the pricing formulas below are identical in case of a constant recovery rate and in case of a random recovery rate that is independently drawn upon the occurrence of a credit event. But only the latter case is consistent with the interpretation of R_i^* as an expected (risk neutral) recovery rate.

⁵ The integral term in Equation (A.5) is the present value of the accrual amount on default.

⁶ Note that both $Prem_t(C; C_{i,t}, R_i^*)$ and $Prot_t(C_{i,t}, R_i^*)$ depend on the par spread, $C_{i,t}$, and the expected recovery rate, R_i^* , via the survival probabilities, $S_i(t, t_j)$.

Appendix A. Appendix to Chapter 1

zero, which means that (A.3) can be equivalently expressed as

$$PV_t(C; C_{i,t}, R_i^*) = (C - C_{i,t}) \left(PVBP_t(C_{i,t}, R_i^*) - \frac{t - t_0}{360} \right) + C \frac{t - t_0}{360}. \quad (\text{A.7})$$

We compute the excess return from t to t' from the protection seller's perspective. At date t , we assume that the protection seller posts initial collateral equal to the notional of the contract. This collateral earns the risk-free rate. In addition, the protection seller pays the protection buyer the present value of the contract, but under standard margining rules, this amount is immediately refunded to the protection seller as variation margin.⁷ Moreover, to simplify matters, we assume that the contract is initiated at the par spread.⁸

If there is no credit event between dates t and t' , the contract is marked to market on date t' and the protection seller receives an amount equal to the change in the present value of the contract from t to t' . In this case, the excess return is

$$r_{i,t,t'}^e = -(C_{i,t'} - C_{i,t}) \left(PVBP_{t'}(C_{i,t'}, R_i^*) - \frac{t' - t}{360} \right) + C_{i,t} \frac{t' - t}{360} \quad (\text{A.8})$$

on the notional amount of the contract.

If there is a credit event between dates t and t' , the excess return is

$$r_{i,t,t'}^e = -(1 - R_i) + C_{i,t} \frac{\tau_i - t}{360}, \quad (\text{A.9})$$

where R_i is the actual recovery rate and the second term of Equation (A.9) is the accrual amount on default (where τ_i is the credit event date).

We use Markit five-year mid spreads and the corresponding expected recovery rates to construct one-week realized excess returns (we denote by $r_{i,t}^e$ the realized excess return over a one-week period ending on date t). Risk free discount factors are bootstrapped from the term structure of LIBOR/swap rates. For each reference name that triggered a credit event, we compute the realized excess return over the one-week period that contains the credit event date, using the actual recovery rate determined in the credit event auction. In case of failure to pay and restructuring credit events, we resume return computations from the first week following the credit event auctions and delete all intermediate data. Our sample includes a total of 22 credit events and losses per dollar of notional range from 23.38% for the Governor and Company of the Bank of Ireland to 98.75% for Landsbanki.

⁷ The interest that is paid on variation margin varies with contract terms. For simplicity, we assume that it is zero.

⁸ Alternatively, we could assume that contracts are traded with upfront amounts and fixed spreads (as we do below when computing realized excess returns on credit indices), which is the convention for trading standardized single-name CDS contracts since the implementation of the ISDA's "Big Bang" Protocol. Assuming that contracts are traded at their par spreads has the advantage that those quotations are available throughout the sample period.

A.2.2 Expected CDS Excess Return

We follow Bongaerts et al. (2011) in defining the date- t conditional expected excess return over the life of a five-year CDS contract as

$$\hat{E}_t[r_{i,t,t_j}^e] = C_{i,t} PVBP_{i,t}^P - EL_{i,t}^P, \quad (\text{A.10})$$

where

$$PVBP_{i,t}^P = \sum_{j=1}^J \left(\frac{(t_j - t_{j-1})}{360} D(t, t_j) P_i(t, t_j) - \int_{t \vee t_{j-1}}^{t_j} \frac{(u - t_{j-1})}{360} D(t, u) dP_i(t, u) \right) - \frac{t - t_0}{360}, \quad (\text{A.11})$$

and

$$EL_{i,t}^P = - \int_t^{t_j} (1 - R_i^*) D(t, u) dP_i(t, u), \quad (\text{A.12})$$

in which the physical survival probability of entity i , $P_i(t, u)$, integrates payoffs instead of the risk neutral survival probability.⁹ Physical survival probabilities are extracted from Moody's KMV one-year and five-year EDFs through

$$P_i(t, t + 1Y) = 1 - EDF1Y_{i,t} \quad \text{and} \quad P_i(t, t + 5Y) = (1 - EDF5Y_{i,t})^5 \quad (\text{A.13})$$

and intermediate values are obtained by interpolation based on the assumption of piecewise constant instantaneous physical default intensities. Conditional expected excess returns for a holding period shorter than five years are obtained by assuming that returns scale proportionally with time-to-maturity. In particular,

$$\hat{E}_t[r_{i,t,t+1}^e] = \frac{7}{t_j - t} \hat{E}_t[r_{i,t,t_j}^e]. \quad (\text{A.14})$$

A.2.3 Portfolio Excess Returns

Because we consider equally weighted portfolios of reference names, the realized excess return on a portfolio p of five-year CDS contracts over the one-week period from $t - 1$ to t is the average realized excess return on the $n_{p,t-1}$ CDS contracts that constitute portfolio p on date $t - 1$; that is,

$$r_{p,t}^e = \frac{1}{n_{p,t-1}} \sum_{i \in \mathcal{S}_{p,t-1}} r_{i,t}^e, \quad (\text{A.15})$$

where $\mathcal{S}_{p,t-1}$ denotes the set of reference names. Similarly, the date- t conditional expected excess return on the portfolio is the average conditional expected excess return on the CDS

⁹By using the same expected recovery rate under the risk neutral and physical probability measure, we implicitly assume that there is no recovery risk premium.

contracts that constitute the portfolio; that is,

$$\widehat{E}_t[r_{p,t,t'}^e] = \frac{1}{n_{p,t}} \sum_{i \in \mathcal{I}_{p,t}} \widehat{E}_t[r_{i,t,t'}^e] = C_{p,t} PVBP_{p,t}^P - EL_{p,t}^P \quad (\text{A.16})$$

where the portfolio level quantities in Equation (A.16) are defined as

$$C_{p,t} = \frac{1}{n_{p,t}} \sum_{i \in \mathcal{I}_{p,t}} C_{i,t}, \quad (\text{A.17})$$

$$EL_{p,t}^P = \frac{1}{n_{p,t}} \sum_{i \in \mathcal{I}_{p,t}} EL_{i,t}^P, \quad (\text{A.18})$$

$$PVBP_{p,t}^P = \frac{1}{n_{p,t}} \sum_{i \in \mathcal{I}_{p,t}} \frac{C_{i,t}}{C_{p,t}} PVBP_{i,t}^P. \quad (\text{A.19})$$

A.2.4 Realized Credit Index Excess Return

Finally, consider a five-year credit index contract with one dollar notional amount. The contract trades with fixed spread C and date- t upfront amount, $UF_{i,t}^{IDX}(C)$, that is received by the seller of credit protection. As in the case of single-name CDS contracts, we compute “unlevered” realized excess returns on credit index contracts, $r_{i,t,t'}^{IDX}$, assuming that contracts are covered by collateral agreements and standard margining rules apply; that is,

$$r_{i,t,t'}^{IDX} = - \left(UF_{i,t'}^{IDX}(C) - UF_{i,t}^{IDX}(C) \right) + \frac{I_t}{I} C \frac{t' - t}{360} - \frac{1}{I} (L_{i,t'} - L_{i,t}), \quad (\text{A.20})$$

where $L_{i,t}$ is the cumulative loss due to credit events among index constituents on date t .¹⁰ Replacing the upfront amounts in Equation (A.20) with those on the replicating basket of single-name CDSs gives the “unlevered” realized excess return on the replicating basket, $r_{i,t,t'}^{CDS}$. Hence, realized credit index excess returns can be readily computed from Markit’s credit index data described in Section 1.2.3 (as before we denote by $r_{i,t}^{IDX}$ and $r_{i,t}^{CDS}$, respectively, the realized excess returns over a one-week period ending on date t). Whenever an index roll date, t_{roll} , falls between dates t and t' , the realized excess return is obtained by first computing the realized excess return on series S_i of index i between t and t_{roll} and then adding to it the realized excess return on series S_{i+1} over t_{roll} to t' .

¹⁰As in case of CDS excess returns, losses accumulated in $L_{i,t}$ are given as one minus the recovery value determined in a credit event auction. The expression in Equation (A.20) presumes that all defaults between dates t and t' occur exactly on date t' . In our implementation, we account for the fact that defaults may happen in between t and t' and adjust the accrual term in Equation (A.20) accordingly.

A.3 CDS Spread Decomposition

From Equation (A.10), the five-year CDS spreads can be expressed in terms of conditional expected excess returns and expected default losses; that is,

$$C_{i,t} = \frac{EL_{i,t}^P + \widehat{E}_t[r_{i,t,t_J}^e]}{PVBPP_{i,t}^P}. \quad (\text{A.21})$$

Conditional expected one-week excess returns are decomposed according to Equation (1.8) and are converted to a five-year holding period by multiplication with $(t_J - t)/7$, see Equation (A.14). Replacing $\widehat{E}_t[r_{i,t,t_J}^e]$ in Equation (A.21) with the resulting expression gives the following decomposition of reference entity i 's five-year CDS spread on date t :

$$C_{i,t} = \frac{EL_{i,t}^P}{PVBPP_{i,t}^P} + \frac{(t_J - t)\widehat{\beta}_i^{DEF}\lambda_{DEF,t}}{7 \times PVBPP_{i,t}^P} + \frac{(t_J - t)\widehat{\beta}_i^{LIQ}\lambda_{LIQ,t}}{7 \times PVBPP_{i,t}^P} + \frac{(t_J - t)u_{i,t}}{7 \times PVBPP_{i,t}^P}. \quad (\text{A.22})$$

The first term on the right hand side of Equation (A.22) is the expected default loss, while the second and third terms are the default and liquidity risk premia, respectively, and the last term is the pricing error. The components of our CDS spread decomposition are the sample means of the terms in Equation (A.22). Note that the same decomposition holds for portfolio level CDS spreads using the respective expressions in Appendix A.2.

A.4 Robustness Checks: Factor Constructions

Conditional volatility weighting. The construction of our tradable liquidity factor is similar to that of Moskowitz et al.'s (2012) time series momentum factor in that it aggregates signed returns. To account for the considerable cross-sectional variation in volatilities across assets, Moskowitz et al. (2012) scale returns by their conditional volatilities. We construct a tradable liquidity factor in a similar way with weights inversely proportional to conditional volatilities; that is,

$$LIQ_t^{CVW} = \sum_{i=1}^{n_t} w_{i,t-1}^{CVW} \text{sgn}(B_{i,t-1}) (r_{i,t}^{IDX} - r_{i,t}^{CDS}), \quad (\text{A.23})$$

where $w_{i,t}^{CVW} = (1/\sigma_{i,t})/(\sum_{j=1}^{n_t} 1/\sigma_{j,t})$ and $\sigma_{i,t}^2$ is an estimate of annualized conditional variance of $r_{i,t}^{IDX} - r_{i,t}^{CDS}$ that is obtained from *daily* returns as in Equation (1) of Moskowitz et al. (2012). Because we use the first six-month period to estimate the conditional volatilities for the computation of the alternative liquidity factor's first observation, its time series consists of 252 weekly observations from March 28, 2007 to February 1, 2012. The alternative liquidity factor has a correlation of 0.97 with the benchmark liquidity factor indicating that our index-level-based weighting scheme is effectively a weighting by conditional volatilities.

Corporate bond market illiquidity factor. We use transaction data from the Financial In-

dustry Regulatory Authority's Trade Reporting and Compliance Engine (TRACE) to construct the corporate bond market illiquidity factor. In particular, we obtain transaction data for plain-vanilla fixed-rate bullet bonds issued by U.S. corporations. The data are filtered for erroneous transactions using Dick-Nielsen's (2009) methodology and, as in Dick-Nielsen et al. (2012), transactions with par volume below 100,000 USD are discarded. Bond-specific Amihud (2002) illiquidity measures are obtained each day by averaging absolute returns of consecutive transactions per million dollar of par volume traded. These are converted to a weekly frequency by taking the within-week median of daily measures. Each week the market-wide measure is obtained as the weighted average (by amount issued) of bond-specific measures. The corporate bond market illiquidity factor is then given as the residual of an AR(2) specification of the market-wide illiquidity measure. When converting bond-specific Amihud (2002) illiquidity measures to a monthly rather than a weekly frequency, the resulting corporate bond illiquidity measure has a correlation of 0.93 in levels and 0.79 in first differences with Dick-Nielsen et al.'s (2012) λ .

Stock market illiquidity factor. To construct the stock market illiquidity factor, we obtain price, return, and volume data for NYSE- and AMEX-traded ordinary common shares of U.S. companies from the CRSP daily stock file. Individual-stock Amihud (2002) illiquidity measures are given as weekly averages of absolute one-day returns per million dollar of daily trading volume. By construction the individual-stock measures are very noisy and outliers may have a nonnegligible impact when aggregating them into a market-wide illiquidity measure. We, therefore, follow Korajczyk and Sadka (2008) and "Winsorize" the individual-stock measures for a given week at the 1st and 99th percentiles of their distribution. Each week the market-wide illiquidity measure is obtained as the cross-sectional mean of "Winsorized" individual-stock measures, and the stock market illiquidity factor is the residual of an AR(2) specification of the market-wide measure.

A.5 Standard Error Computation

A.5.1 Standard Errors of Factor Price of Risk Estimates

We describe the standard error computation for a general K -dimensional vector of factors, $f_t = [f_{1,t}, \dots, f_{K,t}]'$, and the most general case that we consider in the paper, namely the case of a cross-sectional regression with an intercept, a characteristic, and an additional univariate beta. In this case, the counterparts of Equations (1.4) and (1.5) in vector notation are

$$r_t = \alpha + \beta f_t + \epsilon_t, \tag{A.24}$$

and

$$\mu_\xi = 1_N \lambda_0 + \mu_c \lambda_c + \beta \lambda + \beta_k^* \lambda_k = X\gamma, \tag{A.25}$$

where $r_t = [r_{1,t}^e, \dots, r_{N,t}^e]'$ is the N -dimensional vector of realized excess returns, α denotes the N -dimensional vector of regression intercepts, β denotes the $N \times K$ matrix of factor betas, $\epsilon_t = [\epsilon_{1,t}, \dots, \epsilon_{N,t}]'$ is the N -dimensional vector of mean zero error terms, μ_ξ denotes the mean of the N -dimensional vector of conditional expected excess returns, $\xi_t = [\xi_{1,t}, \dots, \xi_{N,t}]'$ with $\xi_{i,t} = \widehat{E}_t[r_{i,t+1}^e]$, μ_c denotes the mean of the N -dimensional vector of characteristics, $c_t = [c_{1,t}, \dots, c_{N,t}]'$, β_k^* denotes the N -dimensional vector of univariate betas of \bar{y}_t with respect to the k -th factor, $f_{k,t}$ with $1 \leq k \leq K$, y_t is an N -dimensional vector of exogenous variables, $y_t = [y_{1,t}, \dots, y_{N,t}]'$, $\bar{y}_t = y_t - \beta_{yg}^* g_t$ is the N -dimensional vector of exogenous variables orthogonalized with respect to an additional factor g_t , the $N \times (K + 3)$ matrix X and the $(K + 3)$ -dimensional vector γ are defined by $X = [1_N, \mu_c, \beta, \beta_k^*]$ and $\gamma = [\lambda_0, \lambda_c, \lambda', \lambda_k]'$, respectively, and 1_N denotes an N -dimensional vector of ones. Note that in contrast to the standard two-pass cross-sectional regression method, there is a distinction between expected excess returns, μ_ξ , and the mean of realized excess returns, μ_r .

Moreover, we define the $d = (K + 1 + 4N)$ -dimensional vector $Y_t = [f_t', g_t', r_t', c_t', y_t', \xi_t']'$ and denote its mean and covariance matrix by $\mu = [\mu_f', \mu_g, \mu_r', \mu_c', \mu_y', \mu_\xi']'$ and V , respectively. In what follows, we will use the following convenient partition of V ,

$$V = \begin{bmatrix} V_f & V_{gf}' & V_{rf}' & V_{cf}' & V_{yf}' & V_{\xi f}' \\ V_{gf} & V_g & V_{rg}' & V_{cg}' & V_{yg}' & V_{\xi g}' \\ V_{rf} & V_{rg} & V_r & V_{cr}' & V_{yr}' & V_{\xi r}' \\ V_{cf} & V_{cg} & V_{cr} & V_c & V_{yc}' & V_{\xi c}' \\ V_{yf} & V_{yg} & V_{yr} & V_{yc} & V_y & V_{\xi y}' \\ V_{\xi f} & V_{\xi g} & V_{\xi r} & V_{\xi c} & V_{\xi y} & V_\xi \end{bmatrix}, \quad (\text{A.26})$$

and express factor betas and univariate betas in terms of the elements of V . The matrix of factor betas is given by $\beta = V_{rf} V_f^{-1}$ and the vector of univariate betas of \bar{y}_t with respect to the k -th factor is given by

$$\beta_k^* = \beta_{\bar{y}f}^* \iota_k = (\beta_{yf}^* - \beta_{yg}^* \beta_{gf}^*) \iota_k = V_{yf} D^{-1} \iota_k - V_{yg} V_g^{-1} V_{gf} D^{-1} \iota_k, \quad (\text{A.27})$$

where $D = \text{diag}(V_f)$, $\beta_{yf}^* = V_{yf} D^{-1}$ denotes the $N \times K$ matrix of univariate betas of y_t with respect to f_t , $\beta_{yg}^* = V_{yg} V_g^{-1}$ denotes the N -dimensional vector of univariate betas of y_t with respect to g_t , $\beta_{gf}^* = V_{gf} D^{-1}$ denotes the $1 \times K$ matrix of univariate betas of g_t with respect to f_t , and ι_k denotes the K -dimensional unit vector whose k -th element is nonzero. As in Kan et al. (2013), we assume that Y_t is stationary and ergodic with finite fourth moment.

Under a potentially misspecified model, there is no γ such that Equation (A.25) is satisfied and γ is chosen to minimize the sum of squared population pricing errors, $e = \mu_\xi - X\gamma$; that is,

$$\gamma = \underset{\delta}{\text{argmin}} (\mu_\xi - X\delta)' (\mu_\xi - X\delta) = (X'X)^{-1} X' \mu_\xi. \quad (\text{A.28})$$

Appendix A. Appendix to Chapter 1

Note that with e defined as above, γ satisfies the first-order conditions

$$X'e = 0_{K+3} \Leftrightarrow 1'_N e = 0, \quad \mu'_c e = 0, \quad \beta' e = 0_K, \quad \text{and} \quad \beta_k^{*'} e = 0, \quad (\text{A.29})$$

where 0_m denotes an m -dimensional vector of zeros. From the final expression in Equation (A.28) an estimate of γ can be obtained by replacing population moments with their sample counterparts; that is,

$$\hat{\gamma} = (\hat{X}'\hat{X})^{-1}\hat{X}'\hat{\mu}_\xi, \quad (\text{A.30})$$

where $\hat{X} = [1_N, \hat{\mu}_c, \hat{\beta}, \hat{\beta}_k^*]$, $\hat{\beta}$ and $\hat{\beta}_k^*$ are given by $\hat{V}_{rf}\hat{V}_f^{-1}$ and $\hat{V}_{yf}\hat{D}^{-1}l_k - \hat{V}_{yg}\hat{V}_g^{-1}\hat{V}_{gf}\hat{D}^{-1}l_k$, respectively, and $\hat{\mu}$.s and \hat{V} .s are the corresponding elements of

$$\hat{\mu} = \frac{1}{T} \sum_{t=1}^T Y_t, \quad (\text{A.31})$$

and

$$\hat{V} = \frac{1}{T} \sum_{t=1}^T (Y_t - \hat{\mu})(Y_t - \hat{\mu})', \quad (\text{A.32})$$

respectively.

Note that $\hat{\theta} = [\hat{\mu}', \text{vec}(\hat{V})']'$ is the method of moments estimator of $\theta = [\mu', \text{vec}(V)']'$. Under the above assumptions,¹¹

$$\sqrt{T}(\hat{\theta} - \theta) \xrightarrow[T \rightarrow \infty]{d} N(0_{d(1+d)}, S_0), \quad (\text{A.33})$$

where $S_0 = \sum_{j=-\infty}^{\infty} E[\psi(Y_t; \theta)\psi(Y_{t+j}; \theta)']$ and $\psi(Y_t; \theta)$ is the moment function,

$$\psi(Y_t; \theta) = [(Y_t - \mu)', \text{vec}((Y_t - \mu)(Y_t - \mu)' - V)']'. \quad (\text{A.34})$$

Since γ is a smooth function of θ , an application of the delta method yields

$$\sqrt{T}(\hat{\gamma} - \gamma) \xrightarrow[T \rightarrow \infty]{d} N(0_{K+3}, (\partial\gamma/\partial\theta')S_0(\partial\gamma/\partial\theta)'). \quad (\text{A.35})$$

Using the expression for S_0 from above, the asymptotic covariance matrix of $\hat{\gamma}$, i.e., $(\partial\gamma/\partial\theta')S_0(\partial\gamma/\partial\theta)'$, becomes $\sum_{j=-\infty}^{\infty} E[h_t h_{t+j}']$, with $h_t = (\partial\gamma/\partial\theta')\psi(Y_t; \theta)$.

In order to find an explicit expression for h_t it remains to compute $\partial\gamma/\partial\theta'$. Using the above

¹¹As noted by Kan et al. (2013), S_0 is a singular matrix. This is due to the fact that \hat{V} is symmetric, i.e., it contains linearly dependent elements. One could alternatively consider the parameter vector $\tilde{\theta} = [\mu', \text{vech}(V)']'$, in which case the covariance matrix of the limiting normal distribution would be nonsingular.

partition of θ , we have $\partial\gamma/\partial\theta' = [\partial\gamma/\partial\mu', \partial\gamma/\partial\text{vec}(V)']$ and

$$h_t = \frac{\partial\gamma}{\partial\theta'} \psi(Y_t; \theta) = \frac{\partial\gamma}{\partial\mu'} (Y_t - \mu) + \frac{\partial\gamma}{\partial\text{vec}(V)'} \text{vec}((Y_t - \mu)(Y_t - \mu)' - V). \quad (\text{A.36})$$

With $H = (X'X)^{-1}$ and $A = HX'$, the Jacobian matrices $\partial\gamma/\partial\mu'$ and $\partial\gamma/\partial\text{vec}(V)'$ are given by

$$\frac{\partial\gamma}{\partial\mu'} = \left[0_{(K+3) \times (K+1+N)}, \frac{\partial\gamma}{\partial\text{vec}(X)'} \frac{\partial\text{vec}(X)}{\partial\mu'_c}, 0_{(K+3) \times N}, A \right], \quad (\text{A.37})$$

and

$$\frac{\partial\gamma}{\partial\text{vec}(V)'} = \frac{\partial\gamma}{\partial\text{vec}(X)'} \frac{\partial\text{vec}(X)}{\partial\text{vec}(V)'}, \quad (\text{A.38})$$

respectively, where $0_{m \times n}$ denotes an $m \times n$ matrix of zeros and

$$\frac{\partial\gamma}{\partial\text{vec}(X)'} = (H \otimes e') - (\gamma' \otimes A). \quad (\text{A.39})$$

Now, note that $\text{vec}(X) = [1'_N, \mu'_c, \text{vec}(\beta)', \beta_k^{*'}]'$. Thus,

$$\frac{\partial\text{vec}(X)}{\partial\mu'_c} = [0_{N \times N}, I_N, 0_{N \times N \cdot (K+1)}]' = ([0, 1, 0'_{K+1}]' \otimes I_N), \quad (\text{A.40})$$

where I_N denotes the N -dimensional identity matrix, and, consequently,

$$\frac{\partial\gamma}{\partial\mu'} (Y_t - \mu) = A(\xi_t - \mu_\xi) - A(c_t - \mu_c)\lambda_c + H[0, c'_t e, 0'_{K+1}]'. \quad (\text{A.41})$$

Similarly,

$$\frac{\partial\text{vec}(X)}{\partial\text{vec}(V)'} = \left[0_{d^2 \times 2N}, \left(\frac{\partial\text{vec}(\beta)}{\partial\text{vec}(V)'} \right)', \left(\frac{\partial\beta_k^*}{\partial\text{vec}(V)'} \right)' \right]'. \quad (\text{A.42})$$

For the remaining expressions, we get

$$\frac{\partial\text{vec}(\beta)}{\partial\text{vec}(V)'} = (V_f^{-1} \otimes I_N) \frac{\partial\text{vec}(V_{rf})}{\partial\text{vec}(V)'} - (V_f^{-1} \otimes \beta) \frac{\partial\text{vec}(V_f)}{\partial\text{vec}(V)'}, \quad (\text{A.43})$$

and

$$\begin{aligned} \frac{\partial\beta_k^*}{\partial\text{vec}(V)'} &= (l'_k D^{-1} \otimes I_N) \frac{\partial\text{vec}(V_{yf})}{\partial\text{vec}(V)'} - (l'_k \beta_{gf}^*{}' V_g^{-1} \otimes I_N) \frac{\partial V_{yg}}{\partial\text{vec}(V)'} \\ &\quad + (l'_k \beta_{gf}^*{}' V_g^{-1} \otimes \beta_{yg}^*) \frac{\partial V_g}{\partial\text{vec}(V)'} - (l'_k D^{-1} \otimes \beta_{yg}^*) \frac{\partial\text{vec}(V_{gf})}{\partial\text{vec}(V)'} \\ &\quad - (l'_k D^{-1} \otimes \beta_{yf}^*) \Theta \frac{\partial\text{vec}(V_f)}{\partial\text{vec}(V)'}, \end{aligned} \quad (\text{A.44})$$

Appendix A. Appendix to Chapter 1

where Θ is a $K^2 \times K^2$ matrix such that $\text{vec}(D) = \Theta \text{vec}(V_f)$. Thus, the Jacobian matrix $\partial \text{vec}(X)/\partial \text{vec}(V)'$ can be expressed as the following sum of Kronecker products

$$\begin{aligned} \frac{\partial \text{vec}(X)}{\partial \text{vec}(V)'} &= ([0_{K \times 2}, V_f^{-1}, 0_K]' \otimes I_N) \frac{\partial \text{vec}(V_{rf})}{\partial \text{vec}(V)'} - ([0_{K \times 2}, V_f^{-1}, 0_K]' \otimes \beta) \frac{\partial \text{vec}(V_f)}{\partial \text{vec}(V)'} \\ &+ ([0_{K \times (K+2)}, D^{-1} \iota_k]' \otimes I_N) \frac{\partial \text{vec}(V_{yf})}{\partial \text{vec}(V)'} - ([0'_{K+2}, V_g^{-1} \beta_{gf}^* \iota_k]' \otimes I_N) \frac{\partial V_{yg}}{\partial \text{vec}(V)'} \\ &+ ([0'_{K+2}, V_g^{-1} \beta_{gf}^* \iota_k]' \otimes \beta_{yg}^*) \frac{\partial V_g}{\partial \text{vec}(V)'} - ([0_{K \times (K+2)}, D^{-1} \iota_k]' \otimes \beta_{yg}^*) \frac{\partial \text{vec}(V_{gf})}{\partial \text{vec}(V)'} \\ &- ([0_{K \times (K+2)}, D^{-1} \iota_k]' \otimes \beta_{yf}^*) \Theta \frac{\partial \text{vec}(V_f)}{\partial \text{vec}(V)'} \end{aligned} \quad (\text{A.45})$$

Moreover,

$$\frac{\partial \text{vec}(V_{rf})}{\partial \text{vec}(V)'} = ([I_K, 0_{K \times (4N+1)}] \otimes [0_{N \times (K+1)}, I_N, 0_{N \times 3N}]), \quad (\text{A.46})$$

$$\frac{\partial \text{vec}(V_f)}{\partial \text{vec}(V)'} = ([I_K, 0_{K \times (4N+1)}] \otimes [I_K, 0_{K \times (4N+1)}]), \quad (\text{A.47})$$

$$\frac{\partial \text{vec}(V_{yf})}{\partial \text{vec}(V)'} = ([I_K, 0_{K \times (4N+1)}] \otimes [0_{N \times (K+1+2N)}, I_N, 0_{N \times N}]), \quad (\text{A.48})$$

$$\frac{\partial V_{yg}}{\partial \text{vec}(V)'} = ([0'_K, 1, 0'_{4N}] \otimes [0_{N \times (K+1+2N)}, I_N, 0_{N \times N}]), \quad (\text{A.49})$$

$$\frac{\partial V_g}{\partial \text{vec}(V)'} = ([0'_K, 1, 0'_{4N}] \otimes [0'_K, 1, 0'_{4N}]), \quad (\text{A.50})$$

$$\frac{\partial \text{vec}(V_{gf})}{\partial \text{vec}(V)'} = ([I_K, 0_{K \times (4N+1)}] \otimes [0'_K, 1, 0'_{4N}]). \quad (\text{A.51})$$

Substituting the expressions in Equations (A.39) and (A.45) into Equation (A.38), and using Equations (A.46)–(A.51) as well as the first-order conditions yields

$$\begin{aligned} \frac{\partial \gamma}{\partial \text{vec}(V)'} \text{vec}((Y_t - \mu)(Y_t - \mu)' - V) &= H \tilde{z}_t \\ &+ A \{ \beta(f_t - \mu_f)(f_t - \mu_f)' - (r_t - \mu_r)(f_t - \mu_f)' \} V_f^{-1} \lambda \\ &+ A \{ \beta_{yf}^* D_t - (y_t - \mu_y)(f_t - \mu_f)' + \beta_{yg}^* (g_t - \mu_g)(f_t - \mu_f)' \} D^{-1} \iota_k \lambda_k \\ &- A \{ \beta_{yg}^* (g_t - \mu_g)^2 - (y_t - \mu_y)(g_t - \mu_g) \} V_g^{-1} \beta_{gf}^* \iota_k \lambda_k, \end{aligned} \quad (\text{A.52})$$

where

$$\begin{aligned} \tilde{z}_t &= [0'_2, u_t(f_t - \mu_f)' V_f^{-1}, v_t(f_t - \mu_f)' D^{-1} \iota_k - e' \beta_{yg}^* (g_t - \mu_g)(f_t - \mu_f)' D^{-1} \iota_k \dots \\ &+ e' \{ \beta_{yg}^* (g_t - \mu_g)^2 - (y_t - \mu_y)(g_t - \mu_g) \} V_g^{-1} \beta_{gf}^* \iota_k - e' \beta_{yf}^* D_t D^{-1} \iota_k]' \end{aligned} \quad (\text{A.53})$$

$u_t = e'(r_t - \mu_r)$, $v_t = e'(y_t - \mu_y)$, and $D_t = \text{diag}((f_t - \mu_f)(f_t - \mu_f)')$.

Finally, adding up the terms in Equations (A.41) and (A.52), h_t can be explicitly expressed as

$$\begin{aligned}
 h_t = & (\gamma_t - \gamma) - A(c_t - \mu_c)\lambda_c + A\{\beta(f_t - \mu_f)(f_t - \mu_f)' - (r_t - \mu_r)(f_t - \mu_f)'\}V_f^{-1}\lambda \\
 & + A\{\beta_{yf}^*D_t - (y_t - \mu_y)(f_t - \mu_f)'\}D^{-1}l_k\lambda_k \\
 & - A\{\beta_{yg}^*(g_t - \mu_g)^2 - (y_t - \mu_y)(g_t - \mu_g)\}V_g^{-1}\beta_{gf}^*l_k\lambda_k \\
 & - A\beta_{yg}^*\{\beta_{gf}^*D_t - (g_t - \mu_g)(f_t - \mu_f)'\}D^{-1}l_k\lambda_k + Hz_t,
 \end{aligned} \tag{A.54}$$

where $\gamma_t = A\xi_t$ and

$$\begin{aligned}
 z_t = & [0, c_t'e, u_t(f_t - \mu_f)'V_f^{-1}, v_t(f_t - \mu_f)'D^{-1}l_k - e'\beta_{yg}^*(g_t - \mu_g)(f_t - \mu_f)'D^{-1}l_k \dots \\
 & + e'\{\beta_{yg}^*(g_t - \mu_g)^2 - (y_t - \mu_y)(g_t - \mu_g)\}V_g^{-1}\beta_{gf}^*l_k - e'\beta_{yf}^*D_tD^{-1}l_k]'.
 \end{aligned} \tag{A.55}$$

Applying the Newey and West (1987) method, a heteroscedasticity and autocorrelation consistent estimator for the asymptotic covariance matrix of $\hat{\gamma}$ is given by

$$\frac{1}{T} \sum_{t=1}^T \hat{h}_t \hat{h}_t' + \frac{1}{T} \sum_{l=1}^m \sum_{t=l+1}^T \left(1 - \frac{l}{m+1}\right) (\hat{h}_t \hat{h}_{t-l}' + \hat{h}_{t-l} \hat{h}_t'), \tag{A.56}$$

where \hat{h}_t is given by Equation (A.54) with population parameters replaced by their sample estimates. In particular, $\hat{e} = \hat{\mu}_\xi - \hat{X}\hat{\gamma}$. The finite sample approximation of $\hat{\gamma}$'s covariance matrix is then obtained as $1/T$ times the estimate of the asymptotic covariance matrix.

Based on Equation (A.54) it is straightforward to break down asymptotic variation of $\hat{\gamma}$ into three components. The first one, $\gamma_t - \gamma$, is variation of $\hat{\gamma}$ in case that the model is correctly specified and estimated using population values, i.e., there is no error associated with the estimation of the characteristic, μ_c , and betas, β and β_k^* . The second source of variation are errors-in-variables (EIV). The second term of Equation (A.54) captures variation associated with the estimation of the characteristic, μ_c , the third term of Equation (A.54) captures variation associated with the estimation of factor betas, β , the fourth, fifth, and sixth terms of Equation (A.54) capture variation associated with the estimation of the univariate betas, β_{yf}^* , β_{yg}^* , and β_{gf}^* , respectively. Variation from the first two sources is, e.g., accounted for by generalized-method-of-moments-based inference. The third source of variation is due to potential model misspecification and captured by $H z_t$. Note that this term vanishes when the model is correctly specified, i.e., when $e = \mu_\xi - X\gamma = 0_N$. Thus, setting $e = 0_N$ gives the asymptotic variance of $\hat{\gamma}$ in a generalized method of moments estimation of μ_c , β , β_{yf}^* , and γ . As mentioned above, this asymptotic variance takes into account EIV but ignores potential model misspecification.

In case that the intercept is restricted to zero, $\gamma = [\lambda_c, \lambda', \lambda_k]'$, $X = [\mu_c, \beta, \beta_k^*]$, and h_t is given

by Equation (A.54), where

$$z_t = [c_t' e, u_t(f_t - \mu_f)' V_f^{-1}, v_t(f_t - \mu_f)' D^{-1} \iota_k - e' \beta_{yg}^* (g_t - \mu_g)(f_t - \mu_f)' D^{-1} \iota_k \dots \\ + e' \{\beta_{yg}^* (g_t - \mu_g)^2 - (y_t - \mu_y)(g_t - \mu_g)\} V_g^{-1} \beta_{gf}^* \iota_k - e' \beta_{yf}^* D_t D^{-1} \iota_k]' , \quad (\text{A.57})$$

and A , H , and e are defined as above. In case that the model specification does not include the characteristic, $\gamma = [\lambda_0, \lambda', \lambda_k]'$, $X = [1_N, \beta, \beta_k^*]$, and h_t is given by

$$h_t = (\gamma_t - \gamma) + A \{\beta(f_t - \mu_f)(f_t - \mu_f)' - (r_t - \mu_r)(f_t - \mu_f)'\} V_f^{-1} \lambda \\ + A \{\beta_{yf}^* D_t - (y_t - \mu_y)(f_t - \mu_f)'\} D^{-1} \iota_k \lambda_k \\ - A \{\beta_{yg}^* (g_t - \mu_g)^2 - (y_t - \mu_y)(g_t - \mu_g)\} V_g^{-1} \beta_{gf}^* \iota_k \lambda_k \\ - A \beta_{yg}^* \{\beta_{gf}^* D_t - (g_t - \mu_g)(f_t - \mu_f)'\} D^{-1} \iota_k \lambda_k + H z_t, \quad (\text{A.58})$$

where

$$z_t = [0, u_t(f_t - \mu_f)' V_f^{-1}, v_t(f_t - \mu_f)' D^{-1} \iota_k - e' \beta_{yg}^* (g_t - \mu_g)(f_t - \mu_f)' D^{-1} \iota_k \dots \\ + e' \{\beta_{yg}^* (g_t - \mu_g)^2 - (y_t - \mu_y)(g_t - \mu_g)\} V_g^{-1} \beta_{gf}^* \iota_k - e' \beta_{yf}^* D_t D^{-1} \iota_k]' , \quad (\text{A.59})$$

and A , H , and e are defined as above. Finally, in case that the model specification does not include the univariate beta, $\gamma = [\lambda_0, \lambda_c, \lambda']'$, $X = [1_N, \mu_c, \beta]$, and h_t is given by

$$h_t = (\gamma_t - \gamma) - A(c_t - \mu_c) \lambda_c + A \{\beta(f_t - \mu_f)(f_t - \mu_f)' - (r_t - \mu_r)(f_t - \mu_f)'\} V_f^{-1} \lambda \\ + H z_t, \quad (\text{A.60})$$

where $z_t = [0, e' c_t, u_t(f_t - \mu_f)' V_f^{-1}]'$, and A , H , and e are defined as above.

A.5.2 Standard Error of the Cross-Sectional R^2

The standard error computation of the cross-sectional R^2 is based on the same principle as that of the factor price of risk estimates. Again, we derive standard errors for the most general case that we consider in the paper and we discuss less general cases at the end of this section.

Let ρ^2 denote the population value of the R^2 ; that is,

$$\rho^2 = 1 - \frac{Q}{Q_0} = 1 - \frac{e' e}{e_0' e_0}, \quad (\text{A.61})$$

where $e_0 = (I_N - (1/N)1_N 1_N') \mu_\xi$ are population deviations of expected excess returns from their cross-sectional average. Replacing population values in Equation (A.61) by their sample estimates, obviously, gives the R^2 .

Assume

unable to explain any cross-sectional variation in expected excess returns. As in the previous

section, ρ^2 is a smooth function of θ and an application of the delta method yields

$$\sqrt{T}(R^2 - \rho^2) \xrightarrow[T \rightarrow \infty]{d} N(0, (\partial \rho^2 / \partial \theta') S_0 (\partial \rho^2 / \partial \theta)'), \quad (\text{A.62})$$

where S_0 is defined as in the previous section, $(\partial \rho^2 / \partial \theta') S_0 (\partial \rho^2 / \partial \theta')' = \sum_{j=-\infty}^{\infty} E[\eta_t \eta_{t+j}]$, and $\eta_t = (\partial \rho^2 / \partial \theta') \psi(Y_t; \theta)$. Thus, it remains to compute $\partial \rho^2 / \partial \theta'$ in order to obtain an explicit expression for η_t . Using the above partition of θ , we have $\partial \rho^2 / \partial \theta' = [\partial \rho^2 / \partial \mu', \partial \rho^2 / \partial \text{vec}(V)']$ and

$$\eta_t = \frac{\partial \rho^2}{\partial \theta'} \psi(Y_t; \theta) = \frac{\partial \rho^2}{\partial \mu'} (Y_t - \mu) + \frac{\partial \rho^2}{\partial \text{vec}(V)'} \text{vec}((Y_t - \mu)(Y_t - \mu)' - V). \quad (\text{A.63})$$

The Jacobian matrices $\partial \rho^2 / \partial \mu'$ and $\partial \rho^2 / \partial \text{vec}(V)'$ are given by

$$\frac{\partial \rho^2}{\partial \mu'} = \left[0'_{K+1+N}, \frac{\partial \rho^2}{\partial \text{vec}(X)'} \frac{\partial \text{vec}(X)}{\partial \mu'_c}, 0'_{N'}, \frac{2}{Q_0} \{(1 - \rho^2) e'_0 - e'\} \right], \quad (\text{A.64})$$

and

$$\frac{\partial \rho^2}{\partial \text{vec}(V)'} = \frac{\partial \rho^2}{\partial \text{vec}(X)'} \frac{\partial \text{vec}(X)}{\partial \text{vec}(V)'}, \quad (\text{A.65})$$

respectively, with

$$\frac{\partial \rho^2}{\partial \text{vec}(X)'} = -\frac{2}{Q_0} (\gamma' \otimes e'). \quad (\text{A.66})$$

Replacing $\partial \text{vec}(X) / \partial \mu'_c$ and $\partial \text{vec}(X) / \partial \text{vec}(V)'$ by the expressions derived in the previous section and making use of the first-order conditions yields

$$\begin{aligned} \eta_t = & \frac{2}{Q_0} \{ \{(1 - \rho^2) e'_0 - e'\} (\xi_t - \mu_\xi) + e' c_t \lambda_c + u_t (f_t - \mu_f)' V_f^{-1} \lambda \\ & - e' \{ \beta_{yf}^* D_t - (y_t - \mu_y) (f_t - \mu_f)' \} D^{-1} l_k \lambda_k \\ & + e' \{ \beta_{yg}^* (g_t - \mu_g)^2 - (y_t - \mu_y) (g_t - \mu_g) \} V_g^{-1} \beta_{gf}^* l_k \lambda_k \\ & + e' \beta_{yg}^* \{ \beta_{gf}^* D_t - (g_t - \mu_g) (f_t - \mu_f)' \} D^{-1} l_k \lambda_k \}, \end{aligned} \quad (\text{A.67})$$

where, as before, $u_t = e' (r_t - \mu_r)$ and $D_t = \text{diag}((f_t - \mu_f)(f_t - \mu_f)')$. As in the previous section, the Newey and West (1987) method applied to η_t 's sample analog, $\hat{\eta}_t$, gives a heteroscedasticity and autocorrelation consistent estimate of the asymptotic variance of the R^2 .

In case that the intercept is restricted to zero, the expression of η_t for the standard error computation does not change.¹² The expressions for η_t in case that the model specification does not include the characteristic or the univariate beta can be obtained from Equation (A.67)

¹²Note that we do not redefine ρ^2 in case that the intercept is restricted to zero. Therefore, ρ^2 is not necessarily nonnegative. Nevertheless, $\rho^2 \leq 1$ and $\rho^2 = 1$ if and only if $e = 0_N$.

Appendix A. Appendix to Chapter 1

Credit Rating	Expected Excess Returns				Realized Excess Returns			
	Q1	Price Impact			Q1	Price Impact		
		Q2	Q3	Q4		Q2	Q3	Q4
AAA-AA	0.36	0.39	0.44	0.61	-0.50	-0.36	-0.86	0.65
	[6.28]	[6.23]	[6.71]	[6.98]	[-0.70]	[-0.42]	[-0.85]	[0.42]
A	0.37	0.45	0.58	0.95	-0.44	-0.34	0.31	-0.25
	[7.82]	[6.91]	[4.91]	[4.43]	[-0.62]	[-0.33]	[0.19]	[-0.08]
BBB	0.50	0.70	1.01	1.58	-0.72	-0.12	0.96	1.71
	[8.92]	[6.80]	[6.23]	[5.26]	[-0.75]	[-0.08]	[0.39]	[0.39]
BB	1.26	1.84	2.31	3.07	-0.87	1.19	2.80	5.23
	[9.72]	[10.41]	[8.30]	[8.81]	[-0.46]	[0.34]	[0.58]	[0.79]
B-CCC	2.58	3.11	4.34	5.74	-1.33	1.04	12.13	23.82
	[11.24]	[10.73]	[6.48]	[6.00]	[-0.29]	[0.14]	[0.96]	[1.41]
Credit Rating	CDS Spreads				Standard Deviations			
	Q1	Price Impact			Q1	Price Impact		
		Q2	Q3	Q4		Q2	Q3	Q4
AAA-AA	0.42	0.54	0.75	1.17	1.16	1.79	2.72	4.83
A	0.51	0.65	0.90	2.08	1.23	1.80	2.65	6.20
BBB	0.70	0.98	1.44	2.79	1.57	2.36	3.65	6.09
BB	1.82	2.85	3.67	5.97	3.30	6.02	7.75	11.02
B-CCC	4.52	6.29	9.31	17.01	8.58	12.40	19.95	27.94

Table A.1: Descriptive Statistics of Price-Impact-Sorted Portfolios.

The table displays descriptive statistics for the 20 portfolios formed by first sorting CDS contracts according to credit ratings and then according to price impact. The upper part of the table reports sample means of conditional expected excess returns (in % per year) and realized excess returns (in % per year). In brackets are t -statistics based on Newey and West (1987) heteroscedasticity and autocorrelation consistent standard errors with 24 lags. The lower part of the table reports sample means of average five-year CDS spreads across portfolio constituents (in % per year) and standard deviations of realized excess returns (in % per year). Portfolio time series consist of 276 weekly observations from October 11, 2006 to February 1, 2012.

by setting the respective parameters, i.e., λ_c or λ_k , equal to zero.

A.6 Additional Figures and Tables

Figure A.1 depicts monthly time series of the explanatory variables of CDS market illiquidity (thin black lines) and the CDS market illiquidity measure (thick gray lines).

Figure A.2 displays the CDS spread decomposition for the 20 price-impact-sorted portfolios.

Table A.1 displays descriptive statistics for the 20 price-impact-sorted portfolios.

Table A.2 summarizes index rules for the main indices of the CDX North American and iTraxx Europe credit index families as well as their sub-indices.

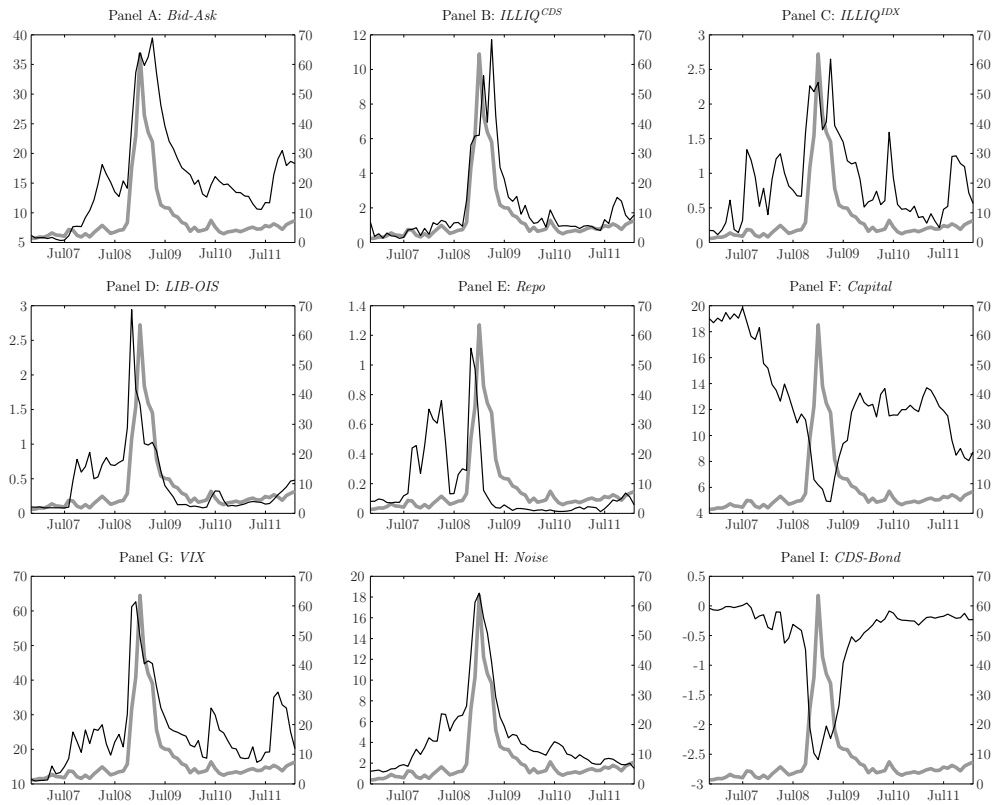


Figure A.1: Explanatory Variables of CDS Market Illiquidity.

The figure displays monthly observations of the explanatory variables of CDS market illiquidity (thin black lines, left hand scales) and the CDS market illiquidity measure (thick gray lines, right hand scales). The explanatory variables are: the average bid-ask spread of single-name CDSs (*Bid-Ask*), the average absolute spread change per quote contributed across single-name CDSs ($ILLIQ^{CDS}$), the average absolute change in the index level per quote contributed across on-the-run credit indices ($ILLIQ^{IDX}$), the spread between three-month LIBOR and OIS rates (*LIB-OIS*), the spread between three-month Agency MBS and Treasury general collateral repo rates (*Repo*), the aggregate market capitalization of financial institutions that make up the G14 group of major credit derivatives dealers (*Capital*), the VIX index (*VIX*), the Hu, Pan, and Wang (2013) “Noise” measure (*Noise*), and the average CDS-bond basis across U.S. investment-grade bonds (*CDS-Bond*). *CDS-Bond*, *LIB-OIS*, *Repo*, and *VIX* are in %. *Bid-Ask*, *CDSILLIQ*, and *Noise* are in basis points. $ILLIQ^{CDS}$ and $ILLIQ^{IDX}$ are in basis points per contributed quote. *Capital* is in 100 billion USD. The time series consist of 64 monthly observations from October 2006 to January 2012.

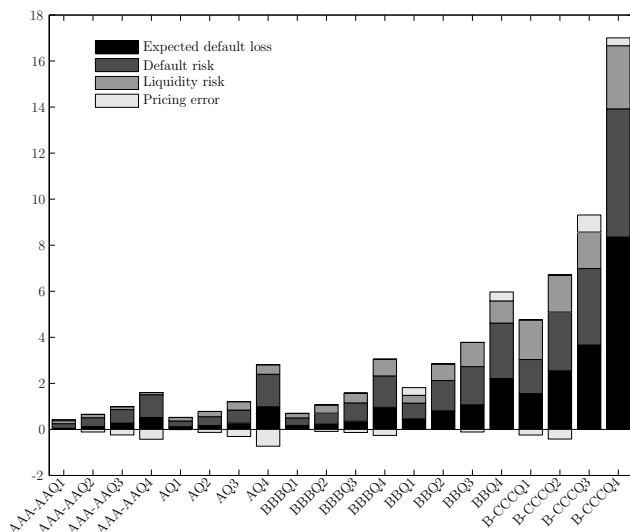


Figure A.2: CDS Spread Decomposition of Price-Impact-Sorted Portfolios. The figure displays five-year CDS spreads (in % per year) of the price-impact-sorted portfolios. CDS spreads are decomposed into expected default losses, factor risk premia, and pricing errors implied by the benchmark model specification. The horizontal axis displays portfolio identifiers.

Main index (No. of constituents) Sub-indices (No. of constituents)	CDX.NA.IG (125) CDX.NA.IG.HVOL (30)	CDX.NA.HY (100) CDX.NA.HY.BB (≤100) CDX.NA.HY.B (≤100)	iTraxx Eur (125) iTraxx Eur HVol (30) iTraxx Eur Sr.Finls (25) iTraxx Eur Sub.Finls (25)	iTraxx Xover (≤50)
Eligible reference names (ref. names) Domicile of eligible ref. names Credit rating of eligible ref. names Index roll dates Inclusion & exclusion (main index)	Corporate & Financial ^a North America Investment grade March and September 20 th / h	Corporate & Financial ^a North America High yield or not rated March and September 27 th / h	Corporate & Financial ^a Europe Investment grade March and September 20 th / h	Corporate Europe High yield or not rated March and September 20 th / h
Inclusion (sub-indices)	Eligible ref. names that are not members of the current index series and that rank among the most liquid 20% in terms of market risk activity in the Depository Trust & Clearing Corporation's (DTCC's) Trade Information Warehouse (TIW) over the six-month period preceding an index roll date are to be included in the next index series. Eligible ref. names that are members of the current index series and that rank among the most illiquid 30% in terms of market risk activity in the DTCC's TIW over the six-month period preceding an index roll date are excluded from the new series along with ref. names that are members of the current index series but no longer meet eligibility criteria.	Eligible ref. names that are not members of the current index series and that rank among the most illiquid 30% in terms of market risk activity in the DTCC's TIW over the six-month period preceding an index roll date are excluded from the new series along with ref. names that are members of the current index series but no longer meet eligibility criteria. In addition, most illiquid eligible ref. names are removed from and most liquid ref. names are added to the new series until the weight of industry sectors and the weight of credit rating categories in the index approximately corresponds to that of the most recent series of the Markit iBoxx USD Liquid High Yield Index.	Eligible ref. names have to meet the following additional eligibility criterion: ref. names have to show weekly transaction activity in the DTCC TIW over the last eight weeks prior to the last Friday of the month preceding an index roll date. Eligible ref. names are ranked within their industry sectors in terms of market risk activity in the DTCC's TIW over the six-month period preceding an index roll date. Then the most liquid ref. names of each industry sector are included in the new series of the index which consists of 30 ref. names from the Autos & Industrials sector, 30 ref. names from the Consumers sector, 20 ref. names from the Energy sector, 20 ref. names from the Telecommunications, Media, and Technology sector, and 25 ref. names from the Financial sector.	Eligible ref. names have to meet the following additional eligibility criteria: (i) ref. names have to show weekly transaction activity in the DTCC TIW over the last eight weeks prior to the last Friday of the month preceding an index roll date and (ii) ref. names' average five-year CDS spreads over the last ten London trading days of the month preceding an index roll date have to be at least twice the average spread of the iTraxx Eur Non Finls index over this period. Eligible ref. names are ranked in terms of their market risk activity in the DTCC's TIW over the six-month period preceding an index roll date. Then the 50 most liquid eligible ref. names are included in the new index series. In case that less than 50 eligible ref. names meet the inclusion criteria the number of index constituents will be adjusted to a multiple of five by including non-eligible ref. names that violate eligibility criteria (i) and/or (ii) and that have issued a bond with maturity date comparable to that of a five-year CDS contract whose credit spread satisfies criteria (ii) from above.
Currency of index contract	USD	USD	EUR	EUR
Tier of index contract	Senior unsecured debt	Senior unsecured debt	Senior unsecured debt ^b	Senior unsecured debt
Documentation clause of index contract	No restructuring	No restructuring	Modified-modified restructuring	Modified-modified restructuring
Maturity of index contract in years	1, 2, 3, 5, 7, 10 ^c	3, 5, 7, 10	3, 5, 7, 10 ^d	3, 5, 7, 10
Maturity dates of index contract	June and December 20 th / h	June and December 20 th / h	June and December 20 th / h	June and December 20 th / h
Quotation of index contract	Spread	Price	Spread	Spread

Table A.2: Summary of Index Rules.

^aMarket makers of the index are not eligible for inclusion. ^bTier for the Sub Finls sub-index is subordinated or lower Tier 2 debt. ^cThe one- and two-year contract maturities of CDX.NA.IG were introduced later than the other contract maturities. ^dContract maturities of the Sr Finls and Sub Finls sub-indices are five and ten years.

B Appendix to Chapter 2

B.1 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 index CDSs in three implementation phases. In the first phase, index CDS dealers and private funds active in the index CDS market (so-called Category 1 Entities) are required to clear their index CDS transactions. In the second phase, financial entities other than Category 3 Entities (so-called Category 2 Entities) are required to clear their index CDS transactions. In the third phase, investment managers and pension plans (so-called Category 3 Entities) are required to clear their index CDS transactions. End-users, i.e., non-financial entities hedging business risk, are exempt from mandatory central clearing.
Dec 31, 2012	Real-time public reporting of index CDS transactions becomes mandatory for index CDS dealers.
Feb 28, 2014	Real-time public reporting of index CDS transactions becomes mandatory for major index CDS market participants.
Mar 11, 2013	Central clearing becomes mandatory for Category 1 Entities trading CDX.IG or CDX.HY (for transactions in the five-year tenor, mandatory central clearing applies to series 11 and all subsequent series).
Apr 10, 2013	Real-time public reporting of index CDS transactions becomes mandatory for any index CDS market participant.
May 31, 2013	The CFTC publishes the final block-trade rules. ¹

¹Block trades are exempt from the trade execution requirement and may be publicly disseminated with delay.

Appendix B. Appendix to Chapter 2

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 methods of execution for swaps that are subject to mandatory trade execution on SEFs (either against orders resting on a SEF's order book or against a response to a RFQ facilitated by the SEF); and (ii.) the process that SEFs can initiate (via so-called made available to trade determinations) to get CFTC approval for mandatory trade execution of certain swaps on SEFs. ²
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 transactions of notional amounts exceeding certain spread- and 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). ³
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 and CDX.HY with five-year tenors.
Feb 26, 2014	The approved made available to trade determination becomes effective and all transactions in the above-mentioned index CDSs (not qualifying as block trades or being end-user exempt) must be executed on SEFs.

B.2 Data Processing

This section gives a detailed account of the data that we use in our empirical analysis, the procedures that we use to account for outliers in the data, and the algorithms that we apply to identify swap execution facilities (SEFs) and package transactions.

²Swaps eligible for "made available to trade" determinations have to be subject to mandatory central clearing.

³Prior to July 30, 2013, all index CDS transaction were publicly disseminated with delay and for transactions with notional amount exceeding USD 100 million, the disseminated notional amounts were capped at USD 100 million.

B.2.1 On-SEF Trade Report History

We collect all trade reports of credit asset class swaps executed on or before October 16, 2015 from three of the four operating swap data repositories (SDRs): the Bloomberg Swap Data Repository (BSDR), the Depository Trust & Clearing Corporation Data Repository (DDR), and the Intercontinental Exchange Trade Vault (ICETV).⁴ The DDR started operating on the effective date of the Commodity Futures Trading Commission's (CFTC's) real-time public reporting requirement for swap dealers, December 31, 2012, and the ICETV and BSDR started operating on February 9, 2013 and May 12, 2014, respectively. Nevertheless, there are trade reports of transaction executed prior to December 31, 2012 because the DDR trade report history contains some historical swap transactions that fall under the CFTC's recordkeeping requirement.⁵

There are three types of trade reports that can be submitted to SDRs: new trade reports, cancellations, and corrections. New trade reports are used to submit transaction data, cancellations are used to cancel a previously submitted trade report that contains erroneous transactions data, and corrections are used to submit the correct transaction data of a previously canceled trade report. The CFTC's real-time public reporting requirement specifies that cancellations and corrections should be submitted in the above order by the party that submitted the erroneous trade report (the so-called reporting party).⁶

From each SDR's trade report history, we remove canceled trade reports and the corresponding cancellations. We also remove duplicate corrections (in case that a correction was not only submitted by the reporting party) and corrections that cannot be traced back to the trade reports that they are supposed to correct. In case of corrections of non-canceled trade reports, we remove both the corrections and the non-canceled trade reports.

We then remove all trade reports of transactions that were not executed on SEFs and all trade reports of on-SEF transactions that were executed prior to the CFTC's SEF compliance date, October 2, 2013, on which temporarily registered SEFs started operating. We also remove trade reports of non-price-forming transactions, such as amendments, novations, and terminations, and trade reports of transactions in contracts other than index CDSs, index swaptions, and index tranche swaps.

Next, we remove all trade reports for which we cannot identify the underlying. These include trade reports of transaction in which the underlying is not a standardized credit index of corporate, municipal, or sovereign creditors, and trade reports with missing or incomplete data items (or fields) that we use to identify the underlying.⁷ For BSDR trade reports this

⁴The trade report history of the fourth SDR, the Chicago Mercantile Exchange Swap Data Repository, consists of a total of 65 trade reports. None of the trade reports would be included in our sample because the transactions were not executed on SEFs.

⁵See Part 43 and Part 46 of Chapter I of Title 17 of the Code of Federal Regulations (17 CFR).

⁶See §43.3(e) of Chapter I of 17 CFR.

⁷Standardized credit indices are uniquely identified by the index's name, the index's series number (which uniquely identifies the creditors in the index), and the index's version number (which keeps track of the creditors

Appendix B. Appendix to Chapter 2

concerns the index's Bloomberg ticker (BSDR field name "ticker") and its version number (BSDR field name "cvs version"),⁸ for DDR trade reports this concerns the index's Reference Entity Database (RED) code (DDR field name "UNDERLYING_ASSET_1"; specifically, the last nine-digits of the item),⁹ and for ICETV trade reports this concerns the index's ICETV product mnemonic (ICETV field name "TVProductMnemonic"; the latter can be mapped via ICETV product definitions to the "TvProductName" field from which the index can be identified).

Finally, we remove trade reports with incomplete transaction data (such as, execution timestamps, transaction prices, and trade sizes), and trade reports of transactions in which the underlying is neither a CDX.IG nor a CDX.HY index. We then merge the trade reports of the three SDRs which amounts to mapping field names used by each of the three SDRs individually to those commonly defined by the CFTC.¹⁰ Where applicable, we augment trade reports that do not provide data for certain fields with the respective standardized contract terms. For example, ICETV trade reports do not specify the day-count convention (ICETV field name "DayCountConvention") which we populate by ACT/360, the day-count convention of standardized index CDSs on CDX.IG and CDX.HY indices.

B.2.2 Identification of SEFs

Each of the transactions that remain in the trade report history must have been executed on one of the eleven SEFs that are registered with the CFTC and offer trading in index CDSs. These are BGC Derivative Markets, Bloomberg SEF, DW SEF, GFI Swaps Exchange, ICAP SEF, ICE Swap Trade, MarketAxess SEF, TeraExchange, tpSEF, Tradition SEF, and TW SEF. However, according to volume data that SEFs have to make publicly available (usually on their websites) on a daily basis,¹¹ no transactions in index CDSs, index swaptions, or index tranche swaps on CDX.IG or CDX.HY indices have been executed on TeraExchange during the sample period.¹²

that have been removed from the index due to credit events).

⁸During the first two weeks of operation, the BSDR disseminated trade reports with missing Bloomberg tickers and version numbers. Nevertheless, the underlying of transactions in indices of the CDX North American family can be identified under the assumption that (i.) the underlying credit index is not a sub-index of the index whose name is contained in the trade report (BSDR field name "underlying asset 1"; this may not be an unreasonable assumption given that Bloomberg SEF—the only SEF that reports transaction data to the BSDR—does not list sub-indices of the CDX North American family, and given that in the period following the first two weeks of operation, there are no trade reports in the BSDR history in which the underlying is such a sub-index); and (ii.) the version number of the underlying credit index series is that of the version with the maximum number of contributors to Markit's end-of-day composite computation on the trade date (the number of contributors is a reliable indicator of trading activity because it usually shifts from one version to the next after the respective credit event auctions have taken place which is consistent with anecdotal evidence that activity shifts on these days because dealers prefer to continue trading in versions that include the defaulted names until recovery values are determined; this is because attachment and detachment points of new versions of a tranche swap can only be set once recovery rates are known).

⁹RED codes uniquely identify the index's name, the index's series number, and the index's version number.

¹⁰See Appendix A to Part 43 of Chapter I of 17 CFR for the definition of data fields for public dissemination.

¹¹See Part 16 of Chapter I of 17 CFR.

¹²Clarus FT, the standard data source for SEF market shares and volumes, collects these data and makes them available historically (see Section B.3 for details).

As explained in §43.3(a)(2) of Chapter I of 17 CFR, SEFs are responsible for reporting transaction data of trades executed on their platforms to SDRs. This in particular means that the choice to which SDR the transaction data are reported is with the SEF and, in general, not with the counterparties to the transaction (see 77 Federal Register 1198 (Jan. 9, 2012) for a clarifying discussion). Bloomberg SEF reports cleared transactions to the BSDR while non-cleared transactions are reported to the DDR, the SDR to which Bloomberg SEF reported all transaction data before the BSDR started operating.¹³ ICE Swap Trade states in its rulebook that it generally reports cleared transactions to the ICETV and non-cleared transactions to the DDR.¹⁴ All other SEFs seem to report transaction data to the DDR.¹⁵ Therefore, we identify the transactions of on-SEF trade reports disseminated by the BSDR as being executed on Bloomberg SEF, and we identify the transactions of on-SEF trade reports disseminated by the ICETV as being executed on ICE Swap Trade. For transactions of on-SEF trade reports disseminated by the DDR it is possible to identify the SEF on which trade execution took place based on the format the SEF used for trade reporting.

Specifically, the SEF that submitted a trade report to the DDR can be identified based on the format in which the underlying and the price notation type is reported (the corresponding DDR field names are “UNDERLYING_ASSET_1” and “PRICE_NOTATION_TYPE”).¹⁶ The different formats that SEFs use for reporting transactions in index CDSs on CDX.IG or CDX.HY indices are exhibited in Table B.1. The table’s entries are based on an identification strategy that is illustrated by means of the following example:

1. For a given index CDS contract (characterized in terms of the underlying index and the contract’s tenor) and date, search for the unique underlying and (case-sensitive) price notation type formats among DDR trade reports that are executed on SEFs. For each such pair of formats, sum up the notional amount of non-block trades. The result of such a search for five-year CDX.IG.21 on February 19, 2014 is, for instance:

Underlying	Price Notation Type	Non-Block Notional (USD MM)
2I65BY:2I65BYCX1	Percentage	100.1
CDX.NA.IG.21:2I65BYCX1	Basis points	200
CDX.NA.IG.21:2I65BYCX1	BasisPoints	250
Dow Jones CDX Investment Grade21 V1:2I65BYCX1	Basispoints	850

¹³See rulebook of Bloomberg SEF and “Change in SDR Reporting,” *Notice to Bloomberg SEF Participants*, June 23, 2014.

¹⁴See rulebook of ICE Swap Trade.

¹⁵For instance, BGC Derivative Markets, GFI Swaps Exchange, ICAP SEF, MarketAxess SEF, and tpSEF state in their rulebooks that they are reporting credit asset class transaction data to the DDR. Tradition SEF leaves the choice to the counterparties of the transaction, and DW SEF and TW SEF do not further specify the SDR to which they report transaction data.

¹⁶In some cases, the “CLEARED” field, the “INDICATION_OF_COLLATERALIZATION” field, or the “DAY_COUNT_CONVENTION” field has to be taken into account as well when identifying the SEF. SEF identification based on collateralization seems economically insensible because collateralization should be transaction-specific rather than SEF-specific but could be consistent with SEFs failing to report collateralization details to SDRs or choosing not to do so in case of trades that are centrally cleared.

Appendix B. Appendix to Chapter 2

IG.21:2I65BYCX1	Percentage	136
MARKIT CDX.NA.IG.21 12/18 CME:2I65BYCX1	Basis points	0.2
MARKIT CDX.NA.IG.21 12/18 ICE:2I65BYCX1	Basis points	2,924.1

2. Identify the underlying and (case-sensitive) price notation type format with the SEF that has reported approximately the same non-block notional amount for the index CDS contract (characterized in terms of the security name and tenor) on that date.¹⁷ The following excerpt is a screenshot of Clarus FT’s SEF View for the above example.

SEF ^	Reporting Date	Asset Class	Product Type	Tenor	CCY Pair	Security	Amount	CCY	Open	Close	High	Low	Block Notional	Block USD Equiv MM	Non-Block USD Equiv MM
BBG	19-Feb-2014	CRD	CRD: Index	5YR		CDX.NA.IG-521	5,599,600,000	U...	64.5	66.1...	66.3...	63.8492	2,675,000,000	2,675	2,925
GFI	19-Feb-2014	CRD	CRD: Index	5YR		CDX.NA.IG-521	850,000,000	U...	64.8125	63.5	66.8...	65.75			850
ICAP	19-Feb-2014	CRD	CRD: Index	5YR		CDX.NA.IG-521	250,000,000	U...	64.5	66.1...	66.1...	63.9375			250
ICE	19-Feb-2014	CRD	CRD: Index	5YR		CDX.NA.IG-521	50,000,000	U...	64.75	65	65	64.75			50
MarketAxess	19-Feb-2014	CRD	CRD: Index	5YR		CDX.NA.IG-521	100,100,000	U...	64.69	65.7					100
TP	19-Feb-2014	CRD	CRD: Index	5YR		CDX.NA.IG-521	200,000,000	U...	64.125	66	66	64.125			200
TW	19-Feb-2014	CRD	CRD: Index	5YR		CDX.NA.IG-521	259,775,000	U...					124,075,000	124	136
Total							7,309,475,000							2,799	4,510

In this example, “MARKIT CDX.NA.IG.21 12/18 CME:2I65BYCX1” and “MARKIT CDX.NA.IG.21 12/18 ICE:2I65BYCX1” share a common underlying format of the form “MARKIT CDX.NA.IG.[Series] [Term as mm/yy] [CCP]:[Nine-digit RED code]” (where CCP denotes the central clearing party) and the non-block notional amounts of transactions with these formats sum up to USD 2,924.3 million, which is approximately the value of USD 2,925 million reported for BBG (i.e., Bloomberg SEF) on Clarus FT’s SEF View. In similar vein, “Dow Jones CDX Investment Grade21 V1:2I65BYCX1” corresponds to GFI (i.e., GFI Swaps Exchange), “2I65BY:2I65BYCX1” corresponds to MarketAxess (i.e., MarketAxess SEF), and “IG.21:2I65BYCX1” corresponds to TW (i.e., TW SEF). Trade reports whose underlying asset field is populated with “CDX.NA.IG.21:2I65BYCX1” have to be further differentiated by their price notation type field values; “Basis points” and “BasisPoints”. The former corresponds to TP (i.e., tpSEF) and the latter corresponds to ICAP (i.e., ICAP SEF). Note that among DDR trade reports on February 19, 2014 there are no further trade reports of on-SEF transactions in five-year CDX.IG.21 that have a non-block notional amount of USD 50 million and that could be identified with ICE (i.e., ICE Swap Trade). This is because CDS index transactions executed on ICE Swap Trade are reported to the ICETV. Indeed, the non-block notional amount of on-SEF transactions in five-year CDX.IG.21 in the ICETV trade report history is USD 50 million on February 19, 2014.

The identification strategy focuses on non-block trades because these tend to have non-capped notional amounts in the transaction data. This is because minimum applicable block sizes are by definition less than or equal to cap sizes.¹⁸ However, apart from having a notional

¹⁷Because trade reports contain rounded notional amounts (see §43.4(g) of Chapter I of 17 CFR), non-block notional amounts cannot be expected to be identical.

¹⁸See §43.4(h) of Chapter I of 17 CFR.

amount above the minimum applicable block size, the CFTC's block trade definition includes additional conditions that a trade has to meet in order to qualify as a block. Thus, there might be non-block trades with notional amounts beyond cap sizes and these trades will complicate identification due to cap-induced mismatches of transaction-data-based non-block notional amounts and SEF-reported non-block notional amounts. However, the incidence of mismatches is lower among non-block trades than among all trades and, therefore, we focus on non-block trades. The issue is discussed in more detail in Section B.2.3.

Table B.1 does not contain an entry for BGC Derivative Markets because the SEF only reported index swaption trades (according to Clarus FT data). Index swaptions (as well as index tranche swaps) are typically traded "with delta," i.e., two counterparties that enter into an index swaption simultaneously enter into the underlying index CDS with a notional amount that makes the overall position approximately neutral to changes of the index CDS spread. The index CDS transaction is typically referred to as a "delta exchange" because it is non-price-forming in that the index swaption is already quoted relative to a reference level at which the index CDS will be traded. BGC Derivative Markets does not seem to report the notional amount of delta exchanges as index CDS trading volume and, therefore, we lack data to verify the trade report format which we would associate with the SEF. Given that all of the trades that we would associate with the SEF are delta exchanges, they would not be part of the sample on which most of our analyses are based.

Table B.1 does not contain an entry for Tradition SEF because the SEF does not seem to use a consistent reporting format for its transaction data. Therefore, we hand match index CDS trades on dates on which the SEF reported index CDS, index swaption, or index tranche swap trading (according to Clarus FT data). Specifically, we look for trades in the respective index CDS contract with transaction prices in the high-low range reported by the SEF, whose aggregate notional amount is consistent with the non-block figure reported by the SEF, and that have previously not been assigned to another SEF.¹⁹ We also hand match ICE Swap Trade transaction that have been reported to the DDR by looking for trades in the respective index CDS contract with transaction prices in the high-low range reported by the SEF, whose aggregate notional amount is consistent with the non-block figure reported by the SEF on dates for which the SEF reports trading of uncleared index CDSs (when reporting SEF volumes, ICE Swap Trade differentiates between clearing houses and missing values indicate uncleared trades).

Despite being able to identify trades executed on DW SEF (the interdealer broker SEF operated by Tradeweb), we will not consider them in our analyses. This is because, we are confronted with contradicting information regarding index CDS trading on DW SEF. On the one hand, we were told by Tradeweb personnel that the SEF never offered index CDS trading. On the other hand, the SEF seems to have filed for listing index CDSs in a class certification submitted to

¹⁹A "notional amount ... consistent with the non-block figure" may be below the SEF-reported figure due to capping.

Swap Execution Facility	Index Name	Underlying Asset	Price Notation Type	Additional Conditions
Bloomberg SEF	CDX.IG	MARKIT CDX.NA.IG.[Series]...	—	—
	CDX.HY	MARKIT CDX.NA.HY.[Series]...	—	—
DW SEF	CDX.IG	IG.[Series] [Term]:...	—	—
	CDX.HY	HY.[Series] [Term]:...	—	—
GFI Swaps Exchange	CDX.IG	Dow Jones CDX Investment Grade [Series] V [Version]....	—	—
	CDX.HY	Dow Jones CDX HY [Series] 500 V [Version]....	—	—
ICAP SEF	CDX.IG	CDX.NA.IG.[Series]:...	Basis Points	Day Count Convention="ACT/360"
	CDX.HY	CDX.NA.IG.[Series]-V [Version]....	Percentage	if Cleared="C" then Collateralization=""
MarketAccess SEF	CDX.HY	CDX.NA.HY.[Series]:...	Basis Points	Day Count Convention="ACT/360"
	CDX.IG	CDX.NA.HY.[Series]-V [Version]....	Percentage	if Cleared="C" then Collateralization=""
	CDX.IG	[Six-digit RED code]:...	—	—
	CDX.HY	CDX.NA.IG.[Series]- [Term]:...	—	—
tpSEF	CDX.IG	[Six-digit RED code]:...	—	—
	CDX.HY	CDX.NA.IG.[Series]- [Term]:...	Basis points	Collateralization="PC"
TW SEF	CDX.HY	CDX.NA.HY.[Series]:...	Price	Collateralization="PC"
	CDX.IG	IG.[Series]:...	—	—
	CDX.HY	IG.[Series]-V [Version]....	—	—
		HY.[Series]:...	—	—
		HY.[Series]-V [Version]....	—	—

Table B.1: Swap Execution Facilities' Reporting Formats.

The table shows formats of Depository Trust & Clearing Corporation Data Repository (DDR) trade reports that are identified with different swap execution facilities. DDR trade report fields referenced in the table are cleared (DDR field name "CLEARED"), collateralization (DDR field name "INDICATION_OF_COLLATERALIZATION"), day count convention (DDR field name "DAY_COUNT_CONVENTION"), price notation type (DDR field name "PRICE_NOTATION_TYPE"), and underlying asset (DDR field name "UNDERLYING_ASSET_1"). [Series] denotes the index's series number, [Version] denotes the index's version number, [Term] denotes the index's term (as the number of years followed by a "Y", e.g., "5Y" for a five-year term), [Six-digit RED code] denotes the index's six-digit RED code, and three dots denote any string of zero or more characters. In the additional conditions column capitalized words indicate field names and expressions in quotation marks indicate field values (" " is an empty field). The price notation type and day count convention are case sensitive. Multi-row entries in the underlying asset and price notation type columns should be read as or conditions whereas multi-row entries in the additional condition column should be read as and conditions.

the CFTC on September 30, 2013,²⁰ and it seems to have reported some volume data in the first five weeks following the SEF compliance date (according to Clarus FT data). Given that only very few trades (26 in total) have been executed on DW SEF during this period, omitting them has no material impact on our results.

B.2.3 Assessing the Identification Algorithm's Performance

Assessing the performance of the SEF identification algorithm is rather difficult because, on the one hand, notional amounts in the transaction data are rounded and capped while aggregate volumes reported by SEFs are not, and, on the other hand, most SEFs do not report the number of transactions executed on their platforms. The exception are tpSEF and Tradition SEF. Because Tradition SEF trades are hand-matched, tpSEF is the only SEF for which the performance of the SEF identification algorithm can be assessed in terms of the number of executed transactions. During the sample period tpSEF reported 2,899 transactions in CDX.IG index CDSs and 4,283 transactions in CDX.HY index CDSs. This compares to 2,897 and 4,285 transactions, respectively, that the algorithm identified as being executed on tpSEF.

Figure B.1 compares the aggregate non-block notional amount reported by SEFs (using Clarus FT data) with the one constructed from the transaction data separately for each SEF. Because minimum block sizes are by definition less than or equal to cap sizes, focusing on the non-block notional amount seems sensible at first sight; these transactions should have non-capped notional amounts below minimum block sizes and, therefore, allow for a meaningful comparison with (uncapped) SEF-reported volumes. The caveat is that block status depends not only on trade size but also on the means of trade execution. By definition a block trade has to occur “away from the registered [SEF]’s ... trading system or platform” and be “executed pursuant to the registered [SEF]’s ... rules and procedures.”²¹ This part of the definition was temporarily overruled on September 19, 2014 by a no-action relief of the CFTC which further specifies that “SEFs are permitted to use [request for quote (RFQ)] functionalities to facilitate the execution of a block trade” but “block trades may not be facilitated through a SEF’s Order Book functionality.”²² To further clarify matters, the CFTC notes that “trades above the minimum block size may occur on the SEF’s Order Book however they will not receive treatment as block trades and will not be afforded a reporting time delay.” Thus, for SEFs on which a significant portion of trades is executed on the order book (anecdotally, these are the SEFs operated by interdealer brokers; that is, GFI Swaps Exchange, ICAP SEF, tpSEF, and Tradition SEF) even the comparison of non-block notional amounts may not be meaningful because block-sized (and eventually capped) transactions do not qualify as block trades and, therefore, render the comparison of non-block notional amounts vulnerably to cap-induced mismatches.

²⁰See “DW SEF LLC Certification to CFTC for listing IRS and CDS” on Tradeweb’s website.

²¹See §43.2 of Chapter I of 17 CFR.

²²See CFTC Letter No. 14-118 (Sep. 19, 2014) and its extension CFTC Letter No. 15-60 (Nov. 2, 2015). For instance, block trades can be executed by a RFQ to one other participant on the SEF (because block trades are exempt from the CFTC’s trade execution requirement, this also applies to made available to trade index CDSs).

Appendix B. Appendix to Chapter 2




Figure B.1: Non-Block Notional Amount Identified (in USD Billion).

The figure shows the aggregate non-block notional amount traded on different swap execution facilities (SEFs) between October 2, 2013 and October 16, 2015. For each SEF, the dark gray bar is the SEF-reported non-block notional amount traded (from Clarus FT) aggregated over all index CDSs on CDX.IG and CDX.HY, and the light gray bar is the aggregate notional amount of non-block trades in index CDS contracts referencing CDX.IG and CDX.HY identified as being executed on the respective SEF. The percentages of the SEF-reported notional amounts that are identified in the transaction data are indicated to the right of the arrows at the end of each light gray bar.

This limits the true information content of Figure B.1 to the four SEFs where the majority of trades are dealer-to-customer (D2C) and executed via RFQs (i.e., Bloomberg SEF, ICE Swap Trade, MarketAxess SEF, and TW SEF) because non-block trades on these SEFs tend to be disseminated with uncapped notional amounts. For all of these SEFs but ICE Swap Trade the aggregate non-block notional amount in the transaction data is close to 100% of the amount reported by the SEF, indicating very good performance of the SEF identification algorithm. Moreover, ICE Swap Trade generally reports its cleared index CDS trades to the ICETV and, therefore, the comparably low fraction of the aggregate non-block notional amount in the transaction data is most likely not due to the identification algorithm. Instead it could be due to a non-negligible fraction of block-sized trades executed on the SEF's order book.

B.2.4 Identification of Package Transactions

We impose additional structure on the data by accounting for the fact that some index CDS transactions may be part of packages, i.e., trades that involve more than one index CDS or an index CDS and a related instrument, such as, an index swaption or an index tranche swap. Specifically, we account for four popular packages: index rolls, curve trades, index swaptions with delta exchange, and index tranche swaps with delta exchange. Packages are typically quoted in relative terms. For instance, quotes of index rolls and curve trades are understood to be the difference between quotes on the individual legs of the package and quotes of index

swaptions and index tranche swaps are relative to the reference level of the delta exchange. Therefore, when a package is executed, the transaction prices on the individual legs of the package or the delta exchange may be different from those of non-package transactions (hereafter referred to as outright trades) in the respective index CDSs that are executed at about that time.

Delta Exchanges of Index Tranche Swaps

In order to identify delta exchanges of index tranche swaps, we first apply a reporting-format-based SEF identification algorithm to on-SEF index tranche swap transactions disseminated by the DDR.²³ There are a total of 575 on-SEF index tranche swap transactions and for 571 of those we are able to identify the SEF on which the transaction was executed. About half of the index tranche swap transactions (273) are executed on ICAP SEF. The other half of index tranche swap transactions is executed on ICE Swap Trade (139), GFI Swaps Exchange (86), and Tradition SEF (73).

For each of the index tranche swap transactions, we first look for a simultaneously executed index CDS transaction on the same SEF which references the same index as the tranche swap and which has the same maturity as the latter. This results in only 243 matching delta exchanges, none of which takes place on ICAP SEF, i.e., the SEF that accounts for almost half of the index tranche swap transactions. This suggests that not all SEFs have the ability to simultaneously execute, confirm, and/or report index tranche swaps and the respective delta exchanges. Therefore, we hand match to most of the remaining index tranche swap transactions an index CDS transaction executed on the same trading day and SEF which references the same index as the tranche swap and which has the same maturity as the latter. Frequently, the match is unique because the index tranche swap references a far off-the-run index series that trades infrequently. However, in order to deal with non-unique matches, we resort to Credit Market Analysis (CMA) intraday tranche swap quotes from the same trading day on which the index tranche swap was executed, and, in addition, require that index CDS and tranche swap notional amounts are consistent with some delta that was quoted on the trading day (i.e., the index CDS's notional amount is approximately equal to delta times the tranche swap's notional amount) and that the transaction price of the index CDS equals the quote's reference level. This results in an additional 297 delta exchanges. The remaining index tranche swaps are assumed to be traded without delta. Since tranche swaps without delta usually trade at less favourable prices, investors may find it beneficial to trade index tranche swaps with delta and unwind the delta exchange themselves. We identify such delta offsets as

²³DDR trade reports of index tranche swaps executed on GFI Swaps Exchange have "UNDERLYING_ASSET_1" field values "Dow Jones CDX Investment Grade[Series]..." or "Dow Jones CDX HY([Series])..." or "DJ CDX IG[Series]..." ([Series] denotes the index's series number). Those of index tranche swaps executed on ICAP SEF have "PRICE_NOTATION_TYPE" field value "Percentage" and those of index tranche swaps executed on ICE Swap Trade have "PRICE_NOTATION_TYPE" field values "Price" or "UpfrontPoints". Trade reports of index tranche swaps executed on Tradition SEF are hand-matched on dates on which the SEF reported index tranche swap trading (according to Clarus FT data). The four SEFs are the only SEFs with index tranche swap trading activity during our sample period according to Clarus FT data.

Appendix B. Appendix to Chapter 2

transactions with the same transaction price and notional amount as a delta exchange of an index tranche swap which occur on the same trading day and SEF.

We assign to each of the identified transactions a index tranche swap delta exchange trade type and a more detailed trade type specifying the tranche swap's attachment and detachment points, the reference level, and the delta.²⁴

Delta Exchanges of Index Swaptions

In order to identify delta exchanges of index swaptions, we first apply a reporting-format-based SEF identification algorithm to on-SEF index swaption transactions disseminated by the DDR.²⁵ On-SEF index swaption transactions disseminated by the ICETV are identified as being executed on ICE Swap Trade. There are a total of 1,250 on-SEF index swaption transactions and for 1,169 of those we are able to identify the SEF on which the transaction was executed.²⁶

²⁴Attachment and detachment points are usually not part of the trade report but they can be inferred from the tranche swap quote used to identify the delta hedge.

²⁵DDR trade reports of index swaptions executed on Bloomberg SEF have "UNDERLYING_ASSET_1" field values "MARKIT CDX.NA.IG.[Series] mm/yy:..." or "MARKIT CDX.NA.HY.[Series] mm/yy:..." and (case-sensitive) "PRICE_NOTATION_TYPE" field value "Basis points". Bloomberg SEF index swaptions are cleared (see the respective contract specifications on Bloomberg SEF's website). Trade reports of index swaptions executed on GFI Swaps Exchange either have "UNDERLYING_ASSET_1" field values "CDX.NA.IG.[Series]v[Version]...", "CDX.NA.IG.[Series]:...", "CDX.NA.HY.[Series]v[Version]...", or "CDX.NA.HY.[Series]:..." and (case-sensitive) "PRICE_NOTATION_TYPE" field values "Basis points" or "Price" and effective dates after trade execution dates, or "UNDERLYING_ASSET_1" field values "IG..." or "HY...", or "UNDERLYING_ASSET_1" field values that contain the word "Swaption", or "EMBEDDED_OPTION" field value "EMBED1". Trade reports of index swaptions executed on ICAP SEF have "UNDERLYING_ASSET_1" field value "...:Bespoke Basket", or "UNDERLYING_ASSET_1" field values "CDX.NA.IG.[Series]:..." or "CDX.NA.HY.[Series]:..." and (case-sensitive) "PRICE_NOTATION_TYPE" field value "Percentage" and "DAY_COUNT_CONVENTION" field value "30/360" and effective dates after trade execution dates. Trade reports of index swaptions executed on ICE Swap Trade have "UNDERLYING_ASSET_1" field values that contain the words "Call", "Payer", "Put", or "Receiver", or "UNDERLYING_ASSET_1" field values "CDX.NA.IG.[Series]:..." or "CDX.NA.HY.[Series]:..." and (case-sensitive) "PRICE_NOTATION_TYPE" field value "BasisPoints", or "UNDERLYING_ASSET_1" field values "MARKIT CDX.NA.IG.[Series] mm/yy:..." or "MARKIT CDX.NA.HY.[Series] mm/yy:..." and (case-sensitive) "PRICE_NOTATION_TYPE" field value "Price", or "UNDERLYING_ASSET_1" field values "CDX.NA.IG.[Series]:..." or "CDX.NA.HY.[Series]:..." and (case-sensitive) "PRICE_NOTATION_TYPE" field value "Percentage" and "DAY_COUNT_CONVENTION" field value "ACT/360" and effective dates at or before trade execution dates. Trade reports of index swaptions executed on MarketAxess SEF have "UNDERLYING_ASSET_1" field value "[Six-digit RED code]:...". Trade reports of index swaptions executed on tpSEF have "UNDERLYING_ASSET_1" field values "CDX.NA.IG.[Series]:..." or "CDX.NA.HY.[Series]:..." and (case-sensitive) "PRICE_NOTATION_TYPE" field value "Percentage" and "DAY_COUNT_CONVENTION" field value "30/360" and effective dates at or before trade execution dates. Trade reports of index swaptions executed on Tradition SEF have "UNDERLYING_ASSET_1" field values "CDX.NA.IG.[Series]:..." or "CDX.NA.HY.[Series]:..." and (case-sensitive) "PRICE_NOTATION_TYPE" field value "Percentage" and "DAY_COUNT_CONVENTION" field value "ACT/360" and effective dates after trade execution dates. According to Clarus FT data all of the above SEFs but ICAP SEF have index swaption trading activity during our sample period. The reason for ICAP SEF not showing up in the Clarus FT data is the fact that, in many cases, the SEF does not seem to declare index swaption (as well as index tranche swap) trading activity in an explicit manner (see Section B.3 for details). In addition, the Clarus FT data shows index swaption trading activity by BGC Derivative Markets which we neglect because the SEF does not seem to report the notional amount of delta exchanges as index CDS trading activity.

²⁶Among the remaining 81 on-SEF index swaption transactions are 65 transactions that we would identify as being executed on BGC Derivative Market and nine transactions that we would identify as being executed on Tradition SEF but for which Clarus FT data does not show trading activity.

Most of the index swaption transactions are executed on GFI Swaps Exchange (383), ICE Swap Trade (368), and Tradition SEF (208). Index swaptions are relatively rarely executed on the remaining four SEFs; Bloomberg SEF (14), ICAP SEF (81), MarketAxess SEF (45), and tpSEF (70).

For each of the index swaption transactions, we first look for a simultaneously executed transaction in the underlying index CDS on the same SEF. This results in only 321 matching delta exchanges. Therefore, we hand match to many of the remaining index swaption transactions a transaction in the underlying index CDS that is executed on the same trading day and SEF. The matches have transaction prices at which reference levels tend to be set and non-round-lot notional amounts. This results in an additional 640 delta exchanges. The remaining index swaptions are assumed to be traded without delta. As in case of delta exchanges of index tranche swaps, we identify delta offsets as transactions with the same transaction price and notional amount as a delta exchange of an index swaption which occur on the same trading day and SEF.

We assign to each of identified transactions an index swaption delta exchange trade type and a more detailed trade type specifying the swaption's underlying, type (payer or receiver swaption), expiry, and strike price.

Index Rolls

We identify index rolls as simultaneously executed index CDS transactions on the same SEF that reference two different series of the same index and that have the same contract tenor. In addition, we require that the ratio of notional amounts on the two legs of an index roll (i.e., the notional amount of the index CDS referencing the older index series divided by the notional amount of the index CDS referencing the newer index series) is within certain bounds.²⁷ In case of non-unique matches, we hand match trades such that disparity between the notional amounts on both legs of the trade is minimal.

The bounds serve to prevent index roll identification from transactions that have been simultaneously executed by chance. Broadly speaking, an index roll may serve one of two purposes: either maintaining an exposure of a given notional amount in a more current index series or getting exposure to a more current index series at a minimal outlay. In the former case, the notional amount of the index CDS referencing the newer index series will be equal to that of the index CDS referencing the older index series. In the latter case, the notional amount of the index CDS referencing the newer index series will typically be smaller than that of the index CDS referencing the older index series because the risky duration (or risky present value of

²⁷The bounds are 1 and 1.2 for index rolls from the immediate off-the-run series to the on-the-run series, and 1 and 4.1667 for all other index rolls. The latter bound is calibrated to volume data reported by GFI Swaps Exchange (from Clarus FT). For some part of our sample period, the SEF reported index roll volumes separately. During this period, the index roll with most displaced index series is a "HY14/HY22 Roll" on April 30, 2014. In the transaction data, a USD 50 million notional amount in five-year CDX.HY.14 is rolled into a USD 12 million notional amount in five-year CDX.HY.22. This determines the bound in that $4.1667 = 50/12$.

a basis point) of the former index CDS is larger than that of latter index CDS. Therefore, we set the lower bound for index roll identification to unity. For index rolls from the immediate off-the-run series to the on-the-run series the upper bound for index roll identification is set to 1.2 because simplified risky duration considerations suggest that the notional amount ratio of zero-outlay index rolls is of the order 1.18 ($\approx 3.25/2.75$) or below for contract tenors of three or more years.²⁸

We assign to each of these transactions an index roll trade type and a more detailed trade type specifying the index CDSs involved in these trades.

Curve Trades

We identify curve trades as simultaneously executed index CDS transactions on the same SEF that have different contract tenors and that reference the same index (but not necessarily the same series of an index). In case of non-unique matches, we hand match trades such that disparity between the notional amounts on both legs of the trade is minimal.

We assign to each of these transactions a curve trade trade type and a more detailed trade type specifying the index CDSs involved in these trades.

B.2.5 Trade Size Aggregation

Given the detailed trade types, we aggregate trade sizes of simultaneously executed non-block transactions in the same index CDS contract and on the same SEF that have the same detailed trade type and the same transaction price. Simultaneously executed transactions may occur in case that a trade aggressor hits or lifts limit orders that different market participants placed on the order book.²⁹ For instance, the best bid on the order book may show USD 50 million depth and, in fact, be composed of two limit buy orders for USD 25 million. When the full depth is hit, two simultaneous transactions for USD 25 million occur although the aggressor traded USD 50 million. Alternatively, simultaneously executed transactions may be duplicates (in case that they have equal trade size as in the previous example) or they may occur in case that two RFQ initiators execute an identical response at the same time. Both scenarios seems less likely to us than the aforementioned order book execution that justifies trade size aggregation. Of course, trade size aggregation fails to account for the fact that an aggressor's order may "walk the book" because we require transaction price to be the same. This is likely to be a minor issue because SEF personnel told us that order books are usually shallow with depth concentrating at the best bid and offer.

²⁸Note that index CDSs have a maturity that is one quarter longer than the contract tenor on the index launch date (i.e., the three-year tenor has approximately 3.25 years to maturity on the launch date) and that indices are launched every six months (i.e., the three-year tenor of an index CDS referencing the immediate off-the-run index has approximately 2.75 years to maturity on the launch date).

²⁹As described in Section B.2.3, an otherwise block-eligible transaction that is executed on an order book does not receive block treatment. Therefore, the situation described here may not explain simultaneous execution of two or more block trades, and this is why we require transactions to be non-block when aggregating trade sizes.

B.2.6 Cleaning Transaction Prices

The last stage of data processing assesses whether index CDS contract terms are sufficiently standardized and whether the transaction data and, in particular, the contained pricing information, are accurate. Because we only use transaction prices of trades in five-year on-the-run and immediate off-the-run index CDSs, this stage is only applied to trades in the latter contracts.

To this end, we remove all trade reports with zero transaction prices, all trade reports with 00:00:00 DST (New York daylight saving time) timestamps, and all trade reports with non-standard maturity dates.³⁰ We also remove all trade reports of outright transactions and delta exchanges with fixed-spread transaction prices because of likely reporting errors (index CDS spreads tend to be different from fixed spreads).³¹

Conventionally, the prices of index CDSs referencing CDX.IG indices are expressed in terms of index CDS spreads (in basis points), while those of index CDSs referencing CDX.HY indices are expressed in terms of prices (in percent). SDR trade reports indicate whether transaction prices are expressed in terms of a price or in terms of an index CDS spread,³² but the indications are frequently erroneous or systematically wrong. For instance, almost all DDR trade reports of transactions in CDX.HY index CDSs that were executed on Bloomberg SEF indicate that transaction prices are expressed in terms of spreads (the reported price notation type is “Basis points”) although the reported transaction prices are expressed in terms of prices. In order to account for erroneous indications, we overwrite indications with SEF-specific price types based on our experience with the trade report history. In our experience, all trade reports of CDX.IG transactions other than those executed on MarketAxess SEF prior to March 6, 2014 contain transaction prices that are expressed in terms of spreads.³³ Similarly, all trade reports of CDX.HY transactions other than those executed on ICE Swap Trade between February 21, 2014 and August 1, 2014 contain transaction prices that are expressed in terms of prices. Transaction prices contained in the MarketAxess SEF trade reports seem to be expressed in terms of prices, while transaction prices contained in the ICE Swap Trade trade reports seem to be expressed in terms of index CDS spreads.

Moreover, we divide transaction prices contained in all trade reports of transactions executed on MarketAxess SEF prior to March 6, 2014 by 100. This is because they seem to be expressed in basis points rather than in percent of the notional amount. For the same reason, we divide transaction prices contained in trade reports of transactions executed on MarketAxess SEF on or after March 6, 2014 by 100 if they indicate “Percentage” price notation type.

In order to ensure that overwriting indications does not introduce further errors, we compare

³⁰Standardized index CDSs that were launched in March (September) mature on the 20th of June (December) of the year that follows the index launch by the number of years specified through the contract tenor.

³¹We exclude index rolls and curve trades because they are priced in relative terms.

³²The corresponding field names of BSDR, DDR, and ICETV trade reports are “price notation type”, “PRICE_NOTATION_TYPE”, and “PriceType”, respectively.

³³Here and in what follows, dates refer to DST calendar dates.

Appendix B. Appendix to Chapter 2

transaction prices with Markit's end-of-day composites. To this end, we remove all trade reports of transactions without available Markit end-of-day composite prices and spreads on the date of trade execution (this, amongst others, removes transactions that were executed on weekends) and process the remaining trade reports through the following filter:

1. In case that a transaction price which is expressed in terms of a spread deviates by more than 5% from Markit's end-of-day composite spread, we first check whether there is a scaling factor such that the scaled transaction price does not deviate by more than 5% from the end-of-day composite.³⁴ If this is the case, we replace the transaction price by the scaled transaction price with minimum percentage deviation from the end-of-day composite spread.
2. In case that a transaction price which is expressed in terms of a price deviates by more than 1% from Markit's end-of-day composite price, we first check whether there is a scaling factor such that the scaled transaction price does not deviate by more than 1% from the end-of-day composite. If this is the case, we replace the transaction price by the scaled transaction price with minimum percentage deviation from the end-of-day composite price.
3. In case that a transaction price which is expressed in terms of a spread continues to deviate by more than 5% from Markit's end-of-day composite spread, we next check whether there is a scaling factor such that the scaled transaction price does not deviate by more than 1% from Markit's end-of-day composite price.³⁵ If this is the case, we overwrite the indication such that it indicates a transaction price which is expressed in terms of a price and replace the transaction price by the eventually scaled transaction price with minimum percentage deviation from the end-of-day composite price.
4. In case that a transaction price which is expressed in terms of a price continues to deviate by more than 1% from Markit's end-of-day composite price, we next check whether there is a scaling factor such that the scaled transaction price does not deviate by more than 5% from Markit's end-of-day composite spread. If this is the case, we overwrite the indication such that it indicates a transaction price which is expressed in terms of a spread and replace the transaction price by the eventually scaled transaction price with minimum percentage deviation from the end-of-day composite spread.

Then, we convert transaction prices which are expressed in terms of a price into equivalent expressions in terms of index CDS spreads, and vice versa.³⁶ Once this is done, we detect outliers by transaction prices that deviate by more than 3% from the intraday mid-quote that prevails at trade execution. Both transaction prices and quotes are in terms of index CDS

³⁴The scaling factor may take one of the four values that would be used when converting transaction prices into decimals, percentages, or basis points: 1/1000, 1/100, 100, and 10000.

³⁵In addition to the four above mention values, the scaling factor may be 1 (no scaling).

³⁶When converting between expressions, we use the ISDA CDS Standard Model which is the industry standard.

spreads and, apart from the robustness check in Section B.7, the intraday mid-quote comes from Markit. We remove all trade reports of transactions with outlier transaction prices and, in case of package transactions with outlier transaction prices, we also remove the trade report of the other leg of the package. Finally, we remove all trade reports of transactions executed on Securities Industry and Financial Markets Association (SIFMA) recommended close and recommended early close days, and all all trade reports indicating bespoke terms.

B.2.7 Inference of Capped Trade Sizes in case of ICETV Trade Reports

ICETV trade reports do not indicate whether the contained notional amounts are capped or not but the rules-based approach of CFTC regulations allows to infer capped notional amounts. It should, however, be noted that exact inference is not possible because trade reports contain rounded notional amounts.

CFTC rules define the cap size of an index CDS transaction as the maximum of the appropriate minimum block size and USD 100 million.³⁷ The appropriate minimum block size is transaction-specific in that it depends on the swap category to which the index CDS transaction belongs. The latter is uniquely determined by the transaction's index CDS spread and the contract tenor.³⁸ There are three index CDS spread categories with cutoffs at 175 bps and 350 bps and six tenor categories with cutoffs at two, four, six, eight and a half, and twelve and a half years. Appendix F to Part 43 of Chapter I of 17 CFR contains the initial appropriate minimum block sizes of the resulting 18 swap categories. Given the above definition, cap sizes of all swap categories can be easily deduced from the appendix by the maximum of USD 100 million and the applicable minimum block size contained in the appendix.

Determining swap categories requires transaction price in terms of index CDS spreads. Therefore, the transaction prices of all ICETV trade reports are processed in the same way as transactions in five-year on-the-run and immediate off-the-run index CDSs (the previous section contains the details) but only up to the point where transaction prices are converted. Then capped notional amounts are inferred by a notional amount at or above the cap size applicable to the swap category that is implied by the index CDS spread and the contract tenor of the transaction. There are a few transaction with zero and fixed-spread transaction prices for which inference is based on the previous end-of-day composite spread from Markit rather than the transaction price. This approach appears to be plausible in light of cap sizes that depend on minimum applicable block sizes and industry practice to use previous end-of-day spreads in order to determine block treatment.³⁹

³⁷See §43.4(h) of Chapter I of 17 CFR.

³⁸Unfortunately, the term tenor is not defined in CFTC rules but the following footnote of a Federal Register publication by the CFTC (see 77 Federal Register 15468 (March 15, 2012) at note 101) suggests that tenor refers to days to maturity: "the tenor of a swap refers to the amount of time from the effective or start date of a swap to the end date of such swap. In circumstances where the effective or start date of the swap was different from the trade date of the swap, the Commission used the later occurring of the two dates to determine tenor." Analogously, we define tenor as the date difference between the date on which the transaction was executed and the maturity date.

³⁹SEFs use previous end-of-day spreads in order to determine block treatment because different methods of

B.3 Additional Data Sources

This section briefly describes other data sources that our analyses rely upon and eventual data processing.

B.3.1 Clarus FT SEF Volumes

We obtain data for USD-denominated index CDS contracts referencing CDX.IG and CDX.HY indices from Clarus FT's SEF View. The view compiles data from trading activity reports that SEFs file on a daily basis on their websites.

Data Description

The Clarus FT data contains 26 fields including the SEF, the reporting date, an credit index identifier, the contract tenor, the notional amount traded, and the block and non-block notional amounts traded. In addition, the data contains three "Markup" fields whose content seems to differ for each SEF. The "Markup1" field seems to contain the index CDS contract description that the respective SEF used when filing the report and is populated for all SEFs. For most SEFs this is the only "Markup" field populated. However, for some SEFs the "Markup2" and "Markup3" fields are populated as well. Those include ICE (i.e., ICE Swap Trade), TP (i.e., tpSEF), and Tradition (i.e., Tradition SEF). For ICE the "Markup3" field contains the derivatives clearing organization (the field is populated with "ICE CLEAR CREDIT", "ICE CLEAR EUROPE", or "None"). This allows us to determine the dates on which ICE Swap Trade data is disseminated by the DDR (remember that ICE Swap Trade reports uncleared index CDSs transactions, i.e., those that make up SEF-reported entries with "Markup3" field values "None", to the DDR). For TP (Tradition) the "Markup3" ("Markup2") field contains the trade count.

Data Cleaning

The "Markup1" field enables us to detect erroneous entries. Specifically, the Clarus FT data are cleaned by removing (i.) duplicate entries; (ii.) entries for swaption and tranche swap contracts (i.e., entries for which the "Markup1" field contains "Call", "Payer", "Put", "Receiver", "Tranche", or "%"); (iii.) ICAP (i.e., ICAP SEF) entries with "Markup1" field values that contain six-digit RED codes instead of nine-digit RED codes (these entries seem to be for index swaptions and index tranche swaps because their prices are very different from the ones of the index CDSs on the credit indices displayed in the "Security" field); (iv.) ICAP (i.e., ICAP SEF) entries with "Markup1" field value "5" (these entries seem to be for swaptions because their prices are very different from the ones of the index CDSs on the credit indices displayed in the "Security" field); (v.) ICE (i.e., ICE Swap Trade) entries with "Markup1" fields values that start with "Markit"

trade execution are available for block and non-block trades and applicability of the method of trade execution has to be verified before trade.

and “Markup3” fields values that are “None” (these entries seem to be for non-cleared index tranche swaps); (vi.) entries with “Markup1” field values that contain “2I65BZ” (these entries seem to be for index CDSs on CDX.EM); (vii.) entries with reporting dates prior to the launch date of the index series; and (viii.) three ICAP (i.e., ICAP SEF) entries that have reporting dates after the maturity of the index CDS contract.

B.3.2 Credit Market Analysis Intraday Quotes

We obtain a custom data set of intraday index CDS quotes from CMA. Alike Markit, CMA provides dealer-run-based intraday composite quotes for index and single-name CDSs but there are important methodological differences between the two data sources. The most important of which is knowledge of the quote setter’s identity: when forming a composite quote, Markit is aware of the dealer that initiated the run while CMA is not.⁴⁰

When forming a composite quote, CMA only has access to contributions from a data-sharing consortium of non-dealer market participants. Consortium members use CMA’s quote parsing software (CMA Quotevision) in order to manage the large number of quotes that is communicated to them via dealer runs.⁴¹ Specifically, the software provides each consortium member with a real-time structured overview of the quotes she has received by mail. The overview is specific to each consortium member and only available to her, but for a given index CDS contract summary statistics (such as, the median mid-quote and the median bid-ask spread) of the quotes contained in the overview are contributed to CMA whenever the overview for the particular contract is updated by receipt of a new run.⁴² CMA’s composite mid-quote is a robust (median-like) statistic of the summary statistics contributed by the consortium members.⁴³ Instead of CMA’s composite mid-quote, our custom data set comprises bid and ask quotes based on the average mid-quote and the average bid-ask spread across all contributions received by CMA within a given second. In addition, the data include the number of contributions underlying the computation of averages. We only use averages with at least two underlying contributions.

B.3.3 GFI Market Data

We obtain GFI Market Data from Fenics. GFI is a leading interdealer broker in both cash and

⁴⁰Quotes are the intellectual property of the quote setter and cannot be shared with third persons without the quote setter’s agreement.

⁴¹According to CMA, some of the consortium members receive up to 20,000 quotes via dealer runs per day.

⁴²For a contribution to be made, the overview has to contain quotes of at least two different dealers and satisfy additional proprietary criteria.

⁴³For inactively quoted CDSs, the procedure can give rise to composites that actually coincide with an individual dealer’s quotes. In order to preserve quote anonymity, CMA adds a small random quantity to composite mid-quotes. Unfortunately, the random quantities added to CDX.IG composites turned out to be too large to allow for meaningful inference of trade direction or precise estimation of transaction costs. For CDX.IG and CDX.HY, quote randomization started on January 1, 2014 and was suspended on November 19, 2014 because of sufficiently active quoting by dealers. The custom data are not randomized.

derivative fixed income markets and the operator of the GFI Swaps Exchange SEF. The data come from GFI's CreditMatch trading platform which is also used by its SEF. The data comprise the best bids and offers resting on GFI order books as well as prices at which quantity can be exchanged during designated matching sessions. Matching sessions are either periodic and time-limited or continuous and open-ended. In case of periodic matching sessions, prices can be determined in a variety of ways including price fixing sessions that turn into matching sessions once prices are fixed. In case of continuous matching sessions, prices are broker-determined mid-points.

We remove a few obvious outliers from the data for five-year on-the-run index CDSs on CDX.IG and CDX.HY. In order to determine the beginning and end of continuous mid-point matching sessions in these contracts, we remove mid-point repetitions in case that there are consecutive mid-points at the same level. According to GFI representatives, the repetitions are due to the data collection procedure which records all non-trade events happening during existing matching sessions (such as, the bids and offers made) in the same way as matching session prices. In order to compute profits from liquidity provision, we convert quotes that are expressed in terms of index CDS spreads into quotes that are expressed in terms of prices.

B.3.4 Markit Index Swaptions

The index swaption data come from Markit and comprise end-of-day composite bid and ask prices and implied volatilities. The composites are formed at 6:30 p.m. New York time and based on a collection of individual dealer quotes. Markit parses quotes from dealer runs throughout the trading day, and the collection on which composite computation is based contains the quotes from each dealer's latest run. In addition to composites, the data also comprise the number of dealers whose quotes are used in the composite computation and the number of quotes parsed over the previous 24 hours. Because dealers may use non-identical reference levels when quoting swaption contracts (cross-sectionally and throughout the trading day), composites are formed per swaption contract and reference level.

We use three-month at-the-money implied index swaption volatility as a control variable for market conditions in trade-by-trade regressions that estimate selection-bias-corrected average effective half-spreads, realized half-spreads, and price impacts. The at-the-money swaption has a strike closest to the index's end-of-day composite. To select among swaptions with the same strike and different reference levels, we choose the swaption that is most actively quoted, and that has the largest number of quoting dealers and the tightest average bid-ask spread across payer and receiver swaptions. In case that these criteria do not result in a unique match, we select the swaption with reference level closest to the end-of-day composite.

B.3.5 Markit Intraday Quotes

Markit intraday data comprise bid and ask quotes in terms of both prices and index CDS spreads.

B.4. Trading Protocol Identification for GFI Swaps Exchange Trades

We first remove duplicate entries from the data. Then, we remove quotes for already-matured index CDS and quotes that give rise to negative bid-ask spreads.⁴⁴ Finally, we compress the information of quotes which are made in the same second (quotes are time stamped with second precision) as some other quote on the same index CDS into single quotes such that in each second there is at most one quote on a particular index CDS.⁴⁵ Occasionally, there are short periods of time in which mid-quotes are more volatile than transaction prices. We ignore quotes from these periods when detecting outliers and in all our analyses.

B.4 Trading Protocol Identification for GFI Swaps Exchange Trades

We use GFI Market Data to identify the trading protocols that were used to execute trades on the GFI Swaps Exchange. First, we identify order-book trades by transaction prices that coincide with either the best bid or offer that prevails on the order book of the GFI Swaps Exchange at trade execution.

Then, we identify workup sessions and workups. In identifying workup sessions, we closely follow the description of the workup protocol that is contained in the rulebook of the GFI Swaps Exchange, in particular, with respect to the 40-second duration of workup sessions and the fact that only trades that occur on the order book trigger workup sessions.⁴⁶ Specifically, we sequence all transactions in a given index CDS contract that are of the same detailed trade type.⁴⁷ Whenever two consecutive transactions in such a sequence occur within 40 seconds and at the same price, we infer that the second transaction is part of a workup session that was triggered by the execution of the first one, the so-called workup trigger, provided that the workup trigger is an order-book trade. We assume that the workup session times out 40 seconds after it was triggered and that any transaction in the sequence which occurs before the timeout at a price other than the workup trigger's transaction price terminates the session prematurely. All transactions that occur at the workup trigger's transaction price before the session terminates are identified as workups (for transactions that were previously identified as order-book trades the identified trading protocol is overwritten).

⁴⁴Note that only the quotes in terms of the primary price type, i.e., the price type used in dealer runs, can be expected to be non-negative. This is because, when converting primary price type quotes into secondary price types, Markit does not take into account the inverse relation between prices and index CDS spreads. For instance, the primary price type for CDX.IG index CDSs is the index CDS spread. In this case, bid (ask) quotes in terms of price (i.e., the secondary price type) correspond to converted bid (ask) quotes in terms of the index CDS spread. When the bid quote in terms of the index CDS spread is below the ask quote, then the price that is obtained by conversion of the bid quote is above the price that is obtained by conversion of the ask quote due to the inverse relation between prices and index CDS spreads.

⁴⁵This compression is achieved by taking the quote pair at the 50th percentile of the mid-quote distribution. In case that there is more than one pair of quotes with mid-quote equal to the 50th percentile, the pair at the 50th percentile of the bid-ask spread distribution is taken.

⁴⁶The rulebook leaves open whether or not operation of the order book is suspended during the workup session. Our identification assumes that the order book continues operating while the session is in progress and that the session will be immediately terminated by any transaction at a price other than the one at which trade size is worked up.

⁴⁷We sequence transactions first by their execution timestamps, then by the disseminating SDR, and finally by the numeric part of the SDR's dissemination identifier.

Finally, we identify mid-market matches as trades with previously unidentified trading protocols whose transaction prices coincide with the mid-market level that prevails at trade execution on the GFI Swaps Exchange.

B.5 Trade Size Weighting

In the paper, we choose to focus on sample means because capped trade sizes may render trade-size-weighted averages subject to biases. However, trade-size-weighted average effective half-spreads may be more representative of the actual cost of trading, for instance, because the weighting scheme already accounts for the fact that transaction costs increase with trade size. Therefore, Table B.2 displays trade-size-weighted effective half-spreads, realized half-spreads, and price impacts of outright trades in five-year on-the-run index CDSs and, for comparison, recapitulates the estimates reported in Table 2.3 of the paper in the first row of each panel. Consistent with effective half-spreads that increase with trade size, trade-size-weighted averages are larger than sample means and so are differences of trade-size-weighted average effective half-spreads of D2C and dealer-to-dealer (D2D) trades. Therefore, the differences between D2C and D2D transaction costs that we report in the paper are conservative. Similarly, the differences between D2C and D2D price impacts that we report in the paper are conservative as well.

In order to mitigate eventual biases of trade-size-weighted averages, we proceed as in Section 2.3.4 of the paper and determine the average size by which trades in five-year on-the-run index CDSs exceed cap sizes using SEF-reported volumes from Clarus FT. We then add these averages on top of the disseminated cap sizes in order to obtain cap-adjusted weights for computing trade-size-weighted averages. Similar to the average trade sizes in excess of caps that we report in the paper (which are based on the broader sample of all on-SEF index CDS transactions in CDX.IG and CDX.HY, respectively), we find that D2C trades in five-year on-the-run index CDSs on CDX.IG exceed caps on average by USD 123.21 million USD. The size by which D2D trades in these index CDSs exceed caps is slightly larger than the one reported in the paper, USD 132.53 million. The corresponding averages for D2C and D2D trades in five-year on-the-run index CDSs on CDX.HY are USD 107.33 million and USD 142.28 million, respectively.

The resulting cap-adjusted trade-size-weighted averages are shown in the last row of each panel. Cap-adjusted trade-size-weighted average effective half-spreads tend to be slightly larger than unadjusted averages while the opposite seems to be the case for trade-size-weighted price impacts. This indicates that capped trade sizes induce a small downward bias in trade-size-weighted effective half-spreads and a small upward bias in trade-size-weighted price impacts. Inference regarding differential effective half-spreads and price impacts of D2C and D2D trades seems to be unaffected by the bias.

B.6. Outright Immediate Off-The-Run Index CDS Trades

	<i>Dealer-To-Customer</i>			<i>Dealer-To-Dealer</i>			<i>D2C-D2D</i>		
	<i>Effc Sprd</i>	<i>Rlzd Sprd</i>	<i>Prc Imp</i>	<i>Effc Sprd</i>	<i>Rlzd Sprd</i>	<i>Prc Imp</i>	<i>Effc Sprd</i>	<i>Rlzd Sprd</i>	<i>Prc Imp</i>
Panel A: CDX.IG									
Sample mean	0.137	0.031	0.106	0.088	0.025	0.063	0.049**	0.006	0.043**
Weighted Avg.	0.149	0.033	0.116	0.092	0.034	0.058	0.057**	-0.001	0.058**
Weighted (adj.)	0.156	0.039	0.117	0.094	0.040	0.053	0.062**	-0.002	0.064**
Panel B: CDX.HY									
Sample mean	0.674	0.166	0.508	0.402	0.155	0.246	0.273**	0.011	0.262**
Weighted Avg.	0.763	0.250	0.513	0.410	0.187	0.223	0.353**	0.062**	0.290**
Weighted (adj.)	0.795	0.302	0.493	0.414	0.199	0.215	0.380**	0.102**	0.278**

Table B.2: Effective Half-Spreads, Realized Half-Spreads, and Price Impacts.

Panels A and B show mean and trade-size-weighted averages of effective half-spreads (EfficSprd), realized half-spreads (RlzdSprd), and price impacts (PrcImp) of outright dealer-to-customer (D2C) and dealer-to-dealer (D2D) trades in five-year on-the-run index CDSs on CDX.IG and CDX.HY, respectively. In the computation of the adjusted trade-size-weighted average, the size of trades that are disseminated with capped notional amounts is adjusted by an index- and D2C-(D2D)-specific mean excess-beyond-cap estimate based on SEF-reported actual trading volumes from Clarus FT. EfficSprd 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 index CDS 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.

B.6 Outright Immediate Off-The-Run Index CDS Trades

Table B.3 displays average effective half-spreads, realized half-spreads, and price impacts of outright trades in five-year immediate off-the-run index CDSs. Transaction costs of trades in immediate off-the-run index CDSs are larger than those of trades in on-the-run index CDSs which is consistent with off-the-run index CDSs being less liquidly traded than on-the-run index CDSs. In contrast, the price impact of trades in immediate off-the-run index CDSs is smaller than the one of trades in on-the-run index CDSs. This most likely reflects the fact that many of the trades in off-the-run index CDSs are liquidity motivated in that they close existing positions. Consistent with our results for outright trades in on-the-run index CDSs, we find larger transaction costs and higher price impacts of D2C trades.

Table B.4 displays regression results for outright trades in five-year immediate off-the-run index CDSs.⁴⁸ Results are broadly consistent with those for outright trades in on-the-run

⁴⁸The mismatch in the number of trades between Table B.3 and Table B.4 is due to unavailable quotes on the five-year on-the-run index (i.e., trades for which the BAS and SPRD explanatory variables cannot be computed).

Appendix B. Appendix to Chapter 2

Type	<i>Dealer-To-Customer</i>			<i>Dealer-To-Dealer</i>			<i>D2C-D2D</i>		
	Effc Sprd	Rlzd Sprd	Prc Imp	Effc Sprd	Rlzd Sprd	Prc Imp	Effc Sprd	Rlzd Sprd	Prc Imp
Panel A: CDX.IG									
Outright	0.191	0.118	0.073	0.103	0.118	-0.016	0.088**	-0.001	0.089**
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	1.056	0.650	0.407	0.590	0.432	0.159	0.466**	0.218*	0.248*
Index roll	0.392	0.239	0.153	0.354	0.131	0.223	0.038	0.108*	-0.070

Table B.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-customer (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 immediate off-the-run index CDSs, for index rolls between five-year on-the-run and immediate off-the-run index CDSs, and for delta exchanges of index swaption and index tranche swap trades that reference the five-year immediate off-the-run index. 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 index CDS 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 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 2,861 (85) and 3,875 (149) outright D2C (D2D) trades in five-year immediate off-the-run index CDSs on CDX.IG and CDX.HY, respectively, and 968 (344) and 1,283 (343) D2C (D2D) index rolls between five-year on-the-run and immediate off-the-run index CDSs on CDX.IG and CDX.HY, respectively.

index CDSs but less strong. Transaction costs tend to increase with trade size and quoted bid-ask spread (as well as implied volatility in case of CDX.HY) and the price impact of trades increases when bid-ask spreads widen and liquidity deteriorates. Transaction costs and price impacts of D2C trades are significantly higher than those of D2D trades, even after accounting for trade characteristics and market conditions.⁴⁹

B.7 Robustness Checks

This section contains the results of a variety of robustness checks that we conducted. These include using an alternative mid-quote when computing and decomposing transaction costs, using alternative time frames over which to compute realized half-spreads and price impacts,

Most of these trades come from index roll dates and are executed before the five-year on-the-run index is quoted for the first time.

⁴⁹Significance is marginal in case of the price impact regression for CDX.HY but the difference in price impacts of D2C and D2D trades is of the same order of magnitude as the one for outright trades in five-year on-the-run index CDSs (see Table 2.5 of the paper).

B.7. Robustness Checks

	CDX.IG			CDX.HY		
	EffcSprd	RlzdSprd	PrcImp	EffcSprd	RlzdSprd	PrcImp
D2C	0.161** (21.82)	0.105** (12.13)	0.056** (6.54)	0.909** (27.30)	0.533** (11.72)	0.375** (8.85)
D2D	0.092** (3.85)	0.114** (3.58)	-0.022 (-1.10)	0.511** (5.74)	0.339* (2.25)	0.171 (1.49)
MDM	0.016 (1.65)	-0.030 (-1.82)	0.046** (2.69)	0.033 (0.97)	-0.051 (-0.74)	0.084 (1.38)
LRG	0.009 (0.88)	-0.004 (-0.30)	0.014 (0.89)	0.099* (2.21)	0.045 (0.60)	0.054 (0.79)
BLCK	0.065** (4.54)	0.060** (2.89)	0.004 (0.36)	0.282** (4.39)	0.293** (3.40)	-0.012 (-0.19)
RFRNC	0.056 (1.95)	0.076 (1.94)	-0.020 (-1.18)	0.603** (3.99)	0.596** (3.63)	0.007 (0.09)
BAS	0.785** (4.47)	0.209 (1.37)	0.576* (2.30)	0.339** (2.68)	-0.062 (-0.41)	0.401** (3.87)
SPRD/100	0.156 (1.52)	0.157 (1.37)	-0.001 (-0.01)	-0.082 (-1.25)	-0.004 (-0.04)	-0.079 (-1.02)
VLTLTY	0.102 (0.99)	-0.027 (-0.23)	0.129 (0.97)	1.693** (3.41)	0.942 (1.31)	0.752 (1.26)
<i>N</i>	2,919	2,919	2,919	3,953	3,953	3,953
D2C – D2D	0.069	-0.009	0.078	0.398	0.194	0.204
<i>p</i> -value	<0.01	0.77	<0.01	<0.01	0.16	0.06

Table B.4: 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-customer (D2C) and dealer-to-dealer (D2D) trades in five-year immediate off-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 index CDS 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 20–50 MM for CDX.IG and USD 5–15 MM for CDX.HY), for large-sized trades (LRG; USD 50–100 MM for CDX.IG and USD 15–30 MM for CDX.HY), for block-sized trades (BLCK; +USD 100 MM for CDX.IG and +USD 30 MM for CDX.HY), and for trades with transaction prices at typical reference levels (RFRNC; index CDS 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 2,834 (85) and 3,806 (147) outright D2C (D2D) trades in five-year immediate off-the-run index CDSs on CDX.IG and CDX.HY, respectively.

Appendix B. Appendix to Chapter 2

Type	Dealer-To-Customer			Dealer-To-Dealer			D2C-D2D		
	Effc Sprd	Rlzd Sprd	Prc Imp	Effc Sprd	Rlzd Sprd	Prc Imp	Effc Sprd	Rlzd Sprd	Prc Imp
Panel A: CDX.IG									
Outright	0.139	0.030	0.109	0.090	0.024	0.065	0.049**	0.006	0.043**
Index roll	0.052	0.020	0.031	0.063	0.044	0.019	-0.011	-0.024*	0.012
Panel B: CDX.HY									
Outright	0.679	0.172	0.507	0.407	0.167	0.240	0.272**	0.005	0.267**
Index roll	0.406	0.249	0.157	0.454	0.103	0.351	-0.049	0.146**	-0.194**

Table B.5: 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-customer (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, for index rolls between five-year on-the-run and immediate off-the-run index CDSs, and for delta exchanges of index swaption and index tranche swap trades that reference the five-year on-the-run index. 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 index CDS 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 means are identical at the 1% and 5% level, respectively, with inference is 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 51,108 (9,204) and 74,320 (10,720) outright D2C (D2D) trades in five-year on-the-run index CDSs on CDX.IG and CDX.HY, respectively, and 968 (344) and 1,283 (343) D2C (D2D) index rolls between five-year on-the-run and immediate off-the-run index CDSs on CDX.IG and CDX.HY, respectively.

and using alternative time windows when matching pairs of trades with the same characteristics.

B.7.1 Robustness of Results to Alternative Mid-Quote

As a robustness check, we repeat Section 2.4 analyses using the mid-quote of a custom data set that CMA made available to us. Table B.5 displays average effective half-spreads, realized half-spreads, and price impacts of outright trades, index rolls, and delta exchanges that involve five-year on-the-run index CDSs when using CMA mid-quotes to compute half-spreads and price impacts. For outright trades, transaction costs based on CMA mid-quotes are almost identical to those based on Markit mid-quotes. The main difference in comparison to Table 2.3 of the paper is the significantly lower price impact of D2C index rolls in CDX.HY. This does not seem to be due the different trades used because we obtain the same result when constraining trades to be identical.

Table B.6 breaks down average effective half-spreads, realized half-spreads, and price impacts

B.7. Robustness Checks

Trade Size	<i>Dealer-To-Customer</i>			<i>Dealer-To-Dealer</i>			<i>D2C-D2D</i>		
	Effc Sprd	Rlzd Sprd	Prc Imp	Effc Sprd	Rlzd Sprd	Prc Imp	Effc Sprd	Rlzd Sprd	Prc Imp
Panel A: CDX.IG									
≤ 25	0.123	0.034	0.089	0.083	0.014	0.069	0.040**	0.020**	0.020**
25–50	0.132	0.019	0.112	0.095	0.025	0.071	0.036**	-0.006	0.042**
50–100	0.146	0.016	0.130	0.095	0.049	0.046	0.051**	-0.033**	0.084**
> 100	0.170	0.049	0.121	0.121	0.133	-0.012	0.049**	-0.084**	0.133**
Panel B: CDX.HY									
≤ 5	0.606	0.185	0.421	0.389	0.156	0.233	0.217**	0.029	0.188**
5–10	0.644	0.134	0.511	0.410	0.135	0.275	0.235**	-0.001	0.236**
10–25	0.705	0.107	0.598	0.413	0.248	0.164	0.292**	-0.141**	0.433**
> 25	0.803	0.290	0.513	0.549	0.416	0.133	0.255**	-0.126	0.381**

Table B.6: 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-customer (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 index CDS 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 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 51,108 (9,204) and 74,320 (10,720) outright D2C (D2D) trades in five-year on-the-run index CDSs on CDX.IG and CDX.HY, respectively.

of outright trades in five-year on-the-run index CDSs by quartiles of the trade size distribution. Similar to what we observed for outright trades above, half-spreads and price impacts based on CMA mid-quotes are quantitatively similar to those based on Markit mid-quotes and inference regarding higher transaction costs and price impacts of D2C trades does not depend on the particular mid-quote used.

Table B.7 displays regression results for outright trades. Qualitatively, we obtain the same results as in the paper and, in case of some regression coefficients, even quantitatively similar estimates. Moreover, overall inference is unaltered, confirming our earlier results that D2C trades have both higher transaction costs and larger price impacts. The main difference in comparison to Table 2.5 of the paper is the magnitude of coefficient estimates on the bid-ask spread and the mid-quote which suggest that, in comparison to Markit, CMA's bid-ask spread is a relatively less important explanatory variable while its mid-spread is a relatively more important explanatory variable. This most likely reflects methodological differences.

Table B.8 displays regression results for index rolls. Results for CDX.IG index rolls are fairly

Appendix B. Appendix to Chapter 2

	CDX.IG			CDX.HY		
	EffcSprd	RlzdSprd	PrcImp	EffcSprd	RlzdSprd	PrcImp
D2C	0.122** (60.43)	0.029** (12.31)	0.093** (29.99)	0.604** (77.06)	0.174** (16.29)	0.430** (30.30)
D2D	0.088** (32.30)	0.022** (5.39)	0.066** (14.84)	0.395** (31.14)	0.183** (8.23)	0.213** (10.09)
MDM	0.008** (5.14)	-0.010** (-3.61)	0.017** (5.88)	0.026** (3.83)	-0.053** (-4.38)	0.079** (6.67)
LRG	0.018** (8.45)	-0.011** (-3.15)	0.029** (7.99)	0.074** (9.15)	-0.067** (-5.03)	0.141** (9.24)
BLCK	0.044** (16.45)	0.019** (4.90)	0.025** (6.97)	0.195** (20.15)	0.105** (6.17)	0.090** (5.38)
RFRNC	0.020** (7.00)	0.028** (5.60)	-0.008 (-1.61)	0.123** (6.70)	0.162** (5.50)	-0.040 (-1.40)
BAS	0.284** (5.77)	0.051 (1.03)	0.233** (3.44)	0.242** (9.81)	0.083** (2.64)	0.159** (3.67)
SPRD/100	0.145** (3.97)	0.036 (1.05)	0.108* (2.30)	0.148** (3.72)	0.002 (0.05)	0.146* (2.35)
VLTLY	0.209** (5.19)	-0.134** (-3.11)	0.343** (5.78)	1.257** (7.93)	-0.249 (-1.35)	1.507** (5.53)
<i>N</i>	60,312	60,312	60,312	85,040	85,040	85,040
D2C – D2D	0.034	0.007	0.027	0.209	-0.009	0.218
<i>p</i> -value	<0.01	0.08	<0.01	<0.01	0.67	<0.01

Table B.7: 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-customer (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 index CDS 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; index CDS 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 (VLTLY). 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 51,108 (9,204) and 74,320 (10,720) outright D2C (D2D) trades in five-year on-the-run index CDSs on CDX.IG and CDX.HY, respectively.

B.7. Robustness Checks

	CDX.IG			CDX.HY		
	EffcSprd	RlzdSprd	PrcImp	EffcSprd	RlzdSprd	PrcImp
D2C	0.053** (12.01)	0.025** (5.45)	0.028** (5.02)	0.395** (15.66)	0.255** (9.56)	0.141** (5.12)
D2D	0.066** (8.36)	0.046** (4.37)	0.019* (2.23)	0.479** (6.78)	0.125** (2.86)	0.355** (7.27)
BLCK	-0.003 (-0.63)	-0.007 (-1.03)	0.004 (0.62)	0.005 (0.13)	-0.016 (-0.46)	0.021 (0.65)
BAS	0.193** (2.83)	0.109 (1.45)	0.085 (1.22)	0.081** (2.63)	0.024 (0.50)	0.057 (1.16)
SPRD/100	0.021 (0.50)	0.068 (1.04)	-0.047 (-0.92)	0.069 (1.74)	0.057 (1.09)	0.012 (0.28)
VLTLTY	0.040 (0.61)	-0.106 (-1.61)	0.145* (2.08)	0.630** (2.93)	0.309 (1.05)	0.321 (1.04)
<i>N</i>	1,312	1,312	1,312	1,626	1,626	1,626
D2C – D2D	-0.013	-0.022	0.009	-0.084	0.130	-0.214
<i>p</i> -value	0.13	0.08	0.43	0.14	<0.01	<0.01

Table B.8: Regressions Controlling for Index Roll 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 dealer-to-customer (D2C) and dealer-to-dealer (D2D) index rolls between five-year on-the-run and immediate off-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 difference between on-the-run and immediate off-the-run transaction prices and m_t is the difference between the latest on-the-run and immediate off-the-run mid-quotes 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 index CDS 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), and for block-sized trades (BLCK; +USD 100 MM for CDX.IG and +USD 25 MM 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 968 (344) and 1,283 (343) D2C (D2D) index rolls between five-year on-the-run and immediate off-the-run index CDSs on CDX.IG and CDX.HY, respectively.

consistent with those of the paper, but there seem to be some differences between the index roll mid-quotes that implied by Markit and CMA mid-quotes. Nevertheless, both sets of result are qualitatively consistent, in that we either do not find differences in transaction costs of D2C and D2D index rolls or in that the difference are fully accounted for by trade characteristics and market conditions. For CDX.HY index rolls, we note that the significantly lower price impact of D2C index rolls cannot be explained by roll characteristics and market conditions.

Appendix B. Appendix to Chapter 2

Matching Trade Size	Dealer-To-Customer			Dealer-To-Dealer			D2C-D2D		
	Effc Sprd	Rlzd Sprd	Prc Imp	Effc Sprd	Rlzd Sprd	Prc Imp	Effc Sprd	Rlzd Sprd	Prc Imp
Panel A: CDX.IG									
≤ 25	0.116	0.038	0.078	0.087	0.016	0.071	0.029**	0.022*	0.007
25–50	0.125	0.020	0.105	0.102	0.021	0.082	0.022**	-0.001	0.023*
50–100	0.127	-0.001	0.128	0.105	0.049	0.056	0.022*	-0.050*	0.072**
> 100	0.154	0.096	0.057	0.113	0.129	-0.016	0.041	-0.033	0.074*
Exact	0.122	0.027	0.095	0.098	0.019	0.079	0.024**	0.008	0.016*
Panel B: CDX.HY									
≤ 5	0.567	0.160	0.408	0.393	0.147	0.246	0.175**	0.013	0.162**
5–10	0.586	0.129	0.458	0.444	0.146	0.298	0.142**	-0.018	0.160**
10–25	0.622	0.071	0.551	0.437	0.228	0.209	0.185**	-0.157**	0.342**
> 25	0.621	0.152	0.469	0.484	0.349	0.135	0.137	-0.197	0.334*
Exact	0.595	0.104	0.491	0.438	0.154	0.284	0.157**	-0.049	0.206**

Table B.9: Effective Half-Spreads, Realized Half-Spreads, and Price Impacts of Matched Pairs. Panels A and B show sample means of effective half-spreads (EfficSprd), realized half-spreads (RlzdSprd), and price impacts (Prclmp) of matched pairs of outright dealer-to-customer (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. EfficSprd 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. Prclmp is defined as $q_t \times (m_{t+\Delta} - m_t)$. Both transaction prices and quotes are in terms of index CDS 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 identical trade size) that occur within a 15-minute window bracketing the D2D trade. In case of more than one matching D2C trade, the EfficSprd, RlzdSprd, and Prclmp 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 4,794 (3,441) and 6,730 (5,314) (exactly) matched pairs of outright D2C and D2D trades in five-year on-the-run index CDSs on CDX.IG and CDX.HY, respectively.

Table B.9 displays the results of the matched pair analysis. Results are consistent with both the matched pair analysis based on Markit mid-quotes and the regression-based adjustment for selection biases in Table B.7.

B.7.2 5- and 30-Minute Realized Half-Spreads and Price Impacts

Table B.10 displays average effective half-spreads, realized half-spreads, and price impacts when we use 5- and 30-minute periods instead of 15-minute periods to compute realized half-spreads and price impacts. Consistent with information getting gradually incorporated into prices, shorter period price impacts are smaller than longer period price impacts. Regardless

B.7. Robustness Checks

Trade Size	<i>Dealer-To-Customer</i>			<i>Dealer-To-Dealer</i>			<i>D2C-D2D</i>		
	<i>Effc Sprd</i>	<i>Rlzd Sprd</i>	<i>Prc Imp</i>	<i>Effc Sprd</i>	<i>Rlzd Sprd</i>	<i>Prc Imp</i>	<i>Effc Sprd</i>	<i>Rlzd Sprd</i>	<i>Prc Imp</i>
<i>Panel A1: CDX.IG 5-Minute Period</i>									
≤ 25	0.121	0.060	0.061	0.082	0.041	0.042	0.039**	0.020**	0.019**
25–50	0.131	0.057	0.074	0.096	0.047	0.049	0.035**	0.010**	0.025**
50–100	0.143	0.059	0.084	0.091	0.069	0.021	0.052**	-0.010	0.062**
> 100	0.168	0.092	0.076	0.123	0.139	-0.016	0.044**	-0.047**	0.092**
<i>Panel A2: CDX.IG 30-Minute Period</i>									
≤ 25	0.122	0.032	0.090	0.082	0.011	0.071	0.040**	0.021**	0.019**
25–50	0.132	0.018	0.114	0.096	0.018	0.078	0.036**	0.000	0.036**
50–100	0.143	0.022	0.122	0.090	0.045	0.045	0.053**	-0.024	0.077**
> 100	0.170	0.045	0.125	0.127	0.164	-0.037	0.042**	-0.119**	0.162**
<i>Panel B1: CDX.HY 5-Minute Period</i>									
≤ 5	0.606	0.295	0.311	0.388	0.202	0.186	0.218**	0.093**	0.125**
5–10	0.639	0.277	0.362	0.416	0.233	0.184	0.223**	0.044**	0.179**
10–25	0.702	0.291	0.411	0.397	0.269	0.128	0.305**	0.022	0.283**
> 25	0.803	0.466	0.337	0.505	0.553	-0.048	0.298**	-0.087	0.385**
<i>Panel B2: CDX.HY 30-Minute Period</i>									
≤ 5	0.605	0.167	0.437	0.386	0.131	0.255	0.219**	0.036	0.183**
5–10	0.637	0.118	0.520	0.414	0.139	0.274	0.224**	-0.022	0.245**
10–25	0.700	0.101	0.599	0.394	0.206	0.188	0.306**	-0.105*	0.411**
> 25	0.801	0.235	0.566	0.477	0.401	0.076	0.324**	-0.166	0.490**

Table B.10: 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-customer (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 5-minute (Panels A1 and B1) or 30-minute (Panels A2 or B2) 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 5-minute (Panels A1 and B1) or 30-minute (Panels A2 or B2) period that follows trade execution by 5 minutes (Panels A1 and B1) or 30 minute (Panels A2 and B2). *PrcImp* is defined as $q_t \times (m_{t+\Delta} - m_t)$. Both transaction prices and quotes are in terms of index CDS 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 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 48,316 and 50,084 (8,559 and 8,864) and 68,264 and 71,603 (9,642 and 10,246) outright D2C (D2D) trades in five-year on-the-run index CDSs on CDX.IG and CDX.HY in case of 5- and 30-minute periods, respectively.

of the period used, the price impact of D2C trades is significantly larger than that of D2D trades. When changing period lengths, we also adjust our definition of what constitutes a recent quote in that we base mid-quote computation on the latest quote in the 5- and 30-minute period prior to trade execution. As can be seen from Table B.10, it does not matter whether we require the quote to come from the latest 5- or 30-minute period because quotes are revised frequently

Appendix B. Appendix to Chapter 2

	CDX.IG			CDX.HY		
	EffcSprd	RlzdSprd	PrcImp	EffcSprd	RlzdSprd	PrcImp
D2C	0.120** (67.42)	0.057** (43.49)	0.063** (36.48)	0.612** (72.58)	0.288** (47.64)	0.323** (37.57)
D2D	0.088** (31.55)	0.045** (17.57)	0.043** (16.47)	0.395** (27.72)	0.232** (15.22)	0.163** (13.65)
MDM	0.008** (5.58)	-0.002 (-1.30)	0.010** (6.39)	0.014 (1.90)	-0.020** (-2.74)	0.034** (5.18)
LRG	0.016** (8.29)	0.001 (0.64)	0.014** (7.90)	0.062** (8.10)	-0.007 (-1.06)	0.069** (9.27)
BLCK	0.042** (17.42)	0.033** (13.34)	0.010** (3.80)	0.188** (19.39)	0.171** (15.34)	0.017 (1.88)
RFRNC	0.020** (7.73)	0.022** (7.61)	-0.003 (-1.18)	0.112** (5.94)	0.136** (6.47)	-0.024 (-1.81)
BAS	0.442** (7.78)	0.116** (4.88)	0.325** (6.14)	0.344** (10.49)	0.139** (6.78)	0.205** (7.30)
SPRD/100	0.026 (0.71)	0.044* (2.39)	-0.018 (-0.48)	0.071* (2.12)	0.013 (0.75)	0.058 (1.88)
VLTLTY	0.196** (5.86)	-0.010 (-0.49)	0.206** (6.15)	1.182** (6.95)	0.185* (1.98)	0.997** (5.62)
<i>N</i>	56,875	56,875	56,875	77,906	77,906	77,906
D2C – D2D	0.032	0.013	0.020	0.217	0.056	0.161
<i>p</i> -value	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01

Table B.11: 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-customer (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 5-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 5-minute period that follows trade execution by 5 minutes. PrcImp is defined as $q_t \times (m_{t+\Delta} - m_t)$. Both transaction prices and quotes are in terms of index CDS 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; index CDS 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 48,316 (8,559) and 68,264 (9,642) outright D2C (D2D) trades in five-year on-the-run index CDSs on CDX.IG and CDX.HY, respectively.

B.7. Robustness Checks

	CDX.IG			CDX.HY		
	EffcSprd	RlzdSprd	PrcImp	EffcSprd	RlzdSprd	PrcImp
D2C	0.121** (65.53)	0.028** (7.90)	0.094** (24.20)	0.611** (73.64)	0.156** (8.89)	0.455** (27.49)
D2D	0.088** (31.67)	0.020** (3.44)	0.068** (11.13)	0.393** (28.74)	0.158** (4.16)	0.235** (7.09)
MDM	0.008** (5.49)	-0.010** (-2.62)	0.018** (4.57)	0.013 (1.82)	-0.048** (-2.68)	0.061** (3.49)
LRG	0.015** (7.86)	-0.005 (-1.21)	0.020** (4.55)	0.061** (7.83)	-0.059** (-3.15)	0.119** (6.26)
BLCK	0.043** (17.20)	0.017** (3.18)	0.026** (5.09)	0.188** (19.52)	0.066* (2.56)	0.122** (5.53)
RFRNC	0.021** (7.92)	0.020** (2.62)	0.002 (0.23)	0.108** (5.94)	0.189** (4.70)	-0.081* (-2.35)
BAS	0.444** (8.36)	0.119 (1.11)	0.325** (2.73)	0.346** (10.99)	0.074 (0.80)	0.272** (3.26)
SPRD/100	0.021 (0.60)	0.042 (0.65)	-0.021 (-0.29)	0.065* (2.00)	-0.027 (-0.50)	0.092 (1.76)
VLTLTY	0.205** (6.02)	-0.131 (-1.87)	0.335** (4.02)	1.203** (7.18)	-0.395 (-1.11)	1.599** (4.45)
<i>N</i>	58,948	58,948	58,948	81,849	81,849	81,849
D2C – D2D	0.033	0.008	0.025	0.218	-0.002	0.220
<i>p</i> -value	<0.01	0.20	<0.01	<0.01	0.95	<0.01

Table B.12: 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-customer (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 30-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 30-minute period that follows trade execution by 30 minutes. PrcImp is defined as $q_t \times (m_{t+\Delta} - m_t)$. Both transaction prices and quotes are in terms of index CDS 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; index CDS 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 50,084 (8,864) and 71,603 (10,246) outright D2C (D2D) trades in five-year on-the-run index CDSs on CDX.IG and CDX.HY, respectively.

enough.⁵⁰

Tables B.11 and B.12 show trade-by-trade regression results when we use 5- and 30-minute periods to compute mid-quotes, realized half-spreads, and price impacts. The signs of coefficient estimates and their magnitudes are consistent with those based on a 15-minute period and, generally speaking, inference is not affected either. The exception is the 30-minute price impact of block-sized trades in CDX.HY index CDSs, which is no longer significantly lower than the one of trades with trade size in the third quartile of the trade size distribution. This is due to the fact that many of the block-sized trades are block eligible and, therefore, disseminated with a delay of usually 15 minutes.

B.7.3 Time Window in Matched Pair Analysis

The matched pairs of trades that we consider in the paper consist of a D2D trade and D2C match with trade size in the same quartile of the trade size distribution that occurs within a 15-minute window bracketing the execution of the D2D trade. Table B.13 shows what happens if instead we consider matches that occur within a 5- or 30-minute window bracketing trade execution. In shorter windows there are less matches than in longer windows, but in general the results for different window sizes are quite consistent. For most trade sizes that we consider (and in case that we require trade sizes to match exactly), pairs consist of D2C trades that have significantly higher transaction costs and larger price impacts than the D2D trades whose characteristics they are supposed to reflect.

B.8 Standard Error Computation

This section provides details about how we compute the standard errors of cumulative impulse responses and their long-run limits in the vector autoregressive (VAR) model.

B.8.1 Standard Errors of Cumulative Impulse Responses

In order to simplify the presentation, we express the VAR system in Equations (2.3a) to (2.3c) by a single (implicit) vector-valued equation; that is,

$$y_t = \Phi_0 y_t + \Phi_1 y_{t-1} + \dots + \Phi_p y_{t-p} + \epsilon_t \quad (\text{B.1})$$

where $y_t = (\Delta m_t, x_t^{\text{D2C}}, x_t^{\text{D2D}})' \in \mathbb{R}^K$, $\epsilon_t = (\epsilon_t^{\Delta m}, \epsilon_t^{x, \text{D2C}}, \epsilon_t^{x, \text{D2D}})' \in \mathbb{R}^K$, $K = 3$, and $p = 10$. We also generalize the presentation to allow for vector-valued trade-related variables, $x_t^{\text{D2C}} \in \mathbb{R}^N$ and $x_t^{\text{D2D}} \in \mathbb{R}^N$, in which case $y_t = (\Delta m_t, (x_t^{\text{D2C}})', (x_t^{\text{D2D}})')' \in \mathbb{R}^K$, $K = 2N + 1$, and, similarly, $\epsilon_t = (\epsilon_t^{\Delta m}, (\epsilon_t^{x, \text{D2C}})', (\epsilon_t^{x, \text{D2D}})')' \in \mathbb{R}^K$, with $\epsilon_t^{x, \text{D2C}}$ and $\epsilon_t^{x, \text{D2D}}$ having zero conditional means

⁵⁰On average, there are 457.34 and 402.14 composite quotes per day for five-year on-the-run CDX.IG and CDX.HY, respectively, around 97% of which occur in the ten-and-a-half-hour period from 7:00 a.m. to 5:30 p.m. New York time. This suggest that during this period a composite quote is computed every one and a half minutes, on average.

B.8. Standard Error Computation

Matching Trade Size	<i>Dealer-To-Customer</i>			<i>Dealer-To-Dealer</i>			<i>D2C-D2D</i>		
	Effc Sprd	Rlzd Sprd	Prc Imp	Effc Sprd	Rlzd Sprd	Prc Imp	Effc Sprd	Rlzd Sprd	Prc Imp
<i>Panel A1: CDX.IG 5-Minute Window</i>									
≤ 25	0.117	0.032	0.085	0.085	0.012	0.073	0.031**	0.020*	0.012
25–50	0.128	0.028	0.101	0.111	0.011	0.100	0.017*	0.017	0.001
50–100	0.136	0.014	0.121	0.118	0.084	0.034	0.018	-0.070*	0.087**
> 100	0.162	0.115	0.047	0.092	0.120	-0.028	0.071**	-0.005	0.075**
Exact	0.127	0.028	0.099	0.101	0.013	0.089	0.026**	0.015	0.011
<i>Panel A2: CDX.IG 30-Minute Window</i>									
≤ 25	0.115	0.036	0.079	0.084	0.017	0.067	0.031**	0.019*	0.012
25–50	0.120	0.029	0.091	0.101	0.017	0.084	0.020**	0.012	0.008
50–100	0.130	0.025	0.105	0.092	0.049	0.043	0.039**	-0.024	0.063**
> 100	0.143	0.100	0.043	0.115	0.159	-0.044	0.028	-0.060	0.088**
Exact	0.120	0.026	0.094	0.095	0.022	0.073	0.025**	0.004	0.021**
<i>Panel B1: CDX.HY 5-Minute Window</i>									
≤ 5	0.602	0.114	0.488	0.417	0.088	0.329	0.185**	0.026	0.159**
5–10	0.609	0.126	0.483	0.490	0.199	0.291	0.119**	-0.073	0.193**
10–25	0.601	0.098	0.503	0.463	0.306	0.157	0.137**	-0.208*	0.346**
> 25	0.696	-0.010	0.706	0.422	0.335	0.086	0.274	-0.345*	0.620*
Exact	0.608	0.082	0.526	0.473	0.160	0.313	0.135**	-0.078	0.213**
<i>Panel B2: CDX.HY 30-Minute Window</i>									
≤ 5	0.570	0.160	0.410	0.371	0.091	0.280	0.199**	0.069	0.131**
5–10	0.572	0.135	0.436	0.433	0.161	0.272	0.139**	-0.026	0.164**
10–25	0.602	0.127	0.475	0.398	0.207	0.192	0.203**	-0.080	0.283**
> 25	0.661	0.223	0.438	0.457	0.514	-0.057	0.205**	-0.291*	0.495**
Exact	0.588	0.123	0.465	0.415	0.143	0.272	0.174**	-0.020	0.193**

Table B.13: 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-customer (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 index CDS 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 identical trade size) that occur within a 5-minute (Panels A1 and B1) or 30-minute (Panels A2 and B2) window bracketing 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 2,331 and 6,392 (1,548 and 4,939) and 3,437 and 8,249 (2,502 and 7,047) (exactly) matched pairs of outright D2C and D2D trades in five-year on-the-run index CDSs on CDX.IG and CDX.HY within 5- and 30-minute windows, respectively.

Appendix B. Appendix to Chapter 2

and conditional covariances given by $\Sigma_{x,D2C}$ and $\Sigma_{x,D2D}$, respectively.

The $K \times K$ -dimensional coefficient matrices in Equation (B.1) are given by

$$\Phi_0 = \begin{pmatrix} 0 & \beta'_0 & \gamma'_0 \\ 0_N & 0_{N \times N} & 0_{N \times N} \\ 0_N & \lambda_0 & 0_{N \times N} \end{pmatrix} \quad \text{and} \quad \Phi_j = \begin{pmatrix} \alpha_j & \beta'_j & \gamma'_j \\ \delta_j & \zeta_j & \eta_j \\ \kappa_j & \lambda_j & \rho_j \end{pmatrix}, \quad j = 1, \dots, p, \quad (\text{B.2})$$

where 0_n and $0_{n \times m}$ denote a n -dimensional vector and a $m \times n$ matrix of zeros, respectively, $\beta_j, \gamma_j, \delta_j$, and κ_j are N -dimensional vectors, and $\zeta_j, \eta_j, \lambda_j$, and ρ_j are $N \times N$ matrices.

Let $\theta_{\Delta m} = (\beta'_0, \gamma'_0, \alpha_1, \beta'_1, \gamma'_1, \dots, \alpha_p, \beta'_p, \gamma'_p) \in \mathbb{R}^{n_{\Delta m}}$, $n_{\Delta m} = K - 1 + pK = 2N + pK$, be the coefficients of Equation (2.3a), let $\theta_{x,D2C} = \text{vec}([\delta_1, \zeta_1, \eta_1, \dots, \delta_p, \zeta_p, \eta_p]') \in \mathbb{R}^{n_{x,D2C}}$, $n_{x,D2C} = pKN$, be the coefficients of the vector-valued generalization of Equation (2.3b), and let $\theta_{x,D2D} = \text{vec}([\lambda_0, \kappa_1, \lambda_1, \rho_1, \dots, \kappa_p, \lambda_p, \rho_p]') \in \mathbb{R}^{n_{x,D2D}}$, $n_{x,D2D} = N^2 + pKN$, be the coefficients of the vector-valued generalization of Equation (2.3c). (Note that $\theta_{x,D2C}$ and $\theta_{x,D2D}$ stack coefficients of the N equations of the vector-valued generalizations of Equations (2.3b) and (2.3c), respectively.) We further collect individual coefficients in a n -dimensional coefficient vector $\theta = (\theta'_{\Delta m}, \theta'_{x,D2C}, \theta'_{x,D2D})' \in \mathbb{R}^n$, $n = K - 1 + pK^2 + N^2 = 2N + pK^2 + N^2$.

Coefficient estimates are obtained by separately estimating each equation of the VAR system by OLS. We assume that the stacked single-equation OLS coefficient estimates are asymptotically normal

$$\sqrt{T}(\hat{\theta} - \theta) \xrightarrow[T \rightarrow \infty]{d} N(0_n, V), \quad (\text{B.3})$$

with block-diagonal covariance matrix V given by⁵¹

$$V = \begin{pmatrix} V_{\Delta m} & 0_{n_{\Delta m} \times n_{x,D2C}} & 0_{n_{\Delta m} \times n_{x,D2D}} \\ 0_{n_{x,D2C} \times n_{\Delta m}} & V_{x,D2C} & 0_{n_{x,D2C} \times n_{x,D2D}} \\ 0_{n_{x,D2D} \times n_{\Delta m}} & 0_{n_{x,D2D} \times n_{x,D2C}} & V_{x,D2D} \end{pmatrix}, \quad (\text{B.4})$$

where $V_{\Delta m}$, $V_{x,D2C}$, and $V_{x,D2D}$ are the asymptotic covariance matrices of $\hat{\theta}_{\Delta m}$, $\hat{\theta}_{x,D2C}$, and $\hat{\theta}_{x,D2D}$, respectively.

The impulse response function, Ψ_s , $s = 0, 1, \dots$, of the VAR system in Equation (B.1) tracks how an isolated unit-sized shock to one of the system variables propagates through the system. Specifically, $(\Psi_s)_{i,j}$ is the value of the i -th system variable s periods after a one unit shock of the j -th variable under the assumption that the system is in steady state initially, i.e., $y_{t-1} = \dots = y_{t-p} = 0_K$.⁵² Because Equation (B.1) defines y_t implicitly, contemporaneous

⁵¹ V is block diagonal because we resolve contemporaneous effects between quote changes, Δm_t , and trade-related variables, x_t^{D2C} and x_t^{D2D} , while contemporaneous effects between the elements that make up trade-related variables are not resolved, i.e., $\Sigma_{x,D2C}$ and $\Sigma_{x,D2D}$ are non-diagonal.

⁵² As mentioned above, contemporaneous effects between the elements that make up trade-related variables are not resolved. Thus, the $(\Psi_s)_{i,j}$ in case that $N > 3$ ignore the fact that shocks to elements that make up x_t^{D2C}

responses are given by

$$\Psi_0 = (I_K - \Phi_0)^{-1}. \quad (\text{B.5})$$

First-period responses take into account that the previous-period response is Ψ_0 . Solving for the implicitly defined y_{t+1} s (one for each unit-sized shock) gives

$$\Psi_1 = \Psi_0 \Phi_1 \Psi_0. \quad (\text{B.6})$$

Similarly, second-period responses take into account that previous-period responses are Ψ_1 and Ψ_0 , respectively. Solving for the implicitly defined y_{t+2} s gives

$$\Psi_2 = \Psi_0(\Phi_1 \Psi_1 + \Phi_2 \Psi_0). \quad (\text{B.7})$$

Continuing in this fashion shows that the impulse response function of the VAR system in Equation (B.1) satisfies the following recursive relation

$$\Psi_s = \Psi_0(\Phi_1 \Psi_{s-1} + \Phi_2 \Psi_{s-2} + \dots + \Phi_p \Psi_{s-p}), \quad s = 1, 2, \dots, \quad (\text{B.8})$$

with initial values $\Psi_0 = (I_K - \Phi_0)^{-1}$ and $\Psi_s = 0_{K \times K}$ for all $s < 0$.⁵³

Equations (B.5) to (B.8) show that the elements of each Ψ_s are continuous functions of the parameter vector θ . Given an estimate of the latter, Ψ_s can be estimated and the estimate's asymptotic distribution follows by an application of the delta method. Specifically, the asymptotic distribution of the estimate of $\psi_s = \text{vec}(\Psi'_s)$ is given by

$$\sqrt{T}(\hat{\psi}_s - \psi_s) \xrightarrow[T \rightarrow \infty]{d} N(0_{K^2}, G_s V G'_s), \quad (\text{B.9})$$

where the $K^2 \times n$ matrix $G_s = \partial \psi_s / \partial \theta'$ denotes the Jacobian.

Due to Equation (B.8), Jacobian matrices satisfy a recursive relation as well; that is,

$$\begin{aligned} G_s = & (\Psi_0 \otimes [\Psi'_{s-1}, \Psi'_{s-2}, \dots, \Psi'_{s-p}]) \frac{\partial \text{vec}([\Phi_1, \Phi_2, \dots, \Phi_p]')}{\partial \theta'} \\ & + (\Psi_0 [\Phi_1, \Phi_2, \dots, \Phi_p] \otimes I_K) [G'_{s-1}, G'_{s-2}, \dots, G'_{s-p}]' \\ & + (I_K \otimes (\Psi'_{s-1} \Phi'_1 + \Psi'_{s-2} \Phi'_2 + \dots + \Psi'_{s-p} \Phi'_p)) G_0 \end{aligned} \quad (\text{B.10})$$

for $s > 0$,

$$G_0 = (\Psi_0 \otimes \Psi'_0) \frac{\partial \text{vec}(\Phi'_0)}{\partial \theta'}, \quad (\text{B.11})$$

(x_t^{D2D}) are contemporaneously correlated with shocks to the remaining elements of x_t^{D2C} (x_t^{D2D}). This is not an issue because in our application we only consider simultaneous shocks to all D2C (D2D) trade-related variables.

⁵³Note that Ψ_s , $s = 0, 1, \dots$, are the coefficient matrices of the vector moving average representation of the VAR system in Equations (2.4a) to (2.4c).

and $G_s = 0_{K^2 \times n}$ for $s < 0$.

Finally, note that cumulative impulse responses, $\Lambda_s = \sum_{u=0}^s \Psi_u$, $s = 0, 1, \dots$, are linear combinations of Ψ_u with $0 \leq u \leq s$. Thus, $\ell_s = \text{vec}(\Lambda'_s) = R_s \pi_s$, where $R_s = (1'_{s+1} \otimes I_{K^2})$, 1_n denotes a n -dimensional vector of ones, and π_s denotes the n_s -dimensional vector of stacked ψ_{us} with $0 \leq u \leq s$, i.e., $\pi_s = [\psi'_0, \psi'_1, \dots, \psi'_s]'$, where $n_s = (s+1)K^2$. It follows from Equation (B.9) that

$$\sqrt{T}(\widehat{\pi}_s - \pi_s) \xrightarrow[T \rightarrow \infty]{d} N(0_{n_s}, H_s V H'_s), \quad (\text{B.12})$$

where $H_s = [G'_0, G'_1, \dots, G'_s]'$ and, consequently,

$$\sqrt{T}(\widehat{\ell}_s - \ell_s) \xrightarrow[T \rightarrow \infty]{d} N(0, R_s H_s V H'_s R'_s). \quad (\text{B.13})$$

B.8.2 Standard Error of Price Impact

In the VAR system, price impact is captured by the long-run cumulative impulse response of mid-quote revisions in response to unit-sized shocks of trade-related variables. Long-run cumulative impulse responses of the VAR system in Equation (B.1) are given by

$$\Lambda = \lim_{s \rightarrow \infty} \sum_{u=0}^s \Psi_u = (I_K - \Phi_0 - \Phi_1 - \dots - \Phi_p)^{-1}. \quad (\text{B.14})$$

As before, the elements of Λ are continuous functions of θ and, therefore, an application of the delta method yields the asymptotic distribution of the $\widehat{\theta}$ -based estimate of Λ . Specifically, the asymptotic distribution of the estimate of $\lambda = \text{vec}(\Lambda')$ is given by

$$\sqrt{T}(\widehat{\lambda} - \lambda) \xrightarrow[T \rightarrow \infty]{d} N(0_{K^2}, G V G'), \quad (\text{B.15})$$

with Jacobian

$$G = \frac{\partial \lambda}{\partial \theta'} = (\Lambda \otimes \Lambda' \underbrace{[I_K, I_K, \dots, I_K]}_{=(1'_{p+1} \otimes I_K), p+1 \text{ times}}) \frac{\partial \text{vec}([\Phi_0, \Phi_1, \dots, \Phi_p]')}{\partial \theta'}. \quad (\text{B.16})$$

B.9 VAR Models in Trade Size

In this section, we re-estimate the VAR system in Equations (2.3a) to (2.3c) using additional D2C- and D2D-trade-related variables that take trade size into account. Specifically, we estimate the VAR system in Equations (2.3a) to (2.3c) with vector-valued generalizations of Equations (2.3b) and (2.3c) in which $x_t^\tau = (n_t^\tau, v_t^\tau)'$ or $x_t^\tau = (n_t^\tau, s_t^\tau)'$, $\tau \in \{\text{D2C}, \text{D2D}\}$, where v_t^τ (s_t^τ) is the aggregate signed trade size (square-root trade size) of all τ -type trades that occur between the $t-1$ -th and t -th quote revisions (i.e., v_t^τ and s_t^τ are sums of products of trade direction indicators, q_{ut} , and trade sizes, size_{ut} , or square-root trade sizes, $\sqrt{\text{size}_{ut}}$, with u between the calen-

B.9. VAR Models in Trade Size

	Coefficient Estimates					Granger Causality Tests		
	Δm	n^{D2C}	ν^{D2C}	n^{D2D}	ν^{D2D}	Δm	D2C	D2D
Panel A1: CDX.IG								
Δm_t	0.342 (66.40)	0.018 (24.80)	0.1×10^{-3} (11.58)	0.011 (6.16)	-0.1×10^{-3} (-1.51)		4936.9 [<0.01]	160.8 [<0.01]
n_t^{D2C}	2.030 (21.61)	0.157 (12.04)	0.001 (7.29)	0.020 (0.64)	0.000 (0.06)	1064.5 [<0.01]		23.1 [0.28]
ν_t^{D2C}	119.931 (19.74)	5.464 (6.49)	0.139 (10.55)	1.240 (0.60)	0.005 (0.10)	873.3 [<0.01]		25.8 [0.17]
n_t^{D2D}	-0.175 (-3.41)	0.049 (6.66)	0.1×10^{-3} (0.66)	0.136 (7.87)	0.000 (0.06)	27.8 [<0.01]	117.5 [<0.01]	
ν_t^{D2D}	-10.366 (-4.51)	1.711 (5.15)	0.009 (1.67)	3.753 (4.81)	0.045 (2.54)	44.3 [<0.01]	111.2 [<0.01]	
Panel A2: CDX.HY								
Δm_t	0.254 (42.51)	0.104 (44.26)	0.7×10^{-3} (6.93)	0.038 (5.11)	-0.001 (-1.26)		5419.5 [<0.01]	88.7 [<0.01]
n_t^{D2C}	0.410 (15.26)	0.295 (29.10)	0.002 (3.46)	-0.057 (-1.75)	0.005 (1.69)	846.0 [<0.01]		25.0 [0.20]
ν_t^{D2C}	7.045 (11.02)	3.711 (15.36)	0.096 (9.15)	-1.067 (-1.38)	0.082 (1.26)	383.0 [<0.01]		30.7 [0.06]
n_t^{D2D}	-0.070 (-5.62)	0.041 (8.48)	0.5×10^{-3} (2.20)	0.073 (4.89)	0.006 (4.37)	48.5 [<0.01]	136.3 [<0.01]	
ν_t^{D2D}	-0.869 (-5.86)	0.381 (6.56)	0.007 (2.62)	-0.124 (-0.69)	0.135 (8.93)	52.8 [<0.01]	112.9 [<0.01]	
Panel B: Price Impact								
	CDX.IG			CDX.HY				
	D2C	D2D	D2C – D2D	D2C	D2D	D2C – D2D		
Λ	0.057 (62.74)	0.018 (10.48)	0.039 (20.87)	0.244 (60.91)	0.048 (6.70)	0.196 (23.90)		
Panel C: Price Discovery								
	CDX.IG			CDX.HY				
	D2C	D2D	Trade-Unrelated	D2C	D2D	Trade-Unrelated		
R^2	29.29	0.84	69.87	37.94	0.29	61.77		

Table B.14: VAR Estimates.

The table shows coefficient estimates of event-time vector autoregressive (VAR) models for mid-quote revisions (Δm) and signed numbers (n^{D2C} and n^{D2D} , resp.) and signed volumes (ν^{D2C} and ν^{D2D} , resp.) of dealer-to-customer (D2C) and dealer-to-dealer (D2D) trades, respectively, that occur between quote revisions. Panels A1 and A2 show sums of VAR coefficient estimates (t -statistics 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. Panel B shows 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 response to median-sized protection-buyer-initiated 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 variance). Quotes are in terms of index CDS spreads and expressed in basis points. 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.

dar time of the $t - 1$ -th and t -th quote revision) and n_t^τ is the number of signed τ -type trades, i.e., the trade-related variable that we use in Section 2.5 of the paper. As a consequence of including the additional trade-related variables, price impacts, as captured by the long-run cumulative quote revisions in response to isolated D2C and D2D trades, become affine linear in trade size (square-root trade size).

Table B.14 displays VAR coefficient estimates, estimated price impacts, and contributions to price discovery in case that trade-related variables consists of the number of signed trades and the aggregate signed trade size, i.e., $x_t^\tau = (n_t^\tau, v_t^\tau)'$. The results in Panels A1 and A2 of the table qualitatively mirror those in Table 2.8 of the paper in all aspects. The price impact estimates in Panel B are in response to median-sized protection-buyer-initiated trades, i.e., trades of USD 50 million in CDX.IG and trades of USD 10 million in CDX.HY. As before, D2C trades have significantly larger price impacts than D2D trades. Panel C shows contributions of trade-related and trade-unrelated components to price discovery that are quantitatively similar to those reported in the paper. The coefficients on v_t^{D2C} and v_t^{D2D} in Equation (2.3a) suggest that price impact of D2C trades increases with trade whereas price impact of D2D trades is insensitive to trade size.

Table B.15 displays results in case that trade-related variables consists of the number of signed trades and the aggregate signed square-root trade size, i.e., $x_t^\tau = (n_t^\tau, s_t^\tau)'$. The table confirms that our results do not hinge on the particular specification of the VAR and indicate that not much additional insight can be gained from including trade size among trade-related variables. As before, price discovery shares are quantitatively similar to those reported in the paper which are based on trade-related variables consisting of the number of signed trades only, i.e., $x_t^\tau = n_t^\tau$. This suggests that it is the occurrence of trades rather than their size that accounts for most of the information content of trades. This is consistent with evidence from the equity market where trade occurrence and not trade size generates volatility (see, e.g., Jones et al. (1994)). Moreover, approximating non-linearities in the relation between quote changes and trade-related variables by means of non-linear transformations of trade size seems to be of minor importance.⁵⁴

Finally, Figure B.2 shows that the cumulative quote revisions implied by the VAR models that include additional size-based trade-related variables. In line with the above, the implied cumulative quote revisions are almost identical and similar to the one implied by the VAR specification that does not include the additional trade-related variables.

⁵⁴In a VAR model with trade-related variables consisting of the number of signed trades, the aggregate signed trade size, and the aggregate signed square-root trade size (i.e., $x_t^\tau = (n_t^\tau, v_t^\tau, s_t^\tau)'$, $\tau \in \{\text{D2C}, \text{D2D}\}$), D2D-(D2C-)trade-related variables account for is 29.28% (0.94%) of the variance of CDX.IG efficient price innovations with the remaining 69.78% being trade-unrelated. The corresponding numbers for CDX.HY are 38.45% (0.30%) and 61.25%.

B.9. VAR Models in Trade Size

	Coefficient Estimates					Granger Causality Tests		
	Δm	n^{D2C}	s^{D2C}	n^{D2D}	s^{D2D}	Δm	D2C	D2D
Panel A1: CDX.IG								
Δm_t	0.341 (66.37)	0.013 (11.36)	0.002 (11.68)	0.012 (3.28)	-0.001 (-0.97)		4939.9 [<0.01]	158.8 [<0.01]
n_t^{D2C}	2.030 (21.61)	0.104 (5.13)	0.019 (7.07)	0.006 (0.10)	0.002 (0.24)	1063.3 [<0.01]		21.7 [0.35]
s_t^{D2C}	14.708 (21.23)	0.346 (2.32)	0.189 (9.31)	0.041 (0.09)	0.021 (0.28)	1016.1 [<0.01]		24.7 [0.21]
n_t^{D2D}	-0.174 (-3.40)	0.050 (4.36)	0.000 (0.31)	0.134 (3.77)	0.001 (0.11)	27.7 [<0.01]	118.6 [<0.01]	
s_t^{D2D}	-1.291 (-3.96)	0.283 (3.87)	0.007 (0.71)	0.545 (2.41)	0.049 (1.38)	35.5 [<0.01]	118.8 [<0.01]	
Panel A2: CDX.HY								
Δm_t	0.253 (42.30)	0.080 (19.56)	0.010 (9.42)	0.044 (2.61)	-0.004 (-0.83)		5450.2 [<0.01]	88.9 [<0.01]
n_t^{D2C}	0.407 (15.15)	0.243 (13.74)	0.021 (4.83)	-0.125 (-1.73)	0.037 (1.66)	842.1 [<0.01]		26.2 [0.16]
s_t^{D2C}	1.578 (14.32)	0.637 (8.79)	0.147 (8.22)	-0.475 (-1.61)	0.138 (1.49)	682.8 [<0.01]		29.0 [0.09]
n_t^{D2D}	-0.070 (-5.66)	0.036 (4.20)	0.004 (1.74)	-0.016 (-0.48)	0.047 (4.55)	48.9 [<0.01]	134.2 [<0.01]	
s_t^{D2D}	-0.236 (-5.93)	0.099 (3.65)	0.014 (2.05)	-0.253 (-2.38)	0.208 (6.27)	53.0 [<0.01]	127.7 [<0.01]	
Panel B: Price Impact								
	CDX.IG			CDX.HY				
	D2C	D2D	D2C – D2D	D2C	D2D	D2C – D2D		
Λ	0.060 (65.66)	0.018 (9.83)	0.042 (20.71)	0.243 (62.34)	0.048 (6.73)	0.194 (23.94)		
Panel C: Price Discovery								
	CDX.IG			CDX.HY				
	D2C	D2D	Trade-Unrelated	D2C	D2D	Trade-Unrelated		
R^2	29.29	0.84	69.87	38.12	0.29	61.59		

Table B.15: VAR Estimates.

The table shows coefficient estimates of event-time vector autoregressive (VAR) models for mid-quote revisions (Δm) and signed numbers (n^{D2C} and n^{D2D} , resp.) and signed square-root volumes (s^{D2C} and s^{D2D} , resp.) of dealer-to-customer (D2C) and dealer-to-dealer (D2D) trades, respectively, that occur between quote revisions. Panels A1 and A2 show sums of VAR coefficient estimates (t -statistics 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. Panel B shows 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 response to median-sized protection-buyer-initiated 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). Quotes are in terms of index CDS spreads and expressed in basis points. 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 B.2: VAR-Model-Implied Price Impact.

The panels show cumulative quote revisions in response to either a single median-sized protection-buyer-initiated dealer-to-customer (D2C; solid black lines) trade or a single median-sized 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 (Panels A and C) and CDX.HY (Panels B and D). Cumulative quote revisions in Panels A and B (Panels C and D) are implied by event-time vector autoregressive models for mid-quote revisions, the sum of signed D2C trades that occur between quote revisions and their signed (square-root) volume, and the sum of signed D2D trades that occur between quote revisions and their signed (square-root) volume. Dashed lines mark 95% confidence intervals based on OLS standard errors. Quotes are in terms of index CDS spreads and expressed in basis points (bps). Median size of trades in CDX.IG (CDX.HY) index CDSs is USD 50 million (USD 10 million). 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.

B.10 Additional Figures and Tables

Panels A and B of Figure B.3 display mid-quote changes, $m_{t+\Delta} - m_t$, following the execution of non-block trades as function of Δ . The panels distinguish between D2C and D2D trades as well as protection-buyer- and protection-seller-initiated trades. In particular, the panels show the average mid-quote change in one-minute periods following trade execution where the average is taken over all pairs of trades and quotes for which the difference between trade execution and quotation times falls within the respective one-minute period. Because we

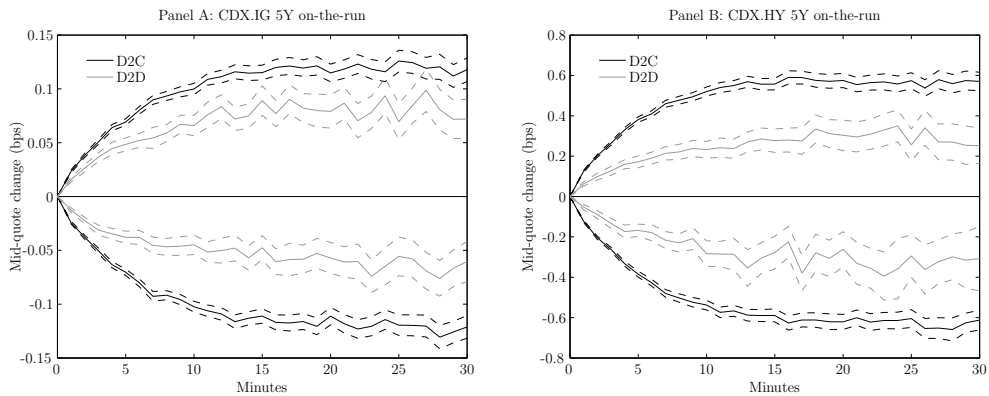


Figure B.3: Mid-Quote Changes as a Function of Time.

Panels A and B show mid-quote changes following non-block outright dealer-to-customer (D2C) and dealer-to-dealer (D2D) trades in five-year on-the-run index CDSs on CDX.IG and CDX.HY, respectively, as a function of time. Mid-quote changes are defined as $c(\Delta) = m_{t+\Delta} - m_t$, where m_t is the latest mid-quote in the 15-minute period prior to trade execution and $m_{t+\Delta}$ is any mid-quote Δ seconds after trade execution. Mid-quotes are in terms of index CDS spreads and expressed in basis points (bps). The upper (lower) of equally colored straight lines is the minute-by-minute average mid-quote change of protection-buyer-initiated (protection-seller-initiated) trades, with trade direction inferred by the Lee and Ready (1991) algorithm. The average mid-quote change in minute i is the sample mean over all $c(\Delta)$ such that $(i - 1) < \Delta/60 \leq i$ (a trade that is followed by multiple quotes in minute i contributes multiples terms to the computation of the sample mean). Dashed lines mark 95% confidence intervals based on standard errors clustered by trade dissemination identifier and quote timestamp. The sample period is October 2, 2013 to October 16, 2015.

separate the mid-quote changes of protection-buyer- and protection-seller-initiated trades, mid-quote changes have the interpretation of Δ -minute price impacts. Consistent with what we observed in Section B.7.2, the Δ -minute price impacts of D2C trades are significantly larger than those of D2D trades for most values of Δ .

Moreover, Figure B.3 shows that price impacts seem to have converged 15 minutes after trade execution and that there seems to be no difference between the price impacts of protection-buyer- and protection-seller-initiated trades. This is the reason why we choose to work with 15-minute price impacts in the paper and why the VAR models that we estimate do not distinguish between protection-buyer- and protection-seller-initiated trades.

C Appendix to Chapter 3

C.1 Sample Construction

I collect all trade reports of credit asset class swaps that were disseminated between December 31, 2012 and June 30, 2015 by the swap data repositories (SDRs) of Bloomberg (the Bloomberg Swap Data Repository, henceforth BSDR), the Depository Trust & Clearing Corporation (the DTCC Data Repository, henceforth DDR), and the Intercontinental Exchange (the ICE Trade Vault, henceforth ICETV), and remove all trade reports with identical dissemination identifiers at the SDR level (i.e., among trade reports disseminated by the same SDR).¹

In the first step, I remove all canceled trade reports and the corresponding cancelations at the SDR level. Occasionally, a trade report is canceled by both counterparties, and there are some cancelations of trade reports that are not contained in the data. This can lead to removal of a larger number of cancelations than canceled trade reports. I also remove duplicate corrections when a trade report is corrected by both counterparties, and corrections of trade reports that are not contained in the data (this includes corrections without dissemination identifier of the original trade report). Finally, I remove corrections and the original trade reports when the original trade report has not previously been canceled as required by CFTC rules.

In the second step, I remove all historical trade reports and trade reports of non-index-CDS transactions, such as, index options, index tranche swaps, and more exotic credit derivatives. Then, I remove all trade reports (at the SDR level) that are disseminated with insufficient information regarding the underlying credit index. These include trade reports where the underlying is a bespoke basket or where the fields used to identify the underlying are missing or incompletely populated. For trade reports disseminated by the BSDR this concerns the “Ticker” and “CDS Version” fields that contain the index’s Bloomberg ticker and version number, respectively.² For trade reports disseminated by the DDR this concerns the last nine digits

¹I do not collect the 65 trade reports of credit asset class swaps that were disseminated by the SDR of the Chicago Mercantile Exchange because most of them are for historical transactions (i.e., transaction executed prior the effective date of the CFTC’s real-time public reporting requirement) or for indices that are not focus of the paper.

²The BSDR did not disseminate Bloomberg tickers and CDS version numbers during the first two weeks of

Appendix C. Appendix to Chapter 3

of the “Underlying_Asset_1” field that contains the index’s Markit RED code, and for trade reports disseminated by the ICETV this concern the “TVProductMnemonic” field that contains the index’s Trade Vault Product Mnemonic. Bloomberg tickers together with version numbers, Markit RED codes, and Trade Vault Product Mnemonics are individually sufficient to uniquely identify the names, series, and versions of synthetic credit indices composed of corporate, municipality, or sovereign reference names. They are also sufficient to uniquely identify non-synthetic indices composed of agency pools (the MBX, IOS, and PO indices), commercial mortgage-backed securities (the CMBX and TRX indices), or prime and sub-prime residential mortgage-backed securities (the PRIMEX and ABX indices). Because the focus of the paper is on synthetic credit indices, I remove the trade reports of non-synthetic credit indices together with all trade reports for which I am unable to identify the mapping between Markit RED codes and names, series, and versions. I also remove all trade reports with non-standard maturities and incomplete transaction data (i.e., all trade reports with missing execution timestamp, missing price notation type, missing price notation, missing currency denomination, missing notional amount, or missing transaction type). Among the remaining trade reports, I focus on those of the main indices of the CDX North American family: CDX North American Investment Grade (CDX.IG) and CDX North American High Yield (CDX.HY).

Before merging the trade reports from the three SDRs, I remove duplicate trade reports at the SDR level. Duplicate trade reports are defined as trade reports of off-SEF transactions with identical transaction data but different dissemination identifiers, and I remove all duplicate trade reports other than the one that was submitted first (i.e., the one with the smallest dissemination identifier).³ I only remove duplicate trade reports of off-SEF transactions because it is possible that multiple transactions with the same terms occur within a second (the precision of trade report timestamps) on the electronic order book of a SEF.

In the third step, I remove trade reports of transactions executed prior to index launch or after maturity, trade reports with zero price notations, trade reports with 00:00:00 timestamps, and trade reports with non-standard currency denominations.^{4,5}

In the fourth step, I account for reporting errors and filter the data for outliers. For both tasks,

operation. Assuming that the underlying index is not a sub-index of the index whose name is contained in the trade report, it is possible to identify the underlying of transactions in CDX family indices. The identifying assumption may not be unreasonable given that in the remainder of the sample period the BSDR did not disseminate a single trade report of a transaction in which the underlying was a sub-index. Assuming that the version of the underlying index is the one with the maximum number of contributing dealers to Markit’s end-of-day composite computation (Markit depth), it is possible to identify the version number. The identifying assumption may not be unreasonable given the industry’s convention to shift trading from one version to the next on the day after the respective credit event auctions have taken place (which tends to be reflected by a shift in Markit depth).

³I exclude the original dissemination identifier when identifying duplicates because these may mismatch due to correction of erroneous transaction data.

⁴To determine whether a transaction is executed prior to index launch (i.e., the start of trading) or after maturity, I convert Universal Coordinated Time (UTC) timestamps into New York time. According to Markit’s index roll timetables, indices of the CDX North American family start trading at 7:30 New York time on the index launch dates.

⁵I remove both 00:00:00 UTC and 00:00:00 New York time timestamps.

I use Markit composite prices and spreads as reference points. When available, the reference point will be the mid-point of the most recent Markit intraday composite bid and ask quotes prior to trade execution. I only consider two-sided quotes that are available in both price notation types and require quotes to be from the same trading day.⁶ Otherwise, the reference point is either the end-of-day composite from the trading day prior to trade execution, the mid-point of the first intraday composite bid and ask quotes from the same trading day that occur after trade execution, or the end-of-day composite from the day of trade execution. Trade reports for which I am unable to find a reference point are removed from the data.

There are three common reporting errors. First, reporting parties frequently submit trade reports that show the index CDS contracts' fixed spreads (i.e., the rates that determine fixed leg payments of the index CDS contracts) instead of the price notations at which counterparties agreed to settle their trades (trades settle at upfront payments that are exchanged between index CDS counterparties at the inception and close of trade; price and spread price notations uniquely determine these upfront payments). Second, reporting parties frequently fail to express prices in percent and spreads in basis points. Third, reporting parties frequently submit trade reports with incorrect price notation types (e.g., indicating that the price notation type is a spread when in fact the reported price notation is a price).

In order to address the first type of reporting error, I remove all trade reports with price notations that are equal to fixed spreads or scaled (by 1/10000, 1/100, 100, or 10000) multiples thereof and all trade reports with price notations equal to 0.01, 1, 100, 10000, 1000000, 0.05, 5, 500, 50000, 5000000 (note that 100 bps and 500 bps are the most common fixed spreads of index CDS contracts).⁷ In order to address the second type of reporting error, I first remove amount price notations and replace the price notation types of the remaining trade reports with those in which the respective index CDSs are conventionally quoted.⁸ This leaves me with only two price notation types in the data, namely, price and spread. For each of the two price notation types, I define an outlier as a price notation with a percentage difference from the respective reference point that exceeds $q\%$ in absolute value; that is,

$$\frac{|P_k - P(n_k, u_k, t_k)|}{P(n_k, u_k, t_k)} > q(n_k), \quad (\text{C.1})$$

where P_k denotes the price notation of the k -th trade report, n_k denotes its price notation type, u_k denotes the transaction's underlying, t_k denotes its execution timestamp, and $P(n, u, t)$ denotes index u 's reference point (Markit intraday or end-of-day) composite of price notation

⁶Unless otherwise specified, trading day refers to a local time trading day.

⁷Scaling by 1/10000, 1/100, 100, or 10000 addresses the second type of reporting error. Note that the scaling factors correspond to those operations that would be used when changing units to decimals, percentages, or basis points. Price notations of 0.000001, 0.0001, 0.000005, and 0.0005 do not appear in the data and are therefore not contained in the above list.

⁸In principle, prices could be backed out from amount price notations. However, I refrain from doing so because disseminated notional amounts are rounded and may be capped, resulting in prices that potentially differ from those agreed upon by index CDS counterparties.

type n associated with timestamp t .⁹ Then, I identify trade reports with incorrectly expressed price notations by outlier price notations (i.e., trade reports with P_k satisfying Equation (C.1)) for which I can find a scaling factor $s \in \Sigma = \{1/10000, 1/100, 100, 10000\}$ such that

$$\frac{|sP_k - P(n_k, u_k, t_k)|}{P(n_k, u_k, t_k)} \leq q(n_k). \quad (\text{C.2})$$

I replace incorrectly expressed price notations by s^*P_k with s^* having minimal percentage distance (as defined by the left hand side of Equation (C.2)) among all $s \in \Sigma$ that satisfy Equation (C.2).

In order to address the third type of reporting error, I identify trade reports with incorrectly expressed price notation types by outlier price notations for which I can find a scaling factor $s \in \{1\} \cup \Sigma$ such that

$$\frac{|sP_k - P(n, u_k, t_k)|}{P(n, u_k, t_k)} \leq q(n), \quad (\text{C.3})$$

for a price notation type n other than n_k (in fact, there is only one such price notation type for each trade report because at this stage the data only contains trade reports with two different price notation types). I replace incorrectly expressed price notation types by n and eventually replace the corresponding price notations by s^*P_k with s^* having minimal percentage distance (as defined by the left hand side of Equation (C.3)) among all $s \in \{1\} \cup \Sigma$ that satisfy Equation (C.3).

For those trade reports without outlier price notations, I proceed with homogenizing price notation types further in that I use the ISDA CDS standard model to convert price notations of spread type into price notations of price type and vice versa.¹⁰ After conversion, I remove all trade reports with outlier price notations with respect to either price notation type and all trade reports for which conversion failed.

Finally, I remove all trade reports of transactions with notional amounts less than USD 10,000 and all trade reports of transactions executed on non-full trading days (i.e., SIFMA recommended full or early close trading days).

I identify the SEF on which the trade was executed from the trade report format (the identification algorithm is described in detail in the Internet Appendix to Collin-Dufresne et al. (2016)). As described in Collin-Dufresne et al. (2016), the structure of the SEF market is such that the

⁹I use cutoffs of 1% and 5% for price notations of price and spread type, respectively.

¹⁰The model input are standardized contract terms (not the ones contained in the trade reports), including the index's effective date, payment frequency, and day count convention (note that neither Bloomberg's nor Markit's converter allows to modify those terms for the sake of standardization—although they might misvalue contracts with other terms, these converters are still valid tools that index CDS counterparties use to agree on upfront amounts; moreover, there are just a few trade reports with non-standard payment frequencies and for the significant number of trade reports with "1/1" day-count convention conversion would not be possible because the converter does not recognize this type of day count convention). The valuation date is the trading day (T) and the protection effective date is T+1.

SEF on which the trade was executed reveals whether the trade was between an end-user and a liquidity providing dealer or whether it was an interdealer trade. For trade reports for which I can identify the SEE, I aggregate trade sizes of transactions with identical terms and the same execution timestamp.

By convention dealers price index-related instruments, such as, index options and index tranche swaps, in reference to the corresponding index CDS and with the implicit understanding that trade execution includes an offsetting index CDS trade of a delta neutralizing notional amount (the so-called “delta exchange”). The transaction price at which the delta exchange takes place is called the “reference level” and included in dealer runs for index options and index tranche swaps.¹¹ But in contrast to the option or tranche swap quote, the reference level does not necessarily reflect the current index level because it is usually fixed at market opening (see, e.g., Hünseler (2013)). When reference levels change throughout the trading day, they tend to change by much coarser increments than quotes of the corresponding index CDSs.¹² I make use of this fact for identifying delta exchanges from the transaction data. Specifically, I collect the reference levels of end-of-day index option composites from Markit and those of intraday tranche swap quotes from Credit Market Analysis and, on a given trading day (UTC trading day in case of the intraday tranche swap quotes), I identify all transactions with reference level transaction prices as delta exchanges.

¹¹Credit derivatives dealers provide their institutional clients with quotes for index CDSs and index-related instruments by instant messaging, e-mails, or via single-dealer screens. A quote update distributed to a wide variety of clients by one of these means is referred to as a dealer run.

¹²For instance, CDX.IG reference levels increment by 0.5 bps while order books in the interdealer market typically employ 0.0625 bps tick sizes. Similarly, CDX.HY reference levels increment by 0.125% while order books typically employ 0.01% increments.

Bibliography

- Acharya, Viral V., Yakov Amihud, and Sreedhar T. Bharath, 2013, Liquidity risk of corporate bond returns: Conditional approach, *Journal of Financial Economics* 110, 358–386.
- Acharya, Viral V., and Lasse H. Pedersen, 2005, Asset pricing with liquidity risk, *Journal of Financial Economics* 77, 375–410.
- Adrian, Tobias, Erkki Etula, and Tyler Muir, 2014, Financial intermediaries and the cross-section of asset returns, *Journal of Finance* 69, 2557–2596.
- Amihud, Yakov, 2002, Illiquidity and stock returns: Cross-section and time-series effects, *Journal of Financial Markets* 5, 31–56.
- Ang, Andrew, Robert J. Hodrick, Yuhang Xing, and Xiaoyan Zhang, 2006, The cross-section of volatility and expected returns, *Journal of Finance* 61, 259–299.
- Arora, Navneet, Priyank Gandhi, and Francis A. Longstaff, 2012, Counterparty credit risk and the credit default swap market, *Journal of Financial Economics* 103, 280–293.
- Asquith, Paul, Thomas R. Covert, and Parag Pathak, 2013, The effects of mandatory transparency in financial market design: Evidence from the corporate bond market, Working paper, National Bureau of Economic Research.
- Bao, Jack, and Jun Pan, 2013, Bond illiquidity and excess volatility, *Review of Financial Studies* 26, 3068–3103.
- 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.
- Bartolini, Leonardo, Spence Hilton, Suresh Sundaresan, and Christopher Tonetti, 2011, Collateral values by asset class: Evidence from primary securities dealers, *Review of Financial Studies* 24, 248–278.
- 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.

Bibliography

- Berndt, Antje, Rohan Douglas, Darrell Duffie, Mark Ferguson, and David Schranz, 2005, Measuring default risk premia from default swap rates and EDFs, Working paper, Stanford University.
- Berndt, Antje, and Iulian Obreja, 2010, Decomposing European CDS returns, *Review of Finance* 14, 189–233.
- 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.
- Bessembinder, Hendrik, and Kumar Venkataraman, 2004, Does an electronic stock exchange need an upstairs market?, *Journal of Financial Economics* 73, 3–36.
- Biswas, Gopa, Stanislava Nikolova, and Christof W. Stahel, 2015, The transaction costs of trading corporate credit, Working paper, University of Nebraska–Lincoln.
- Boehmer, Ekkehart, Charles M. Jones, and Xiaoyan Zhang, 2015, Potential pilot problems: Treatment spillovers in financial regulatory experiments, Working paper, Columbia University.
- Bongaerts, Dion, Frank de Jong, and Joost Driessen, 2011, Derivative pricing with liquidity risk: Theory and evidence from the credit default swap market, *Journal of Finance* 66, 203–240.
- , 2012, An asset pricing approach to liquidity effects in corporate bond markets, Working paper, Tilburg University.
- Brunnermeier, Markus K., and Lasse H. Pedersen, 2009, Market liquidity and funding liquidity, *Review of Financial Studies* 22, 2201–2238.
- Bühler, Wolfgang, and Monika Trapp, 2009, Time-varying credit risk and liquidity premia in bond and CDS markets, Working paper, University of Cologne.
- Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller, 2011, Robust inference with multiway clustering, *Journal of Business & Economic Statistics* 29, 238–249.
- Campello, Murillo, Long Chen, and Lu Zhang, 2008, Expected returns, yield spreads, and asset pricing tests, *Review of Financial Studies* 21, 1297–1338.
- Chen, Kathryn, Michael Fleming, John Jackson, Ada Li, and Asani Sarkar, 2011, An analysis of CDS transactions: Implications for public reporting, Staff report, Federal Reserve Bank of New York.

- Chen, Ren-Raw, Frank Fabozzi, and Ronald Sverdlow, 2010, Corporate credit default swap liquidity and its implications for corporate bond spreads, *Journal of Fixed Income* 20, 31–57.
- Collin-Dufresne, Pierre, Benjamin Junge, and Anders B. Trolle, 2016, Market structure and transaction costs of index CDSs, Working paper, EFPL.
- Crossen, Christopher, and Xu Zhang, 2011, Validating the public EDF model for global financial firms, Moody's Analytics Quantitative Research.
- Dick-Nielsen, Jens, 2009, Liquidity biases in TRACE, *Journal of Fixed Income* 19, 43–55.
- Dick-Nielsen, Jens, Peter Feldhütter, and David Lando, 2012, Corporate bond liquidity before and after the onset of the subprime crisis, *Journal of Financial Economics* 103, 471–492.
- Driessen, Joost, 2005, Is default event risk priced in corporate bonds?, *Review of Financial Studies* 18, 165–195.
- Duffie, Darrell, 2010, Presidential address: Asset price dynamics with slow-moving capital, *Journal of Finance* 65, 1237–1267.
- Duffie, Darrell, Leandro Saita, and Ke Wang, 2007, Multi-period corporate default prediction with stochastic covariates, *Journal of Financial Economics* 83, 635–665.
- Duffie, Darrell, and Haoxiang Zhu, 2015, Size discovery, Working paper, Stanford University.
- Edwards, Amy K., Lawrence E. Harris, and Michael S. Piwowar, 2007, Corporate bond market transaction costs and transparency, *Journal of Finance* 62, 1421–1451.
- Elton, Edwin J., Martin J. Gruber, Deepak Agrawal, and Christopher Mann, 2001, Explaining the rate spread on corporate bonds, *Journal of Finance* 56, 247–277.
- Erlandsson, Ulf, Arup Ghosh, and Graham Rennison, 2008, Systematic CDS index trading, Barclays Capital Quantitative Credit Strategy Research.
- Fama, Eugene F., and Kenneth R. French, 1993, Common risk factors in the returns on stocks and bonds, *Journal of Financial Economics* 33, 3–56.
- Filipović, Damir, and Anders B. Trolle, 2013, The term structure of interbank risk, *Journal of Financial Economics* 109, 707–733.
- Franzoni, Francesco, Eric Nowak, and Ludovic Phalippou, 2012, Private equity performance and liquidity risk, *Journal of Finance* 67, 2341–2373.
- Froot, Kenneth A., and Paul G. J. O'Connell, 1999, The pricing of U.S. catastrophe reinsurance, in Kenneth A. Froot, ed.: *The Financing of Catastrophe Risk* (University of Chicago Press, Chicago, IL).

Bibliography

- 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.
- Gospodinov, Nikolay, Raymond Kan, and Cesare Robotti, 2014, Misspecification-robust inference in linear asset-pricing models with irrelevant risk factors, *Review of Financial Studies* 27, 2139–2170.
- 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.
- Greenwood, Robin, and Dimitri Vayanos, 2014, Bond supply and excess bond returns, *Review of Financial Studies* 27, 663–713.
- Gromb, Denis, and Dimitri Vayanos, 2002, Equilibrium and welfare in markets with financially constrained arbitrageurs, *Journal of Financial Economics* 66, 361–407.
- Grossman, Sanford J., and Merton H. Miller, 1988, Liquidity and market structure, *Journal of Finance* 43, 617–633.
- Harris, Lawrence E., and Michael S. Piwowar, 2006, Secondary trading costs in the municipal bond market, *Journal of Finance* 61, 1361–1397.
- Hasbrouck, Joel, 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.
- He, Zhiguo, and Arvind Krishnamurthy, 2013, Intermediary asset pricing, *American Economic Review* 103, 732–770.
- Heckman, James, Justin L. Tobias, and Edward Vytlacil, 2003, Sample estimators for treatment parameters in a latent-variable framework, *Review of Economics and Statistics* 85, 748–755.
- 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.
- Hu, Grace Xing, Jun Pan, and Jiang Wang, 2013, Noise as information for illiquidity, *Journal of Finance* 68, 2341–2382.

- Huang, Jing-Zhi, and Ming Huang, 2012, How much of the corporate-treasury yield spread is due to credit risk?, *Review of Asset Pricing Studies* 2, 153–202.
- 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.
- Jones, Charles M., Gautam Kaul, and Marc L. Lipson, 1994, Transactions, volume, and volatility, *Review of Financial Studies* 7, 631–651.
- Kan, Raymond, Cesare Robotti, and Jay Shanken, 2013, Pricing model performance and the two-pass cross-sectional regression methodology, *Journal of Finance* 68, 2617–2649.
- Kondor, Péter, and Dimitri Vayanos, 2014, Liquidity risk and the dynamics of arbitrage capital, Working paper, London School of Economics.
- Korablev, Irina, and Shisheng Qu, 2009, Validating the public EDF model performance during the credit crisis, Moody's Analytics Quantitative Research.
- Korajczyk, Robert A., and Ronnie Sadka, 2008, Pricing the commonality across alternative measures of liquidity, *Journal of Financial Economics* 87, 45–72.
- 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.
- Lesplingart, Clothilde, Christophe Majois, and Mikael Petitjean, 2012, Liquidity and CDS premiums on European companies around the subprime crisis, *Review of Derivatives Research* 15, 257–281.
- Li, Haitao, Junbo Wang, Chunchi Wu, and Yan He, 2009, Are liquidity and information risks priced in the Treasury bond market?, *Journal of Finance* 64, 467–503.
- Lin, Hai, Junbo Wang, and Chunchi Wu, 2011, Liquidity risk and expected corporate bond returns, *Journal of Financial Economics* 99, 628–650.
- Lo, Andrew W., 2002, The statistics of Sharpe ratios, *Financial Analysts Journal* 58, 36–52.
- 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.
- Ludvigson, Sydney C., 2013, Advances in consumption-based asset pricing: Empirical tests, in George M. Constantinides, Milton Harris, and René M. Stulz, eds.: *Handbook of the Economics of Finance* (Elsevier Science, Amsterdam, North Holland).

Bibliography

- Managed Funds Association, 2015, Why eliminating post-trade name disclosure will improve the swaps market, Position paper.
- Merton, Robert C., 1974, On the pricing of corporate debt: The risk structure of interest rates, *Journal of Finance* 29, 449–470.
- Mitchell, Mark, and Todd Pulvino, 2012, Arbitrage crashes and the speed of capital, *Journal of Financial Economics* 104, 469–490.
- Moskowitz, Tobias J., Yao Hua Ooi, and Lasse H. Pedersen, 2012, Time series momentum, *Journal of Financial Economics* 104, 228–250.
- Newey, Whitney K., and Kenneth D. West, 1987, A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix, *Econometrica* 55, 703–708.
- O’Kane, Dominic, 2008, *Modelling single-name and multi-name credit derivatives* (John Wiley & Sons, Chichester, West Sussex).
- Pagano, Marco, and Ailsa Röell, 1996, Transparency and liquidity: A comparison of auction and dealer markets with informed trading, *Journal of Finance* 51, 579–611.
- Pástor, Ľuboš, and Robert F. Stambaugh, 2003, Liquidity risk and expected stock returns, *Journal of Political Economy* 111, 642–685.
- Pires, Pedro, Joao Pereira, and Luis Martins, 2014, The empirical determinants of credit default swap spreads: A quantile regression approach, *European Financial Management*, Forthcoming.
- Qiu, Jiaping, and Fan Yu, 2012, Endogenous liquidity in credit derivatives, *Journal of Financial Economics* 103, 611–631.
- Reiss, Peter C., and Ingrid M. Werner, 1998, Does risk sharing motivate interdealer trading?, *Journal of Finance* 53, 1657–1703.
- Sadka, Ronnie, 2006, Momentum and post-earnings-announcement drift anomalies: The role of liquidity risk, *Journal of Financial Economics* 80, 309–349.
- , 2010, Liquidity risk and the cross-section of hedge-fund returns, *Journal of Financial Economics* 98, 54–71.
- Schultz, Paul, 2001, Corporate bond trading costs: A peak behind the curtain, *Journal of Finance* 56, 677–698.
- Siriwardane, Emil N., 2015, Concentrated capital losses and the pricing of corporate credit risk, Working paper, New York University.
- Tang, Dragon Yongjun, and Hong Yan, 2007, Liquidity and credit default swap spreads, Working paper, University of Hong Kong.

Vassalou, Maria, and Yuhang Xing, 2004, Default risk in equity returns, *Journal of Finance* 59, 831–868.

Zhu, Haoxiang, 2012, Finding a good price in opaque over-the-counter markets., *Review of Financial Studies* 25, 1255–1285.

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

Jan Benjamin Junge

Swiss Finance Institute at EPFL
Quartier UNIL-Dorigny, Extranef 128
1015 Lausanne, Switzerland

Email: benjamin.junge@epfl.ch
Phone: +41 21 69 30104
Web: <https://people.epfl.ch/benjamin.junge>

Education

- 09/2010– Ph.D. in Finance
Swiss Finance Institute at EPFL
Swiss Federal Institute of Technology Lausanne (EPFL), Switzerland
Advisor: Prof. Anders B. Trolle
- 10/2005–04/2010 M.Sc. in Mathematical Finance (Diplom-Finanzökonom math.)
University of Constance, Germany

Research Interests

Credit Risk, Empirical Asset Pricing, Liquidity, Market Microstructure

Working Papers

Liquidity Risk in Credit Default Swap Markets (with Anders B. Trolle)
Revise and resubmit at the *Review of Financial Studies*

Market Structure and Transaction Costs of Index CDSs (with Pierre Collin-Dufresne and Anders B. Trolle)

Index CDS Trading Costs around the Introduction of SEFs

Seminars and Conference Presentations

- 2016 SFI Research Days
- 2013 3rd International Conference of the Financial Engineering and Banking Society, IV World Finance Conference, Asian Finance Association Annual Meeting, Northern Finance Association Annual Conference
- 2012 11th Swiss Doctoral Workshop in Finance

Teaching Experience

- 2016, 2015, 2014 Swiss Federal Institute of Technology Lausanne (EPFL), Switzerland
Teaching Assistantship, Financial Econometrics (M.Sc.)
Prof. Florian Pelgrin

Curriculum Vitae

2013, 2012, 2011 Swiss Federal Institute of Technology Lausanne (EPFL), Switzerland
Teaching Assistantship, Econometrics (M.Sc.)
Prof. Lorian Mancini

Work Experience

08/2009–10/2009 Bayerische Hypo- und Vereinsbank AG, Munich, Germany
Corporate & Investment Banking—Structured Products Development
Internship

02/2008–04/2008 Sal. Oppenheim jr. & Cie. KGaA, Frankfurt, Germany
Trading & Derivatives—Structured Solutions
Internship

08/2007–10/2007 Sal. Oppenheim jr. & Cie. KGaA, Frankfurt, Germany
Trading & Derivatives—Quantitative Analysis
Internship

08/2006 F&C Asset Management plc, London, U.K.
Emerging Market Debt
Internship
