

No. 581

Martin D. Gould, Nikolaus Hautsch, Sam D. Howison,
and Mason A. Porter

Counterparty Credit Limits: An Effective Tool for Mitigating Counterparty Risk?

The CFS Working Paper Series

presents ongoing research on selected topics in the fields of money, banking and finance. The papers are circulated to encourage discussion and comment. Any opinions expressed in CFS Working Papers are those of the author(s) and not of the CFS.

The Center for Financial Studies, located in Goethe University Frankfurt's House of Finance, conducts independent and internationally oriented research in important areas of Finance. It serves as a forum for dialogue between academia, policy-making institutions and the financial industry. It offers a platform for top-level fundamental research as well as applied research relevant for the financial sector in Europe. CFS is funded by the non-profit-organization Gesellschaft für Kapitalmarktforschung e.V. (GfK). Established in 1967 and closely affiliated with the University of Frankfurt, it provides a strong link between the financial community and academia. GfK members comprise major players in Germany's financial industry. The funding institutions do not give prior review to CFS publications, nor do they necessarily share the views expressed therein.

Counterparty Credit Limits: An Effective Tool for Mitigating Counterparty Risk?

Martin D. Gould* ^{†1}, Nikolaus Hautsch², Sam D. Howison³, and Mason A. Porter^{3,4,5}

¹CFM–Imperial Institute of Quantitative Finance, Department of Mathematics,
Imperial College, London SW7 2AZ, UK

²Department of Statistics and Operations Research, University of Vienna, Vienna
A-1090, Austria

³Mathematical Institute, University of Oxford, Oxford OX2 6GG, UK

⁴CABDyN Complexity Centre, University of Oxford, Oxford OX1 1HP, UK

⁵Department of Mathematics, University of California, Los Angeles, CA 90095, USA

Abstract

A counterparty credit limit (CCL) is a limit imposed by a financial institution to cap its maximum possible exposure to a specified counterparty. Although CCLs are designed to help institutions mitigate counterparty risk by selective diversification of their exposures, their implementation restricts the liquidity that institutions can access in an otherwise centralized pool. We address the question of how this mechanism impacts trade prices and volatility, both empirically and via a new model of trading with CCLs. We find empirically that CCLs cause little impact on trade. However, our model highlights that in extreme situations, CCLs could serve to destabilize prices and thereby influence systemic risk.

Keywords: Counterparty credit limits; counterparty risk; price formation; market design; systemic risk.

1 Introduction

The recent international financial crisis has underlined the vital importance of understanding counterparty risk. As exemplified by the collapse of Lehman Brothers and the ensuing defaults and near-defaults by AIG, Bear Stearns, Fannie Mae, Freddie Mac, Merrill Lynch, the Icelandic banks, and the Royal Bank of Scotland, the complex and highly interconnected nature of the modern financial ecosystem can cause counterparty failures to propagate rapidly between institutions and can thereby amplify their severity [May et al., 2008]. Consequently, assessing and implementing measures to mitigate the risk of default contagion remains a task of primary importance.

One possible mitigation measure, currently implemented by several multi-institution trading platforms in the foreign exchange (FX) spot market, is the application of counterparty credit limits (CCLs). A CCL is a limit imposed by a financial institution to cap its

*Corresponding author. Email: m.gould@imperial.ac.uk.

[†]Martin D. Gould completed part of this work while at the University of Oxford and as an academic visitor at the Humboldt-Universität zu Berlin.

maximum possible exposure to a specified counterparty. CCLs are designed to complement existing risk-based capital requirements by protecting financial institutions from the large losses that can result from sudden counterparty defaults, such as those that occurred with the failure of Lehman Brothers in 2008.

Despite this clear benefit, the application of CCLs also entails an important drawback. With CCLs in place, institutions can only access the trading opportunities offered by counterparties with whom they possess sufficient bilateral credit. Therefore, CCLs cause otherwise centralized liquidity to fracture into localized liquidity pools, such that different institutions have access to different trading opportunities at the same time. Because it limits individual institutions' access to liquidity, this process could severely destabilize prices, and may therefore influence systemic risk. Indeed, several prominent market commentators¹ made a similar argument in opposition to the proposed introduction of mandatory credit exposure limits between large financial institutions (see Basel Committee on Banking Supervision [2014] and Scott [2014]).

The aim of the present paper is to assess how the implementation of CCLs affects liquidity and trade in financial markets. To address this question, we utilize an unusually rich data set from Hotspot FX, which is a large electronic trading platform that enables institutions to apply CCLs to their trading counterparties. Crucially, this data enables us to quantify how CCLs affect both the liquidity available to given institutions and the prices at which they consequently trade.

Our empirical investigation focuses on two main questions. First, how do CCLs impact the prices that individual institutions pay for their trades? To address this question, we introduce the notion of the "skipping cost" of a trade to measure the additional cost borne by an institution due to the liquidity restrictions imposed by its CCLs. Second, to what extent do these liquidity restrictions influence volatility? We address this question by comparing the realized volatility of trade prices with the corresponding realized volatility in the platform-wide best quotes from all institutions, irrespective of their CCLs. Our work is the first study to address these questions.

We find that more than half of all the trades in our data set have a skipping cost of 0 and that the mean additional cost caused by CCLs is less than half a basis point. Although we identify a handful of trades with large skipping costs, we argue that the existence of such trades is a natural consequence of the substantial heterogeneity in the types and sizes of institutions that trade in the FX spot market. We also find that the realized volatility of trade prices is very similar to the corresponding realized volatility in the platform-wide best quotes.

These empirical results all suggest that CCLs have very little impact on trading on Hotspot FX. However, examining how CCLs impact trade on a specific platform is only one aspect of understanding how they might affect financial markets more generally. In particular, empirical study of historical data does not provide insight into how these results might change if institutions were to make substantial modifications to their CCLs. Therefore, we complement our empirical analysis with an investigation of a third question: how does the structure of the underlying CCL network influence skipping costs and trade-price volatility? To address this question, we introduce a model of trade in which institutions assign CCLs to their trading counterparties. In contrast to our empirical analysis, in which the CCL network is fixed and unobservable to us, this approach makes it possible to investigate how varying a market's CCL network affects trade in our model.

¹See, e.g., Goldman Sachs [2012] and J.P. Morgan [2012], and "Banks urge Fed retreat on credit exposure", Financial Times, 15 April 2012.

Simulations of our model provide valuable insights into the possible impact of CCLs. For example, when the CCL network is dense (in the sense that most institutions can access most trading opportunities), we find that the application of CCLs has very little impact on trade prices and volatility. However, as the CCL network's edge density falls, we find that the skipping costs of trades and the corresponding trade-price volatility both rise sharply. These sharp rises are not accompanied by a corresponding rise in quote-price volatility, which remains approximately constant. This result indicates that when CCLs progressively restrict institutions' access to liquidity, their impact on trade escalates considerably. This finding is consistent with concerns regarding the possible dangers of CCLs in situations where liquidity is scarce. We also find that the CCL network's topology can strongly influence both trade prices and volatility. This result raises important questions regarding regulators' abilities to monitor the impact of CCLs among the large and heterogeneous populations of institutions in real markets.

Together, our results paint an interesting and complex picture regarding the impact of CCLs in financial markets. They also raise several important issues for regulators and policy makers. Both our empirical investigation and our model simulations suggest that when liquidity is plentiful, CCLs may indeed be an effective tool for mitigating counterparty risk. However, our model also illustrates how an aggressive application of CCLs can create large jumps in the trade-price series, even when the quote-price series remain relatively stable. We therefore argue that understanding and monitoring how institutions set and adjust their CCLs is a vital step for regulators in assessing how their implementation might impact market stability and, ultimately, whether they constitute a benefit or a hindrance in combating systemic risk.

Our work contributes to the rapidly growing literature that addresses effective ways to mitigate counterparty risk. We discuss several such publications, and highlight how CCLs differ from other possible approaches to this problem in Section 2.2. More generally, we also contribute to the literature about the causes and consequences of counterparty risk. Jarrow and Yu [2001] noted that market-wide risk factors and institution-specific counterparty risks can interact to generate highly correlated failure probabilities for different institutions. Giesecke and Weber [2004] illustrated how the strength of default-contagion effects depend heavily on the specific counterparty network between different institutions. Jorion and Zhang [2009] noted that counterparty risk contagion could explain the strongly clustered nature of defaults over time, and thus conjectured that a fear of counterparty failures could explain the sudden worsening of the international credit crisis after the collapse of Lehman in 2008. Our paper adds a specific focus on how CCLs feature in this discussion.

Our model contributes to the literature on price formation in non-centralized liquidity pools. Perraudin and Vitale [1996] studied price formation in a model of trade in a decentralized market in which designated market makers attempt to infer information based on the order flow that they experience from dealers. Our paper incorporates a similar framework in which individual institutions trade bilaterally. In contrast to Perraudin and Vitale [1996], our core focus is the relationship between the institutions' interaction topology and the consequent impact on trade prices.

The paper proceeds as follows. In Section 2.3, we provide an introduction to counterparty risk, describe the CCL mechanism in detail, and discuss how CCLs are currently implemented by several large electronic trading platforms in the FX spot market. In Section 3, we describe the data that forms the basis of our empirical study. We present our empirical results in Section 4. In Section 5, we introduce and study our model of how CCLs affect trade. We conclude in Section 6.

2 Counterparty Risk and Counterparty Credit Limits

2.1 Counterparty Risk

Counterparty risk is the risk that one or more of a financial institution's counterparties will default on their agreed obligations (see Gregory [2010]). Counterparty defaults can occur for a wide variety of reasons, ranging from technical issues, such as computer system malfunctions, to serious financial difficulties, such as insolvency. Irrespective of their cause, counterparty defaults can cause significant financial distress and can push other institutions towards their own defaults. This, in turn, can create a cascade of rapidly propagating institutional failures. Therefore, counterparty risk is a key factor in determining whether and with what speed localized shocks escalate to systemic events that impact the global economy [The Counterparty Risk Management Policy Group, 2005].

To date, the vast majority of work on counterparty risk has focused on the counterparty credit risk arising from the possibility that a counterparty defaults on its payment obligations from a derivative contract (see, e.g., Brigo et al. [2013] for a detailed survey). Despite the dominance of counterparty credit risk in the existing literature, financial institutions also face several other important types of counterparty risks [Gregory, 2010]. Prominent examples include liquidity risk, which is the risk of a liquidity shortage arising from a counterparty default, and settlement risk, which is the risk of suffering losses by delivering cash or assets to a counterparty that fails to settle the opposite leg of an agreement. Several historical events, such as the near-catastrophic domino-effect defaults caused by the failure of Bankhaus Herstatt in 1974 [Bank for International Settlements, 2002], underline the severity of these forms of counterparty risk and provide strong motivation for exploring safeguards against them.

2.2 Approaches to Mitigating Counterparty Risk

There are several possible approaches to mitigating counterparty risk, among which two have received particular attention. The first is that of trade novation via a central counterparty (CCP); see Norman [2011] and Rehlon and Nixon [2013] for detailed discussions. The role of a CCP is to guarantee the obligations arising from all contracts agreed between two counterparties. If one counterparty fails, then the other is protected via the default-management procedures and resources of the CCP. During the past decade, several prominent regulatory bodies² have argued that CCPs are an effective tool for mitigating counterparty risk.

The second approach is that of applying a credit valuation adjustment (CVA); see Brigo et al. [2013] and Gregory [2010] for detailed discussions. In this framework, an institution adjusts the price that it offers another institution to account for the risk of trading with it. In other words, an institution may offer each other institution a different price for the same transaction to account for its perceived risk of counterparty failure. In principle, an institution can use the additional revenue generated by a CVA to construct a contingent claim whose payoff is triggered by the default of the given counterparty, such that the resulting net loss is zero.

Despite their clear benefits, these approaches to mitigating counterparty risk also suffer from important drawbacks. Koeppl [2013] noted that CCPs generate moral hazard by removing the incentive for individual institutions to assess the creditworthiness of their trading counterparties. Pirrong [2012] argued that CCP novation does not reduce the aggregate

²See, e.g., The Basel Committee on Banking Supervision [2013] and The Counterparty Risk Management Policy Group [2005].

counterparty risk across all institutions, but rather concentrates all such risk into the CCP, which thus becomes a single point of failure with systemic importance. Biais et al. [2012] noted that although CCPs allow mutualization of the idiosyncratic risk faced by individual institutions, they cannot provide protection against the aggregate risk that affects all institutions together. Menkveld [2015] showed that standard methodologies for calculating default probabilities can greatly underestimate the probability of clustered defaults, which place severe stress on a CCP. Given the historical failures of several CCPs in a wide variety of asset classes — including FX, equities, and futures [Gregory, 2010] — concerns about whether CCPs really mitigate risk, or simply repackage it, seem to be well-founded.

CVA also suffers from important drawbacks. Calculating a CVA requires each institution to estimate a time-varying risk premium for each of its trading counterparties. This risk premium depends heavily on the counterparty’s default probability, which is extremely difficult to estimate in practice. Cesari et al. [2010] noted that even if an institution is able to estimate a risk premium for a given counterparty, performing this estimation does not provide any insight into how it should construct a portfolio to provide the required payoff upon a counterparty default. Indeed, constructing this portfolio is often impossible in practice. It is also worth noting that if one regards CVA as an insurance, then the “premium” should be set aside as a provision against default, yet in reality there is always pressure to use such assets for other purposes. Moreover, CVAs are not suitable for assets traded on an exchange in which many different institutions access the same liquidity pool (such as a limit order book (LOB); see Gould et al. [2013] for a detailed introduction to LOBs), because implementing CVAs would require each institution to set different prices for different counterparties trading the same asset.

These weaknesses suggest that despite their widespread discussion and implementation, neither CCP novation nor CVAs constitute a panacea for the issue of counterparty risk. The failure of these measures to provide a conclusive solution to the problem is strong motivation for exploring alternative avenues.

2.3 Counterparty Credit Limits

One alternative approach to mitigating counterparty risk is the application of counterparty credit limits (CCLs). Consider a financial market populated by a set of institutions $\Theta = \{\theta_1, \theta_2, \dots\}$, in which each institution θ_i assigns a CCL $c_{(i,j)} \geq 0$ to each other institution θ_j . The CCL $c_{(i,j)}$ specifies the maximum level of counterparty credit exposure that θ_i is willing to extend to θ_j .³ Assigning a CCL to a given counterparty does not require posting collateral; instead, it simply involves notifying the exchange of the relevant value $c_{(i,j)}$. Institution θ_i cannot engage in any trading activities with θ_j that would make θ_i ’s total exposure to θ_j exceed $c_{(i,j)}$ or that would make θ_j ’s total exposure to θ_i exceed $c_{(j,i)}$. The maximal amount that θ_i and θ_j can trade is therefore equal to $\min(c_{(i,j)}, c_{(j,i)})$. We call this quantity the *bilateral CCL between θ_i and θ_j* . These bilateral CCLs determine which subset of trading opportunities are available to each institution. This subset changes over time according to the relevant institutions’ trading activity.

As argued by Jarrow and Yu [2001], financial institutions face significant counterparty risks whenever their exposures are concentrated within a small number of counterparties, because the default of any such counterparty would likely cause severe financial distress. CCLs

³In the FX spot market, trades agreed on day d are settled on day $d + 2$. Therefore, each trade by an institution in this market entails exposure to the counterparty during the period between trade agreement and trade settlement.

provide financial institutions with an explicit defense against entering this situation, because they cap the maximum exposure that an institution can face from any other institution.

If an institution θ_i perceives another institution θ_j to be unacceptably likely to default, then θ_i can ensure that it never trades with θ_j by setting $c_{(i,j)} = 0$, because arranging any trade with θ_j would result in a non-zero exposure and would thereby violate this CCL. Alternatively, if θ_i perceives θ_j to be extremely unlikely to default, then θ_i can also assign an unlimited amount of credit to θ_j by setting $c_{(i,j)} = \infty$. Irrespective of the CCL set by θ_i , it still remains open to θ_j to further restrict the bilateral exposure by choosing $c_{(j,i)}$ appropriately.

In contrast to CVAs, the application of CCLs does not require institutions to estimate the market value of their counterparty risk. In contrast to trade novation via a CCP, CCLs do not require a single, centralized clearing node that constitutes a single point of failure for an entire market.⁴ Instead, the application of CCLs enables institutions to specify an upper bound on each of their counterparty exposures, and thereby to mitigate counterparty risk by selective diversification of their exposures.

2.4 Counterparty Credit Limits in the FX Spot Market

Several major multi-institution electronic trading platforms in the FX spot market already offer institutions the ability to implement CCLs. On these platforms, each institution θ_i privately declares to the exchange their CCL $c_{(i,j)}$ for each other institution θ_j . Trade occurs via a mechanism similar to a standard limit order book (LOB), except that institutions are only able to conduct transactions that do not violate their bilateral CCLs. More precisely, when an institution θ_i submits a buy (respectively, sell) market order, the order matches to the highest-priority sell (respectively, buy) limit order that is owned by an institution θ_j such that neither $c_{(i,j)}$ nor $c_{(j,i)}$ are violated by conducting the given trade. We call this market organization a *quasi-centralized LOB* (QCLOB) because different institutions have access to different subsets of the same centralized liquidity pool. For a detailed introduction to QCLOBs, see Gould et al. [2016].

Institutions trading on a QCLOB platform cannot in general see the state of the global LOB. Instead, each institution sees only the active orders that correspond to trading opportunities that it can access (i.e., do not violate any of its bilateral CCLs) at time t .⁵ More precisely, for each $j \neq i$, the volume of each separate limit order placed by θ_j that is visible to θ_i is reduced (if necessary) so that its size does not exceed the bilateral CCL between θ_i and θ_j . Each institution therefore views a filtered set of all limit orders on the platform, according to its CCLs.

In addition to viewing their filtered LOB, each institution in a QCLOB can access a *trade-data stream* that lists the price, time, and direction (buy/sell) of each trade that occurs. All institutions can see all entries in the trade-data stream in real time, irrespective of their bilateral CCLs with the institutions involved in a given trade. Therefore, although institutions in a QCLOB can only see a subset of the trading opportunities available to other institutions, they do have access to a detailed historical record of all previous trades.

⁴We note, however, that the use of CCLs does not exclude the subsequent clearing of trades via a CCP. We return to this discussion in Section 6.

⁵Some QCLOB platforms (such as Reuters and EBS) offer institutions the ability to access an additional data feed that provides snapshots of the global LOB at regular time intervals in exchange for a fee.

3 Data and Analysis

3.1 Sample Selection

Our empirical investigation is based on a data set provided to us by Hotspot FX. The data describes all trading activity on the Hotspot FX platform for the EUR/USD (euro/US dollar), GBP/USD (pounds sterling/US dollar), and EUR/GBP (euro/pounds sterling) currency pairs⁶ for the entire months of May and June 2010. According to the 2010 Triennial Central Bank Survey [Bank for International Settlements, 2010], global trade for these currency pairs constituted about 28%, 9%, and 3% of the total turnover of the FX market, respectively. Therefore, our results enable us to compare and contrast the impact of CCLs across currency pairs with different traded volumes.

During our sample period, three major multi-institution trading platforms dominated electronic trading volumes in the FX spot market: Reuters, EBS, and Hotspot FX.⁷ All three of these platforms use similar trading mechanics, and, in particular, all three implement CCLs via QCLOBs. Importantly, however, EBS and Reuters primarily serve the interbank market, whereas Hotspot FX serves both the interbank market and a broad range of other financial institutions, such as hedge funds, commodity trading advisers, corporate treasuries, and institutional asset managers.

Hotspot FX operates continuous trading, 24 hours per day, 7 days per week. However, the vast majority of activity on the platform occurs on weekdays during the *peak trading hours* of 08:00:00–17:00:00 GMT. We exclude all data from outside these time windows to ensure that our results are not influenced by unusual behaviour during inactive periods. We exclude 3 May (May Bank Holiday in the UK) and 31 May (Spring Bank Holiday in the UK; Memorial Day in the US) because market activity on these days was extremely low. We also exclude any days that included a gap in recording lasting 30 seconds or more.

After making these exclusions, our data set contains the peak trading hours for each of 30 trading days. In Figure 1, we plot the total volumes of market orders and limit orders for each of the three currency pairs on each of these days. In Table 1, we provide the corresponding summary statistics. Consistently with the market-wide volume ratios reported by the Bank for International Settlements [2010], the mean daily volume of market orders for EUR/USD exceeds that of GBP/USD by a factor of about 3 and that of EUR/GBP by a factor of about 9.

3.2 Data Format

For each currency pair and each day, the Hotspot FX data consists of two files. The first file is the *tick-data file*, which lists all limit order arrivals and departures. For each limit order arrival, this file lists the price, size, direction (buy/sell), arrival time, and a unique order identifier. For each limit order departure, this file lists the departure time and the departing order’s unique identifier (which is assigned at its arrival). A limit order departure can occur for two reasons: because the order is matched by an incoming market order or because the order is cancelled by its owner. The second file is the *trade-data file*, which lists all trades. For each trade, this file lists the price, size, direction (buy/sell), and trade time. In both files, all times are recorded in milliseconds. For each of the three currency pairs that we study, the platform’s minimum order size is 0.01 units of the base currency, and the platform’s tick

⁶A price for the currency pair XXX/YYYY denotes how many units of the *counter currency* YYY are exchanged per unit of the *base currency* XXX.

⁷See Bech [2012] for an estimated breakdown of transaction volumes between platforms during this period.




Figure 1: Total daily volumes of (top panel) market orders and (bottom panel) limit orders for (green circles) EUR/USD, (orange squares) GBP/USD, and (purple triangles) EUR/GBP activity on Hotspot FX on each day in our sample, measured in units of the counter currency. See Table 1 for the corresponding summary statistics.

	EUR/USD	GBP/USD	EUR/GBP	EUR/USD	GBP/USD	EUR/GBP
	Panel A: Volume of Market Orders			Panel A: Number of Market Orders		
Min	2.5×10^9	7.4×10^8	1.0×10^8	3.3×10^3	1.6×10^3	1.6×10^2
Median	4.4×10^8	1.5×10^9	3.6×10^8	5.4×10^3	2.9×10^3	4.9×10^2
Max	7.5×10^9	2.5×10^9	1.2×10^9	8.7×10^3	4.3×10^3	1.2×10^3
Mean	4.6×10^9	1.5×10^9	4.0×10^8	5.6×10^3	2.9×10^3	5.1×10^2
St. Dev.	1.2×10^9	4.2×10^8	2.4×10^8	1.4×10^3	6.2×10^2	2.1×10^2
	Panel C: Volume of Limit Orders			Panel D: Number of Limit Orders		
Min	7.2×10^{12}	5.5×10^{12}	3.7×10^{12}	3.5×10^6	3.0×10^6	2.0×10^6
Median	9.4×10^{12}	7.8×10^{12}	6.2×10^{12}	4.4×10^6	4.4×10^6	2.9×10^6
Max	1.4×10^{13}	9.7×10^{12}	7.6×10^{12}	6.0×10^6	5.3×10^6	3.6×10^6
Mean	1.0×10^{13}	7.9×10^{12}	6.2×10^{12}	4.5×10^6	4.4×10^6	2.9×10^6
St. Dev.	1.9×10^{12}	9.9×10^{11}	7.9×10^{11}	6.5×10^5	5.1×10^5	3.5×10^5

Table 1: Summary statistics for the total daily (Panel A) volume of market orders, (Panel B) number of market orders, (Panel C) volume of limit orders, and (Panel D) number of limit orders for EUR/USD, GBP/USD, and EUR/GBP activity on Hotspot FX during May–June 2010. All volumes are in units of the counter currency.

size (i.e., the smallest permissible price interval between different orders) is 0.00001 units of the counter currency. For further details regarding trade on the Hotspot FX platform, see Knight Capital Group [2015].

The Hotspot FX data has several features that are particularly important for our study. First, the tick-data files list all limit order arrivals and departures, irrespective of each order’s ownership. This enables us to determine the complete set of all limit orders (irrespective of their owners’ CCLs) for a given currency pair at any time during the sample period. By determining this set at the time of each trade, we are able to calculate detailed statistics regarding the impact of CCLs on trade prices. Second, the small tick sizes on Hotspot FX enable us to observe market participants’ price preferences (i.e., the prices at which they place orders) with a high level of detail. By contrast, analyzing data from platforms with larger tick sizes (such as Reuters and EBS) would provide a more coarse-grained view of such price preferences and could therefore make the results more difficult to interpret, particularly among trades for which the influence of CCLs is small. Third, all limit orders represent actual trading opportunities that were available in the market. This is not the case on some other FX spot trading platforms, which allow institutions to post indicative quotes that do not constitute a firm commitment to trade. Fourth, the trade-data files include explicit buy/sell indicators, which allow us to identify trades without the need for trade-classification inference algorithms (such as the one introduced by Lee and Ready [1991]), which can produce inaccurate results.

The Hotspot FX data also has some weaknesses. First, the tick-data files do not contain information about hidden orders, so there are some trades listed in the trade-data file for which no corresponding limit order departures are reported in the tick-data file. For each of the three currency pairs, these trades account for approximately 5% of the total traded volume. In the absence of further details about these trades, we choose to exclude them from our study. Second, in some extremely busy periods, several limit order departures can occur at the same price in very rapid succession. Therefore, for some trades, it is not possible to determine exactly which limit order departure corresponds to a given trade. For each such trade, we choose to associate the limit order departure whose time stamp is closest to the

reported trade time. We regard any incorrect associations made in this way to be a source of noise in the data. To ensure that this approach does not influence our conclusions, we also repeated all of our calculations when excluding all trades for which it is not possible to associate exactly one limit order departure, and we found that all of our results were qualitatively the same as those that we report throughout the paper.

If a market order matches to several different limit orders, each partial matching is reported as a separate line in the trade-data file, with a time stamp that differs from the previous line by at most 1 millisecond. In the absence of explicit details regarding order ownership, we regard all entries that correspond to a trade of the same direction and that arrive within 1 millisecond of each other as originating from the same market order, and we record the corresponding statistics for this market order only once. For trades that match at several different prices (i.e., “walk up the book”), we record the volume-weighted average price (VWAP) as the price for the whole trade, and calculate the corresponding skipping cost using this VWAP price.⁸

Although the Hotspot FX data does not include information about market activity on Reuters or EBS, we do not regard this to be an important limitation for the present study. Due to the greater heterogeneity among member institutions on Hotspot FX than on Reuters or EBS (see Section 2.4), it seems reasonable to expect that CCLs have a larger impact on trade and liquidity on Hotspot FX than they do on these other platforms. For example, large banks trading on Hotspot FX may be unwilling to trade with small counterparties, and may therefore assign them a CCL of 0. By contrast, the CCLs between institutions on Reuters and EBS are likely to be much higher, to reflect the confidence in large trading counterparties within the interbank market. By studying data from Hotspot FX, we are able to assess the impact of CCLs among a large and heterogeneous population.

3.3 Bid–Ask Bounce

Bid–ask bounce describes the tendency for consecutive trades of a given asset to alternate between being buyer-initiated and seller-initiated (see Roll [1984]). Because the bid–ask spread is by definition strictly positive, the occurrence of bid–ask bounce can cause subsequent trades to occur at different prices, even in the absence of any change to the market state.

Similarly to the application of CCLs, bid–ask bounce is a microstructural effect that impacts price series at the trade-by-trade level. Studying all buyer-initiated and seller-initiated trades together could cause bid–ask bounce to obscure the impact of CCLs in the trades that we observe. Throughout this paper, we therefore study buyer-initiated and seller-initiated trades separately, in an attempt to disentangle our results about CCLs from the possible impact of bid–ask bounce.

4 Empirical Results

4.1 Daily Activity

As a preliminary assessment of market activity, we first consider the time series of price changes between successive trades. Assume that the k^{th} trade for a given currency pair on

⁸Because each partial matching of a single market order is subject to the same CCLs, we regard it as inappropriate to study each such partial matching as a separate event, as doing so would produce long sequences of correlated data points from single market orders.

a given trading day occurs at time t . Let p_k denote the price of this trade, and let b_k , a_k , and m_k denote, respectively, the bid-, ask-, and mid-prices in the global LOB immediately before this trade occurs. Let $p_{k'}$ denote the price of the most recent trade in the same direction (i.e., buyer- or seller-initiated) as the k^{th} trade. Similarly, let $b_{k'}$, $a_{k'}$, and $m_{k'}$ denote, respectively, the bid-, ask-, and mid-prices in the global LOB immediately before the most recent trade in the same direction as the k^{th} trade. The *change in trade price* f_k is given by

$$f_k = \begin{cases} p_k - p_{k'}, & \text{if the } k^{\text{th}} \text{ trade is a buyer-initiated trade,} \\ p_{k'} - p_k, & \text{if the } k^{\text{th}} \text{ trade is a seller-initiated trade.} \end{cases} \quad (1)$$

Because the prices of trades vary across currency pairs and across time, we normalize each price change by the mid-price (i.e., the mean of the best bid- and ask-prices in the global LOB) immediately before the trade occurred. Specifically, we calculate the *normalized change in trade price*

$$\tilde{f}_k = \frac{f_k}{m_k}, \quad (2)$$

which we measure in basis points (where 1 basis point corresponds to 0.01%). Unlike f_k , the normalized price change \tilde{f}_k is scaled to be independent of the size of the underlying exchange rate, and thereby allows like-for-like comparisons across different currency pairs. Note, however, that this normalization makes the gap between successive values on the pricing grid (i.e., the scaled tick size) different for each currency pair, and time-dependent.

In Figure 2, we plot the \tilde{f}_k series on the first day in our sample. The plots for the other trading days are all qualitatively similar. In these series, it is common to observe normalized trade-price changes of more than 5 or even 10 basis points. Many of these large price changes are isolated events that are not surrounded by other price changes with similar magnitudes, which suggests that they are not a consequence of volatility clustering. A core aim of our empirical investigation is to assess the extent to which these large changes in the trade-price series are attributable to CCLs.

4.2 Skipping Costs

We next address the question of how CCLs impact the prices that institutions pay for individual trades. As we discussed in Section 2.4, when an institution θ_i on Hotspot FX submits a buy (respectively, sell) market order, the order matches to the highest-priority sell (respectively, buy) limit order that is owned by an institution θ_j such that the bilateral CCL between θ_i and θ_j is not violated by the trade. Therefore, due to the impact of CCLs, the price at which a given market order matches is not necessarily the best price available to other institutions at that time.

The Hotspot FX data enables us to calculate the difference between the price at which a buyer-initiated (respectively, seller-initiated) trade occurred and the lowest price among all sell (respectively, highest price among all buy) limit orders at the same instant. It thereby enables us to quantify precisely the additional cost borne by the institution that submitted the market order, as a result of CCLs preventing this institution from accessing better-priced liquidity. We call this additional cost the *skipping cost*,

$$r_k = \begin{cases} p_k - a_k, & \text{if the } k^{\text{th}} \text{ trade is a buyer-initiated trade,} \\ b_k - p_k, & \text{if the } k^{\text{th}} \text{ trade is a seller-initiated trade.} \end{cases} \quad (3)$$




Figure 2: Normalized changes in trade price \tilde{f}_k for (top row) EUR/USD, (middle row) GBP/USD, and (bottom row) EUR/GBP trades on Hotspot FX during 08:00:00–17:00:00 GMT on 4 May 2010.

As in Equation (2), we also calculate the *normalized skipping cost*

$$\tilde{r}_k = \frac{r_k}{m_k}. \quad (4)$$

In Equation (3), the sign difference between buyer-initiated and seller-initiated trades reflects that every trade has a non-negative skipping cost. In the extreme case in which the CCLs are such that all institutions always have access to all trading opportunities, all trades occur at the best quotes at their time of execution, so $p_k = q_k$ for all k . In this case, all trades have a skipping cost of $r_k = 0$, so price formation is equivalent to that in a standard LOB.

In Figure 3, we show the empirical cumulative density functions (ECDFs) of normalized skipping costs \tilde{r}_k . More than half of trades have a normalized skipping cost (and therefore a skipping cost) of 0, which implies that they occurred at the best price available in the global LOB at their time of execution. Among the other trades, the distribution of normalized skipping costs has a similar shape for all three currency pairs. As illustrated by the log-log survivor functions,⁹ this similarity extends to about the 99.9th percentile of the distributions. Beyond this, EUR/USD has a handful of trades with extremely large skipping costs that do not occur for the other two currency pairs.




Figure 3: (Left) Empirical cumulative density function (ECDF) and (right) log-log survivor function (i.e., one minus the ECDF) for the normalized skipping costs \tilde{r}_k of (solid green curves) EUR/USD, (dashed orange curves) GBP/USD, and (dotted-dashed purple curves) EUR/GBP trades on Hotspot FX during May–June 2010.

In Table 2, we list the corresponding summary statistics about normalized skipping costs \tilde{r}_k . For each of the three currency pairs, the mean normalized skipping cost is about 0.2 basis points and the standard deviation of normalized skipping costs is about 0.5 basis points (i.e., 0.005%). Consistently with Figure 3, these results suggest that the statistical properties of normalized skipping costs are similar for each of the three currency pairs.

When considering the raw skipping costs r_k (i.e., without normalization to account for the mid-price), the mean skipping costs range from about 1.8 ticks (for EUR/GBP) to about 3.0 ticks (for GBP/USD). Given that the tick size for each of the three currency pairs is 0.00001 units of the counter currency (see Section 3.2), these skipping costs correspond to a mean additional cost of about £18.00 and \$30.00, respectively, for an institution submitting

⁹The survivor function is given by 1 minus the ECDF.

Table 2: Summary statistics for the normalized skipping costs \tilde{r}_k (in basis points) of EUR/USD, GBP/USD, and EUR/GBP trades on Hotspot FX during May–June 2010.

	EUR/USD	GBP/USD	EUR/GBP
Minimum	0	0	0
Median	0	0	0
Maximum	30.31	9.65	5.62
Mean	0.19	0.21	0.21
Standard Deviation	0.46	0.43	0.45

a market order with a size of 1 million units (which is the modal market order size for each of the three currency pairs). Although these mean skipping costs are relatively small, some trades in our sample have much larger skipping costs. The largest skipping cost that we observe exceeds 30 basis points, and would correspond to incurring a total additional cost of about \$3630.00 when submitting a trade of size 1 million euro.

On Hotspot FX, institutions can infer the approximate skipping cost of their trades by comparing their local bid- or ask-price (which they observe from the filtered set of limit orders that they observe on the platform) to the prices at which other trades have recently occurred (which they observe via their trade-data stream; see Section 2.4). Given that this is the case, why do some institutions perform trades that have extremely large skipping costs? We believe that the answer to this question lies in the fact that Hotspot FX serves a wide variety of institutions with varying levels of access to other trading mechanisms, such as direct telephone trading or voice brokers. At times when submitting a market order would entail a considerable skipping cost, large institutions would likely instead perform the same trade via another mechanism. By contrast, small institutions rarely have access to these other trading mechanisms, so they may have little option other than to accept large skipping costs as a cost of their trading. In essence, we argue that the large heterogeneity in skipping costs that we observe is a consequence of the substantial heterogeneity in the types and sizes of institutions that trade on Hotspot FX.

4.3 Price Changes

Our results in Section 4.2 reveal that the skipping costs borne by institutions vary considerably across the trades in our sample. The existence of some trades with a normalized skipping cost of several basis points suggests that, due to their CCLs, some institutions have access to a relatively small fraction of the liquidity available on the platform. This observation raises the question of how strongly CCLs impact the price changes between successive trades (i.e., the f_k series). This question is important because if different institutions pay considerably different prices for the same asset at a similar time (as our results in Section 4.2 suggest is the case), then the trade-price series could include large fluctuations that do not reflect similar changes in the asset’s fundamental value. Therefore, the price-formation process as observed on a platform that implements CCLs could be rather different to that on another platform in which all institutions can trade with all others.

To assess this question empirically, we introduce the following decomposition of each term in the f_k series into two constituent parts. Using the notation that we introduced in Section 4.1, the *change in quote price* g_k between the k^{th} trade and the previous trade in

the same direction is given by

$$g_k = \begin{cases} a_k - a_{k'}, & \text{if the } k^{\text{th}} \text{ trade is a buyer-initiated trade,} \\ b_{k'} - b_k, & \text{if the } k^{\text{th}} \text{ trade is a seller-initiated trade.} \end{cases} \quad (5)$$

Given a pair of trades k and k' , we can similarly calculate the *change in skipping cost*

$$h_k = \begin{cases} (p_k - a_k) - (p_{k'} - a_{k'}), & \text{if the } k^{\text{th}} \text{ trade is a buyer-initiated trade,} \\ (b_{k'} - p_{k'}) - (b_k - p_k), & \text{if the } k^{\text{th}} \text{ trade is a seller-initiated trade.} \end{cases} \quad (6)$$

The following identity provides a useful relationship between the three price-change series. For buyer-initiated trades, observe that

$$\begin{aligned} g_k + h_k &= (a_k - a_{k'}) + ((p_k - a_k) - (p_{k'} - a_{k'})) \\ &= p_k - p_{k'} \\ &= f_k. \end{aligned}$$

Similarly, for seller-initiated trades,

$$\begin{aligned} g_k + h_k &= (b_{k'} - b_k) + ((b_{k'} - p_{k'}) - (b_k - p_k)) \\ &= p_{k'} - p_k \\ &= f_k. \end{aligned}$$

Therefore, for both buyer-initiated and seller-initiated trades, it holds that

$$f_k = g_k + h_k. \quad (7)$$

Equation (7) enables us to decompose each change in trade price into the corresponding constituent change in quote price and change in skipping cost. Specifically, Equation (7) says that the price change between any pair of trades in the same direction can be expressed as the sum of the price change of the best quotes and the change in skipping costs of the two trades. In this section, we perform several statistical comparisons of the f_k , g_k , and h_k series to quantify the relative impact of CCLs, versus that of quote revisions, on price changes.

In the left panel of Figure 4, we show a quantile–quantile (Q–Q) plot of the f_k series versus the g_k series. In this plot, the quantile points cluster tightly along the diagonal. This implies that the shape of the distribution of the f_k series is very similar to that of the g_k series. In the right panel of Figure 4, we show a Q–Q plot of the f_k series versus the h_k series. As the plot illustrates, the distribution of changes in skipping costs is more tightly concentrated around 0 than is the distribution of changes in trade prices. Intuitively, this suggests that the changes in skipping cost account for only a small fraction of the total price change between successive trades.

To assess whether the similarities and differences highlighted in Figure 4 also hold at the trade-by-trade level (and not only at the level of the unconditional distributions), we also construct scatter plots of the individual terms of the series. In the left column of Figure 5, we show scatter plots of the f_k series versus the corresponding terms of the g_k series. For GBP/USD and EUR/GBP, the points cluster strongly along the diagonal, which indicates that for each trade, the change in trade price is very similar to the change in quote price. Some points in the EUR/USD plot occur away from the diagonal, but the vast majority of




Figure 4: Quantile–quantile (Q–Q) plots for (left) changes in trade price f_k versus changes in quote price g_k , and (right) changes in trade price f_k versus changes in skipping costs h_k for (green circles) EUR/USD, (orange squares) GBP/USD, and (purple triangles) EUR/GBP trades on Hotspot FX during May–June 2010. In each plot, the points indicate the 0.01, 0.02, \dots , 0.99 quantiles of the empirical distributions. The dashed black lines indicate the diagonal.

data points still cluster along the diagonal. In the right column of Figure 5, we show scatter plots of the f_k series versus the corresponding terms of the h_k series. In contrast to the plots of f_k versus g_k , these plots do not reveal any visible relationship between the f_k and h_k series for any of the three currency pairs. This finding suggests that application of CCLs simply manifests as additive noise in the price-formation process.

To examine the relationships between the f_k , g_k , and h_k series across all trades in our sample, we also calculate the sample correlation ρ between these series (see Table 4.3). For each of the three currency pairs, the sample correlation between the f_k and g_k series is very close to 1, with a very small standard error. This quantifies the strong relationship between changes in trade price and changes in quote price (see the left panels of Figure 5) and suggests that changes in trade price are strongly correlated with corresponding changes in the underlying quotes. By contrast, the sample correlations between the f_k and h_k series are very close to 0, and they have similar orders of magnitude to the corresponding standard errors. This provides further evidence that changes in skipping cost are uncorrelated with changes in trade price.

Together, our results in this section suggest the following interpretation of Equation (7). Mechanically, each change in trade price consists of two components: a change in the underlying quotes and a change in the skipping cost. The change in trade price is strongly correlated with the change in quotes, but has little or no correlation with the change in skipping cost. Therefore, although the change in skipping cost sometimes constitutes a considerable fraction of the total change in trade price, this impact manifests itself as uncorrelated noise in the trade-price series.

From an economic perspective, the strong positive correlation between the f_k and g_k series and the absence of significant correlation between the f_k and h_k series suggests that fundamental revaluations in trade prices arise due to corresponding changes in the best

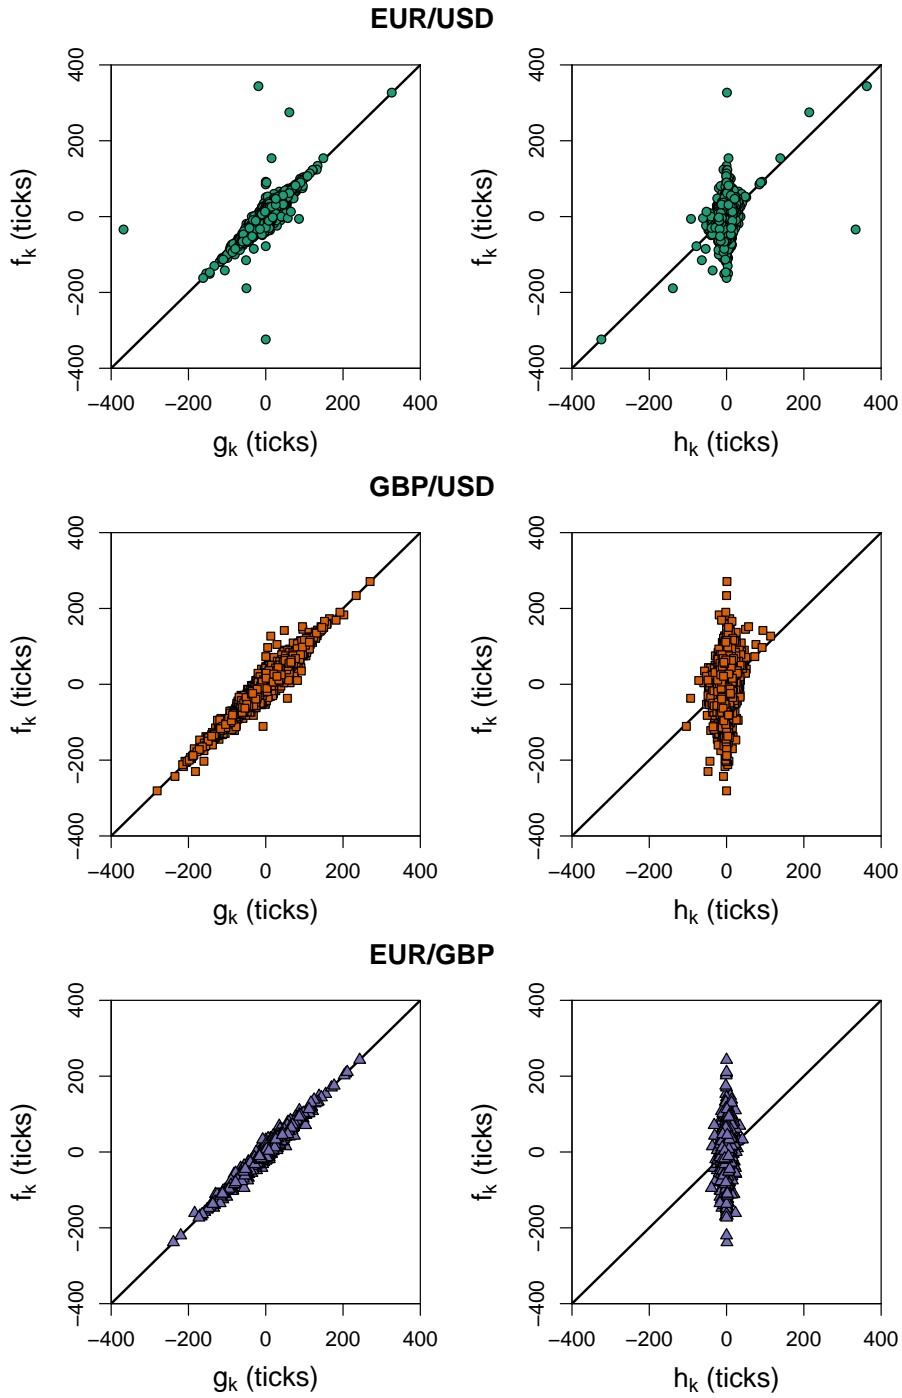


Figure 5: Scatter plots of changes in trade price f_k versus (left column) changes in quote price g_k and (right column) changes in skipping costs h_k for (top row) EUR/USD, (middle row) GBP/USD, and (bottom row) EUR/GBP trades on Hotspot FX during May–June 2010. The solid black lines indicate the diagonal.

	EUR/USD	GBP/USD	EUR/GBP
Panel A: f_k versus g_k			
	0.94	0.97	0.99
ρ	(< 0.01)	(< 0.01)	(< 0.01)
Panel B: f_k versus h_k			
	-0.04	0.00	0.00
ρ	(0.03)	(< 0.01)	(0.02)

Table 3: Sample correlation ρ between (Panel A) changes in trade price f_k and changes in quote price g_k and (Panel B) changes in trade price f_k and changes in skipping cost h_k for EUR/USD, GBP/USD, and EUR/GBP trades on Hotspot FX during May and June 2010. The numbers in parentheses are the corresponding standard errors, which we estimate by calculating the sample standard deviation of ρ across 10000 bootstrap samples of the data.

quotes. The f_k series can be regarded as a noisy observation of g_k series, where the uncorrelated, additive noise is caused by the restriction of institutions' trading activities to their bilateral trading partners. The strength of this effect varies across institutions due to the heterogeneity in their CCLs.

4.4 Volatility

Our results in Section 4.3 suggest that changes in trade price have little or no correlation with changes in skipping cost. However, given that some trades have large skipping costs (see Section 4.2), it is still possible that the volatility in the trade-price series differs significantly from the corresponding volatility of the underlying quotes. In this section, we assess the extent to which this is the case.

Recall from Section 4.2 that if the CCLs on a given platform were such that all institutions could access all trading opportunities, then all trades would occur at the best quotes at their time of execution, so f_k would equal g_k for all k . In this case, any volatility estimate would produce the same result when applied to either of these series. However, because CCLs restrict institutions' access to liquidity, this is not the case in general. By comparing the realized volatility of the trade-price series with the realized volatility of the corresponding quote-price series, it is possible to quantify the difference between the volatility in the prices that institutions actually pay for their trades to the underlying volatility observable in the platform-wide best quotes.

In contrast to studying the prices of individual trades, for which the application of CCLs always creates a non-negative additional cost (due to the addition of the corresponding non-negative skipping cost), it is not clear *a priori* whether the application of CCLs will cause the volatility of the trade-price series to be greater than or less than the volatility of the corresponding quote-price series. On the one hand, it is possible for quote prices to remain stable while trade prices fluctuate, because different institutions have access to different trading opportunities. In this case, the volatility of the trade-price series would be greater than that of the quote-price series. On the other hand, it is possible for quote prices to fluctuate while trade prices remain stable, because not all trade owners have access to the trading opportunities offered at the best quotes. In this case, volatility of the trade-price series would be less than that of the quote-price series. The aim of this section is to estimate the realized volatilities of these series and to determine which of these two possibilities occurs

on Hotspot FX.

For each currency pair each day, we estimate the sell-side trade-price volatility v_A , the sell-side quote-price volatility v_a , the buy-side trade-price volatility v_B , and the buy-side quote-price volatility v_b by calculating the quadratic variation of each process, sampled at regularly spaced intervals in trade time¹⁰ and sub-sampled at regularly spaced offsets. We consider buyer-initiated and seller-initiated trades separately to eliminate bid–ask bounce, which could obscure our results (see Section 3.3). For a detailed discussion of this methodology, see the Appendix.

Following Liu et al. [2015], we present our results for $K = 108$ regularly spaced sampling intervals each day (which corresponds to a mean interval length of 5 minutes, when measured in calendar time) and sub-sampled at $L = 10$ regularly spaced offsets. We also repeated all of our calculations for a variety of different interval numbers, ranging from $K = 50$ to $K = 500$, and a variety of different numbers of sub-sampling offsets, ranging from $L = 5$ to $L = 20$. We found our results to be qualitatively similar in all cases.

In Figure 6, we show scatter plots of the quote-price volatility versus the trade-price volatility for each day in our sample. For all three currency pairs, and for both buyer-initiated and seller-initiated trades, the points on the scatter plots cluster tightly on the diagonal. To help quantify the strength of this relationship, we also calculate the corresponding sample correlation coefficients, measured across all 30 days in our sample (see Table 4.4). In all cases, the sample correlation is very close to 1 and has very small standard error. Together, these results imply that each day’s quote-price volatility is very similar to the corresponding trade-price volatility.

	EUR/USD	GBP/USD	EUR/GBP
Panel A: v_B versus v_b			
ρ	0.997 (< 0.01)	0.985 (< 0.01)	0.986 (< 0.01)
Panel B: v_A versus v_a			
ρ	0.997 (< 0.01)	0.993 (< 0.01)	0.985 (0.01)

Table 4: Sample correlation ρ between realized trade-price volatility versus realized quote-price volatility for (Panel A) seller-initiated and (Panel B) buyer-initiated EUR/USD, GBP/USD, and EUR/GBP trades on Hotspot FX during May and June 2010. The numbers in parentheses are the corresponding standard errors, which we estimate by calculating the sample standard deviation of ρ across 10000 bootstrap samples of the data.

In Figure 6, some points lie slightly above the diagonal, while others lie slightly below the diagonal. To assess the deviation from the diagonal, we calculate the log-ratio

$$z = \begin{cases} \log(v_B/v_b) & \text{for seller-initiated trades,} \\ \log(v_A/v_a) & \text{for buyer-initiated trades,} \end{cases} \quad (8)$$

for each currency pair each day. A positive value of z denotes that the trade-price volatility exceeds the quote-price volatility, while a negative value of z denotes that the quote-price

¹⁰In this way, the number of seconds between successive samplings of the price series varies according to the inter-arrival times of trades. We also repeated all of our calculations by sampling the same series at regularly spaced intervals in calendar time, and we found our results to be qualitatively similar to those that we obtain with regularly spaced intervals in trade time.

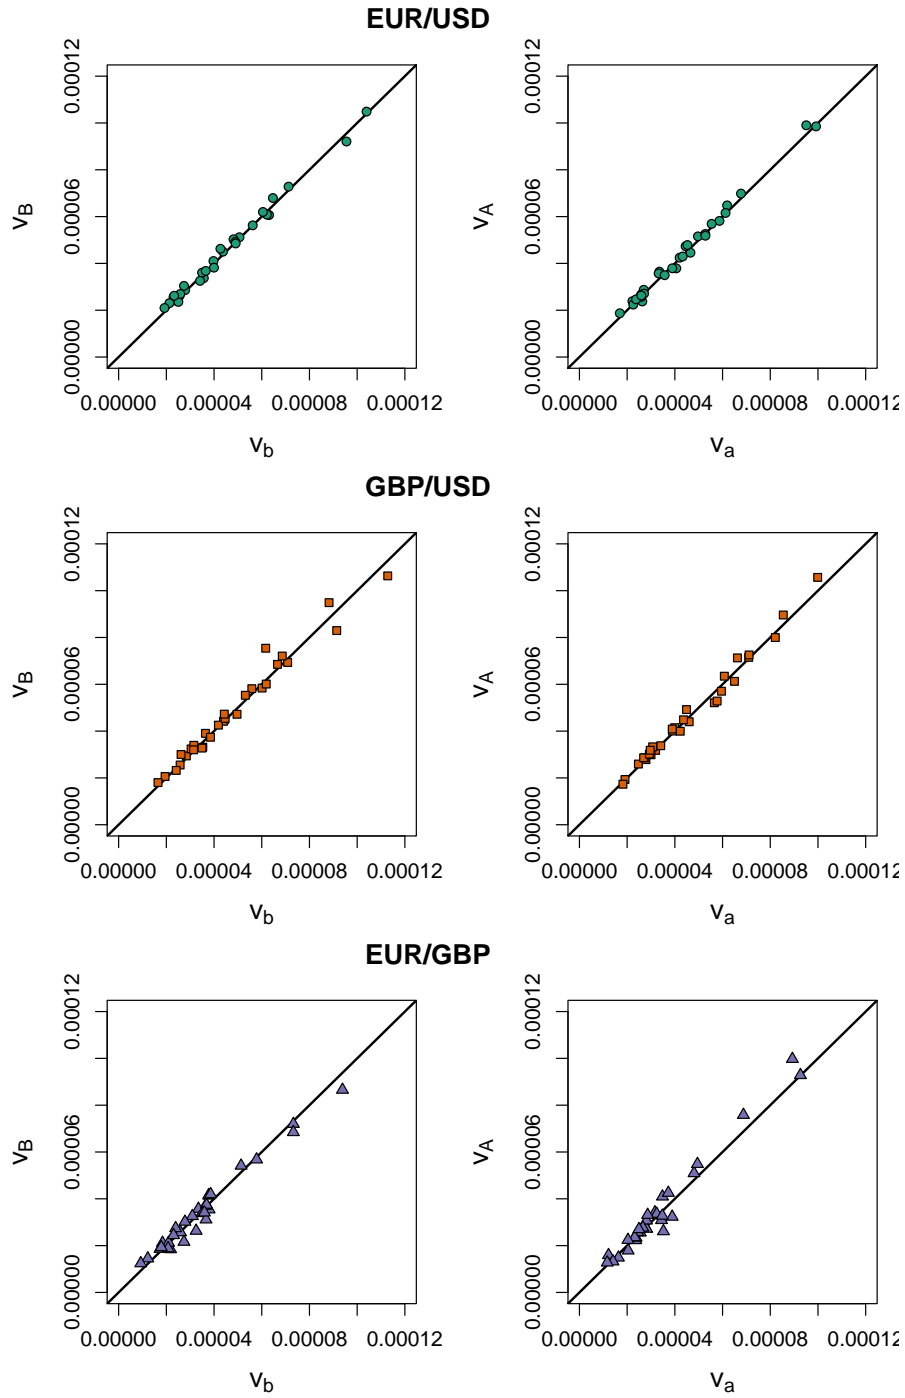


Figure 6: Scatter plots of realized trade-price volatility versus realized quote-price volatility for (left column) seller-initiated and (right column) buyer-initiated (top row) EUR/USD, (middle row) GBP/USD, and (bottom row) EUR/GBP trades on Hotspot FX during May–June 2010. The solid black lines indicate the diagonal.

volatility exceeds the trade-price volatility. In Table 4.4, we list the mean and standard deviation of z for buyer-initiated and seller-initiated trades for each of the three currency pairs. In all cases, the mean value of z is slightly positive, which suggests that, on average, the trade-price volatility on a given day typically exceeds the corresponding quote-price volatility. However, the magnitude of this effect is less than one standard deviation, which suggests that its strength is very weak.

	EUR/USD	GBP/USD	EUR/GBP
Panel A: Seller-Initiated Trades			
z	0.022 (0.049)	0.020 (0.064)	0.010 (0.116)
Panel B: Buyer-Initiated Trades			
z	0.015 (0.045)	0.012 (0.050)	0.022 (0.114)

Table 5: Mean values of z for (Panel A) seller-initiated trades and (Panel B) buyer-initiated trades for EUR/USD, GBP/USD, and EUR/GBP trades on Hotspot FX across all trading days in our sample. The numbers in parentheses are the corresponding standard deviations.

Together, our results in this section reveal that on average, the realized volatility of the trade-price series slightly exceeds that of the quote-price series, but only by a small margin. Therefore, the vast majority of volatility in the trade-price series is also directly observable in the quote-price series. This implies that the volatility observable in both the quote-price and trade-price series is dominated by a common, underlying volatility that supersedes the idiosyncratic impact of CCLs on the trade-price series.

5 A Model of Trade with CCLs

In Section 4, we used data from Hotspot FX to analyze how CCLs affect trade in a real market. Studying historical data is, however, only one aspect of understanding how this mechanism might affect financial markets more generally, because such analysis does not provide insight into how these results might change were institutions to make substantial modifications to their CCLs. Because the underlying network of CCLs in a real market is fixed and unobservable to us, empirical analysis does not provide a way to address this important question. In this section, we therefore complement our empirical work by introducing and studying a model of trade in which institutions assign CCLs to their trading counterparties.

In our model, each institution updates its buy and sell prices for a single asset and performs a trade whenever it identifies a trading counterparty offering to buy or sell at a mutually agreeable price. A crucial feature of our model is that not all institutions can trade with all others; instead, each institution can only trade subject to its CCLs, at prices that depend not only on other institutions' buy and sell prices, but also on the underlying network of bilateral CCLs between them.

Our motivation is to provide a direct link between the institutions' interaction topology (via the network of CCLs) and the corresponding trade prices. By studying quantitative relationships between the CCL network and the trade-price series, while holding all other model parameters constant, we are able to assess how restricting the trading opportunities available to specific institutions affects both the prices of individual trades and the volatility

of the trade-price series. Modelling approaches that ignore the heterogeneous impact of CCLs would not facilitate such investigation, so our approach is a natural framework for studying this problem.

The setting for our model is an infinite-horizon, continuous-time market populated by a set of N institutions $\Theta = \{\theta_1, \theta_2, \dots, \theta_N\}$ that trade a single risky asset. Each institution $\theta_i \in \Theta$ maintains a private buy-valuation B_t^i and a private sell-valuation A_t^i . The values of B_t^i and A_t^i vary across the different institutions to reflect differences in the institutions' opinions regarding the likely future value of the asset, as well as differences in their inventory, cash-flow, financing constraints, and so on. To focus on the impact of CCLs without considering the impact of strategic activity (which could make our results more difficult to interpret), we model these prices with stylized stochastic processes.

For each institution θ_i , we rewrite the buy and sell prices in terms of a mid-price $M_t^i = (B_t^i + A_t^i)/2$ and spread $s_t^i = A_t^i - B_t^i$, so that

$$B_t^i = M_t^i - \frac{s_t^i}{2}, \quad A_t^i = M_t^i + \frac{s_t^i}{2}.$$

We describe the dynamics of the spread in detail below; for now, we only remark that the values of s_t^i are constrained such that they are never smaller than some minimum value $s_0 > 0$.

5.1 Temporal Evolution

Before simulating the temporal evolution of our model, we choose an initial state in which no trading is possible. We provide details of this initialization in Section 5.5. Leaving aside the behaviour of s_t^i and M_t^i at trade times, which we describe in Section 5.3, we assume that between trades the s_t^i are governed by

$$ds_t^i = -\kappa (s_t^i - s_0) dt, \tag{9}$$

for some constant $\kappa > 0$, and that the M_t^i are governed by

$$dM_t^i = \gamma M_t^i dW_t^{M^i}, \tag{10}$$

where $\gamma > 0$ is the mid-price volatility (with units $(\text{time})^{-\frac{1}{2}}$) and $W_t^{M^i}$ are mutually independent Brownian motions.¹¹

Equation (9) causes each institution's spread to approach its minimum value s_0 as time increases after a trade. Were there no trading, the processes M_t^i would be geometric Brownian motions with no drift. This model minimizes the complications associated with the mixing of price and time scales in the model parameters; indeed, the temporal evolution is directly influenced only by the parameter γ (whose inverse defines a timescale) and by trading via the network of CCLs. Although a geometric Brownian motion with no drift has constant mean, its variance increases with time. In the absence of trading, our mid-prices would therefore disperse and spread out progressively and indefinitely over time. As we see in Section 5.3, however, the occurrence of trades ensures that prices remain grouped together. By using the same values of γ for each institution, we ensure that the behavior of each institution is ex-ante identical apart from their different access to trading opportunities because of the heterogeneous CCL network. We now discuss this feature in more detail.

¹¹A possible refinement of the model would be to include a common market factor w_T^m in addition to the idiosyncratic noise terms, writing $dM_t^i = \gamma M_t^i (\rho_i dW_t^m + \sqrt{1 - \rho_i^2} dW_t^{M^i})$, where ρ_i is a (common) correlation coefficient. This adds little to our analysis, however.

5.2 The CCL Network

We assume that each institution θ_i assigns a CCL to each other institution θ_j . In order to perform a time-stationary analysis of our model, we assume that each institution's access to trading opportunities does not vary over time, and, specifically, that it does not vary according to the trading history. For each pair of institutions θ_i and θ_j , we model the bilateral CCL with a binary indicator: either θ_i and θ_j are trading partners or they are not. For simplicity, we allow trading partners to trade arbitrarily large amounts.

The CCLs can thus be represented via an undirected network in which the nodes represent the institutions and the edges represent the extant bilateral credit relationships: θ_i and θ_j can trade with each other if and only if the edge $\theta_i \leftrightarrow \theta_j$ exists in the network. We call such a network a *CCL network*. We show some example CCL networks in Figure 8.

Although using a binary CCL network represents a simplification of the impact of CCLs in real markets, it enables us to focus on the impact of CCLs in a simple yet adequately realistic equilibrium-pricing framework. However, we note one constraint imposed by our model: the CCL network must be connected. If this is not the case, then disconnected components can drift apart as there is no means to link them together.

5.3 Trading

We assume that a trade occurs at each time t^* such that a pair of institutions θ_i and θ_j with a bilateral CCL have prices that satisfy $B_{t^*}^i = A_{t^*}^j$. We call this price the *trade price*.

Recall that in our empirical analysis in Section 4, we classify trades according to whether they are buyer-initiated or seller-initiated. To aid comparisons between the output of our model and our empirical results, we implement the following simple rule to classify trades. Let

$$\bar{M}_t = \frac{1}{N} \sum_{i=1}^N M_t^i \quad (11)$$

denote the empirical mean of the N institutions' mid-prices at time t . Consider a trade that occurs between θ_i and θ_j at time t^* and with trade price $p = B_{t^*}^i = A_{t^*}^j$. If $p > \bar{M}_t$, we label this trade as buyer-initiated, and we call θ_i the *initiator* and θ_j the *acceptor*. Otherwise, we label this trade as seller-initiated, and we call θ_j the initiator and θ_i the acceptor.

For each trade, we think of the initiator as having submitted a market order at the trade price, and we think of the acceptor as having owned a limit order at the trade price, which is then matched by this market order. The initiator trades at a relatively unfavorable price, and the fewer bilateral CCLs this institution has, the further this price is expected to be from \bar{M}_t . We thus mimic the relative competitive disadvantage of poorly connected institutions.

Whenever a buyer-initiated (respectively, seller-initiated) trade occurs, we record the lowest price among all institutions' sell prices (respectively, the highest price among all institutions' buy prices), in order to calculate the skipping cost of the trade. We then adjust the mid-prices of θ_i and θ_j by subtracting $s_0/2$ from $M_{t^*}^i$ and adding $s_0/2$ to $M_{t^*}^j$ (respectively, subtracting $s_0/2$ from $M_{t^*}^j$ and adding $s_0/2$ to $M_{t^*}^i$) and widening each of s_t^i and s_t^j by $s_0/2$. Finally, we re-set the values of $B_{t^*}^i$, $A_{t^*}^i$, $B_{t^*}^j$, and $A_{t^*}^j$ according to these new mid-prices and spreads. This has the effect of separating $B_{t^*}^i = A_{t^*}^j$ by $3s_0/2$. All other prices (including $B_{t^*}^j$ and $A_{t^*}^i$) remain unchanged. This feature models a decrease in trading desire from the initiator and acceptor, due to the execution of the trade. At a technical level, widening the spread removes the undesirable possibility of the initiator's price and acceptor's price being equal infinitely often in an arbitrarily small interval after they first meet.

We now see how trading stops the mid-prices from spreading out. If the mid-prices of θ_i and θ_j diverge by the average of their spreads, then the buy price of one trader meets the sell price of another from below and a trade is triggered. The mid-prices are then moved back together by s_0 , reversing the previous divergence. The spreads then revert towards s_0 , which has the effect of stopping them from growing indefinitely as trades occur.

5.4 Adjustments for Discrete Time-Stepping

We simulate the evolution of our model in discrete time, with a time step $\Delta t > 0$, using a simple explicit (Euler–Maruyama) differencing. In general, this implementation produces an overshoot before we detect that a trade should take place. Therefore, whenever a buyer-initiated trade occurs between a buyer θ_i and a seller θ_j , we actually observe $B_t^i > A_t^j$, rather than $B_t^i = A_t^j$. In the simplest (and, for small spreads, generic) case, no other relevant prices are sandwiched between the buy and sell prices in question. Whenever this happens, we deem a trade to have taken place at the end of the time step and at a price equal to the mean of B_t^i and A_t^j .

In a small number of cases, the overshoot caused by explicit differencing may be so large as to create more than one trading opportunity. For example, the price moves that occur in a discrete time step might cause θ_i 's buy price to exceed the sell prices of both θ_j and a third institution θ_k . In such a case, we first deal with the trade that occurs furthest from \bar{M}_t . After recording this trade and changing the buyer's and seller's mid-price, spread, buy price and sell price (see Section 5.3), if other trading opportunities exist, we then process the one whose trade price is furthest from the updated \bar{M}_t , and so on until there are no further trading opportunities to consider.

5.5 Parameter Choices and Implementation

The primary aim of our model is to understand how CCLs impact liquidity, trade prices, and volatility. We therefore hold the values of γ , κ and s_0 fixed, and study how our model's output varies when we vary the edges in the CCL network.

We first set $\bar{M}_0 = 1$ and $s_0 = \epsilon M_0$, where we take the dimensionless parameter ϵ to be 0.001, implying spreads of about 0.1%. We choose $\kappa = 1$, which sets the (otherwise arbitrary) time unit as $1/\kappa = 1$, and we choose the volatility $\gamma = \epsilon\sqrt{\kappa} = 0.001$ to balance the changes in the spread and the changes in the mid-price.

We initialize the mid-prices M^i by drawing them randomly from a normal distribution with mean \bar{M}_0 and standard deviation ϵ , and we set all the spreads equal to s_0 . We then run the trade-processing algorithm described in Section 5.3 to adjust the mid-prices and spreads of all institutions for whom this initial state would cause trading to occur. We repeat this trade-processing step until no trading opportunities remain (i.e., until $B_{t_0}^i < A_{t_0}^j$ for each pair of institutions such that $\theta_i \leftrightarrow \theta_j$).

The final parameter in our model is the discretization time step. The dominant term in the discrete evolution of the system is the noise term, which in relative terms (i.e., relative to the value of the relevant quantity at the beginning of the time step) is $O(\gamma\sqrt{\Delta t})$. Accurate discretization of the stochastic processes requires that this term be small. Moreover, we wish to avoid the situation where the discrete time steps regularly create multiple simultaneous trading opportunities. If we expect the separation of the mid-prices to be $O(\epsilon M_0/N)$, we require that $\gamma\sqrt{\Delta t}$ is smaller than this. We therefore take $\Delta t = 1/3N^2$. This choice of Δt is also sufficiently small that errors associated with the numerical integration of the stochastic differential equations are negligible.

For the results that we present in Section 5.7, we use $N = 128$ institutions. We also repeated our simulations with several different choices of N in the range 100 to 1000 (and with appropriately modified values of Δt), and we found that our results were qualitatively similar in each case. For each CCL network that we study, we simulate the temporal evolution of our model from $t = 0$ to $t = 10$. We then discard all activity before $t = 2$ as a burn-in period, to remove the transient behaviour that occurs before the model settles into its equilibrium state. We verified that these choices were sufficiently large by examining the numerical stability of our results using a range of different burn-in periods and total time lengths. We found that our results were numerically stable for all burn-in periods longer than about $t = 1$ (which, for the parameter choices that we use in our simulations, is the time scale for reversion of the B_t^i and A_t^i), and for all total lengths greater than about $t = 2$.

In Figure 7, we show a single simulation run of the model, using a simple CCL network and the parameter choices listed in Section 5.5. As the figure illustrates, the prices of subsequent trades can deviate from each other considerably. Therefore, the model does a good job at capturing how heterogeneity in institutions' access to trading opportunities (which arises as a direct consequence of their CCLs) can manifest in the trade-price series. We now turn to a more quantitative analysis of how this effect is related to the CCL network.




Figure 7: An example simulation of the model, using a CCL network with $N = 3$ institutions in which $\theta_1 \leftrightarrow \theta_2$ and $\theta_1 \leftrightarrow \theta_3$, but in which θ_2 and θ_3 cannot trade with each other. The blue curves indicate the institutions' private buy valuations B_t^i and the red curves indicate the institutions' private sell valuations A_t^i . The black crosses indicate trades.

5.6 CCL Networks

For any CCL network with N nodes, the maximum possible number of edges is $N(N-1)/2$. Therefore, a CCL network with N nodes and n edges has *edge density*

$$d = \frac{2n}{N(N-1)}. \quad (12)$$

Our model can be used to investigate any CCL network with any number of nodes and any configuration of edges. However, to highlight the most salient features of our results, we restrict our discussion to two classes of networks with specific topological structures.

The first class of CCL networks that we discuss are *Erdős–Rényi random networks* [Erdős and Rényi, 1960], in which a specified number of edges are placed uniformly at random between pairs of nodes. We use this class of networks to model a market in which institutions choose their trading partners uniformly at random. Although this assumption is likely to be a poor reflection of how institutions set CCLs in a real market, studying this class of networks enables us to investigate the temporal evolution of our model in a simple, stylized framework with no deterministic structure.

To construct our Erdős–Rényi random networks, we fix a choice of d , then use Equation (12) to calculate the required number of edges n to produce this edge density. We place these n edges uniformly at random among the N nodes of the network, then check whether each node is connected to at least one other node. If so, we accept the CCL network; if not, we reject the CCL network and re-draw an alternative network using the same rules.¹² For a given choice of d , we create a sample of 1000 such CCL networks, then simulate 1000 independent runs of the model for each of these 1000 CCL networks, where each run uses a different initial seed for the pseudo-random number generator. The left panel of Figure 8 shows a schematic to illustrate a single random draw of an Erdős–Rényi random network with $N = 12$ institutions and $n = 14$ edges.

The second class of CCL networks that we consider are *core–periphery networks* (see Csermely et al. [2013]). Core–periphery networks consist of two groups of nodes: core nodes and periphery nodes. All core nodes are connected to all other core nodes. All periphery nodes are connected to exactly one core node, such that the degree of no two core nodes differs by more than 1. The right panel of Figure 8 shows a schematic to illustrate our construction of a core–periphery CCL network with $N = 12$ institutions.

We use core–periphery networks to model a market in which a core group of institutions assign each another very high CCLs, but in which all other institutions only have a credit line with one large institution within the core. Several recent studies of market organization have suggested that many large financial markets have an approximate core–periphery structure, in which the core consists of large, international banks and the periphery consists of smaller financial institutions such as small banks, hedge funds, or mutual funds (see, e.g., Craig and von Peter [2014] and Fricke and Lux [2012]). Therefore, our core–periphery structure represents an approximation of the complex structure of real markets, albeit simplified to a convenient deterministic form.

To construct our core–periphery networks, we first fix the fraction ψ of periphery nodes. When $\psi = 0$, all institutions are core institutions, so the CCL network is complete and all institutions are able to trade with all others. For a given choice of ψ (and therefore of d), we create a single CCL network, then simulate 1000 independent runs of the model for this

¹²Recall from Section 5.2 that we restrict our attention to cases in which the CCL network is connected, to prevent disconnected components from drifting apart.




Figure 8: Schematics of CCL networks. (Left panel) An Erdős–Rényi random network with $N = 12$ institutions and $n = 14$ edges, and (right panel) a core–periphery network with $N = 12$ institutions, of which 3 nodes are core nodes (shown in dark grey) and 9 nodes are periphery nodes (shown in light grey).

CCL network, where each run uses a different initial seed for the pseudo-random number generator.

5.7 Results

We now present our simulation results for the two classes of CCL network topologies described in Section 5.6. We studied buyer-initiated and seller-initiated trades separately by implementing the trade-classification algorithm described in Section 5.3. In line with our expectations (due to the symmetry of buyers and sellers in our model), our results are qualitatively the same for buyer-initiated and seller-initiated trades. To increase the size of our sample, we present our results for all trades together (i.e., we aggregate buyer-initiated and seller-initiated trades).

In Figure 9, we plot the number of trades that occur for each edge density d . For both Erdős–Rényi random networks and core–periphery networks, CCL networks with lower edge densities produce fewer trades. The intuition is simple: the lower the edge density, the lower the number of bilateral trading partners in the population, and therefore the lower the number of trades that occur within a given time horizon. Figure 9 also illustrates that the mean number of trades that occur for a given CCL network depends on the network topology, not just its edge density. Specifically, the mean number of trades that occur in a core–periphery network is much smaller than the mean number of trades that occur in an Erdős–Rényi random network with the same edge density. This result is interesting from a practical perspective because it suggests that the influence of CCLs depends not only on how many trading partners each institution has, but also on who those trading partners are.

In Figure 10, we plot the mean skipping costs among the trades that occur for each edge density d . For both Erdős–Rényi random networks and core–periphery networks, CCL networks with lower edge densities produce higher skipping costs. Intuitively, this result illustrates that the more restrictive CCLs are to institutions’ access to trading opportunities, the greater the skipping costs institutions pay for their trades.

Figure 10 again illustrates that the CCL network’s topology, and not just its edge density, plays an important role in market dynamics. For any edge density d , the mean skipping cost among trades for an Erdős–Rényi random network is lower than the mean skipping cost among trades for a core–periphery network.

For both classes of networks, the mean skipping cost decreases rapidly as the edge density




Figure 9: Mean number of trades that occur for the (green) Erdős–Rényi random networks and (purple) core–periphery networks. The solid curves indicate the mean across all independent runs of the model and the dashed curves indicate one standard deviation.




Figure 10: Mean skipping costs of trades for the (green) Erdős–Rényi random networks and (orange) core–periphery networks. The solid curves indicate the mean across all independent runs of the model and the dashed curves indicate one standard deviation.

increases from 0 to about 0.1. For Erdős–Rényi random networks, the mean skipping cost is very close to 0 for all edge densities above about 0.3. In this case, the impact of CCLs on individual trade prices is very small. For core–periphery networks, the mean skipping cost remains much higher before eventually decreasing to 0 as the edge density reaches 1 (in which case the CCL network is complete, so all trades have 0 skipping cost by definition). Moreover, the standard deviation of skipping costs is much larger for core–periphery networks than for the Erdős–Rényi random networks.

In Figure 11, we plot the trade-price and quote-price volatility that occur for each edge density d . We calculate both types of volatility using the same methodology¹³ as described in Section 4.4.




Figure 11: Realized trade-price volatility for the (green) Erdős–Rényi random networks and (orange) core–periphery networks, and realized quote-price volatility for the given (blue) Erdős–Rényi random networks and (pink) core–periphery networks. The solid curves indicate the mean across all independent runs of the model and the dashed curves indicate one standard deviation.

For the Erdős–Rényi random networks, the trade-price volatility exceeds the quote-price volatility when the CCL network’s edge density is very low. As the edge density increases, the trade-price volatility decreases more quickly than the quote-price density. For edge densities larger than about 0.1, the trade-price volatility is approximately equal to the quote-price volatility.

Intuitively, this result makes sense because the quote-price volatility is determined by the maximum among all buy prices (respectively, the minimum among all sell prices) in the market, and therefore is influenced only by the extreme prices among the population. By contrast, the trade-price volatility depends on the prices of all trades conducted by

¹³When studying our model, we only considered event-time estimators of volatility. We did not repeat our calculations using calendar-time volatility estimators, because our choice of timescale is arbitrary.

all institutions. As the edge density across the population increases, the influence on the number of bilateral CCLs for the extremum institution is relatively small, because only a small fraction of all possible edges in the CCL network involve this institution. Therefore, its influence on the quote-price series is much smaller than its influence on the trade-price series, which is affected by all bilateral CCLs among all institutions.

For the core–periphery networks, the quote-price volatility is approximately stable across all edge densities, except for values of d close to 0 and 1. By contrast, the trade-price volatility first increases sharply as d increases slightly above 0, then decreases gradually, then decreases sharply as d increases beyond about 0.95. This result illustrates that for the core–periphery networks, the edge density has a much stronger influence on trade-price volatility than it does on quote-price volatility.

Recall from Equation (7) and Section 4.4 that trade-price volatility can be regarded as the underlying volatility observable in the best quotes, plus an additional contribution caused by CCLs. Therefore, consistently with our results for skipping costs (see Figure 10), Figure 11 suggests that as the edge density of the CCL network decreases, the strength of this additional impact on trade-price volatility from CCLs increases, for both Erdős–Rényi random networks and core–periphery networks.

In Figure 12, we plot the corresponding z ratios of the realized quote-price volatility and the realized trade-price volatility (see Equation (8)). For the Erdős–Rényi random networks, the z ratios are close to 0 (and sometimes even slightly negative) for edge densities greater than about 0.1, but positive for edge densities smaller than about 0.1. For the core–periphery networks, the z ratios are positive for almost all edge densities, and often have considerable a magnitude. In these cases, the application of CCLs causes the volatility in the trade-price series to exceed the volatility observable in the underlying quotes by a considerable margin. When this happens, only a very small fraction of the volatility in the trade-price series is explained by a corresponding volatility in the underlying best quotes.

5.8 Discussion of Simulation Results

Consistently with our expectations, we find that institutions trade less and pay greater skipping costs when their CCLs cause them to experience more severe restrictions in their access to trading opportunities. We also find that CCLs have a considerable impact on trade-price volatility, but much smaller impact on quote-price volatility. Therefore, as the CCLs become progressively more restrictive, the ratio z (which measures the relative magnitudes of the volatility in the trade-price and quote-price series) increases considerably. Moreover, by comparing different network structures, it is apparent that a CCL network’s topology, and not just its edge density, plays an important role in the determining the extent to which CCLs impact trade and volatility.

Aside from the decreased frequency of trading, there is no difference in the dynamic evolution of our model for different choices of CCL networks. Therefore, the strong variability in z ratios that we observe in Figure 12 is directly attributable to the impact of CCLs in our model. This result is very interesting from a practical perspective, because it suggests that even if CCLs do not impact institutions’ trading strategies, they can still strongly influence the price-formation process, simply because they influence institutions’ access to the total available liquidity pool.

Together, our model simulations suggest that when CCLs cause severe restrictions to institutions’ access to liquidity, they can significantly impact both trade prices and volatility. Two features are particularly interesting. First, as the edge density of a CCL network falls,




Figure 12: Mean ratio z of the realized quote-price volatility and the realized trade-price volatility for the (green) Erdős–Rényi random networks and (orange) core–periphery networks. The solid curves indicate the mean across all independent runs of the model and the dashed curves indicate one standard deviation.

both the skipping costs of individual trades and the trade-price volatility increase. This increase is not accompanied by a similar increase in quote-price volatility, so the z ratio falls. Second, the impact of CCLs depends not only on the edge density, but also on the specific topology of the CCL network. Therefore, forecasting how a change in the edge density d will impact skipping costs and trade-price volatility also requires detailed knowledge of the CCL network’s topology. Intuitively, this result implies that understanding the possible impact of CCLs in financial markets requires not only knowledge of how many institutions are trading partners, but also of which institutions are trading partners with which others.

In situations where financial institutions’ CCLs severely restrict their access to liquidity, our results suggest that it is indeed plausible that CCLs could serve to exacerbate systemic risk, by forcing them to either trade at extremely unfavourable prices, or not trade at all. Moreover, our results highlight that the strength of this effect is not solely determined by the edge density of the CCL network, but also by its topology. This presents a difficult question for regulators: how can the state of the CCL network between institutions in real markets be monitored in real time? In a market that utilizes CCLs, any sensible forecast of volatility must incorporate this information, so understanding both the edge density and topology of this network will become an important step in many risk-management processes and option-pricing methodologies.

6 Conclusion

We have investigated how the application of CCLs impacts liquidity and trade in financial markets. We first addressed this issue empirically, by studying a novel data set from a large electronic trading platform in the FX spot market that utilizes CCLs to facilitate trade. Although we found that CCLs have little or no impact on most of the trades in our sample, we also found that CCLs contribute a considerable skipping cost for some trades. We argued that the substantial heterogeneity in skipping costs that we observe is a natural consequence of the heterogeneity among the type and size of institutions trading on the platform. By implementing CCLs, Hotspot FX can facilitate trade for a wide variety of different financial institutions while providing these institutions with the ability to decide for themselves whether or not to trade with specific counterparties. Because of this direct control of counterparty exposures, there is no need for the platform to set high barriers to entry for new participants. Indeed, the two most recent BIS Triennial Central Bank Surveys have both noted that a new trend for direct participation from small, non-bank institutions has been a key driver for sustained growth in FX volumes since the start of the decade [Bank for International Settlements, 2010, 2013]. Our findings are consistent with the hypothesis that a wide variety of different financial institutions, with varying access to liquidity, interact simultaneously on Hotspot FX.

We also considered how CCLs impact volatility. By decomposing price changes into two components – one attributable to changes in the best quotes, and the other attributable to contemporaneous changes in the impact of CCLs – we observed that the volatility observable in the trade-price series can be regarded as the underlying volatility observable in the best quotes, plus an additional contribution caused by CCLs. We found that CCLs contribute a small additional volatility to the trade-price series on Hotspot FX, when compared with the corresponding volatility observable in the best quotes.

To complement our empirical analysis, we also introduced a model of a single-asset market in which institutions assign CCLs to their trading counterparties. In our model, the network of CCLs provides explicit control over the interaction topology between different institutions. By holding the model parameters fixed and varying only this network of CCLs, we were able to study how CCLs impact trade in our artificial market. Our main observation is that both the edge density and topology of the CCL network play an important role in determining the skipping costs of trades and the corresponding volatility in the trade-price series. When the restrictions imposed by CCLs are particularly severe, they can cause large skipping costs and a high level of volatility in the trade-price series. This high trade-price volatility is not accompanied by a similarly high quote-price volatility.

Several possible extensions to our model could help to provide further insight into the impact of CCLs. For example, the temporal evolution of institutions' buy and sell prices could incorporate jumps or stochastic volatility, to more closely reflect the behaviour observed in real markets. There are also several possible ways to incorporate strategic considerations into the model. As one example, different institutions could implement different time-update rules for their buy and sell prices, to reflect heterogeneity in their trading styles. As another example, each institution could also choose how to update its buy and sell prices according to its CCLs. For example, institutions could be less willing to revise their buy price downward or their sell price upward if they can see from the price of recent trades that they are already likely to be paying a large skipping cost. We anticipate that these extensions will provide useful avenues for future research.

We believe that our results help to illuminate several important questions in regulatory

debate surrounding CCLs. To our knowledge, ours is the only empirical or theoretical study to address the use of this mechanism in a quantitative framework. Our empirical results indicate that the liquidity restrictions imposed by CCLs do not strongly impact the prices of the vast majority of trades on Hotspot FX. We therefore argue that the application of CCLs (and the consequent creation of skipping costs) may be regarded as a necessary consequence of providing direct market access to a widening selection of different financial institutions, rather than a weakness of this market design. However, our model simulations also suggest that as CCLs become progressively more restrictive, trade-price volatility can escalate rapidly. In such situations, it seems plausible that the application of CCLs could serve to exacerbate systemic risk by severely restricting institutions' access to trading opportunities. Therefore, much like CVAs and trade novation via a CCP (see Section 2.2), CCLs do not provide a simple solution to the problem of counterparty risk. However, our empirical results suggest that CCLs can constitute a sensible approach to this problem under normal market conditions, when liquidity is plentiful and CCLs are not too restrictive.

An important open question is whether (and how) CCLs might be implemented alongside other measures to mitigate counterparty risk. For example, it is possible that a platform could offer institutions the ability to apply CCLs, even if trades were still novated by a CCP. This configuration could, in principle, provide institutions with a double-layered protection: trades would still be novated by a CCP, but in the event of a CCP failure, institutions could ensure that they were only exposed to specified counterparties, and only up to a pre-specified limit. Before this configuration could be adopted, however, many important questions about the possible interactions between these two mechanisms would need to be addressed. How should trades novated by the CCP count towards a given institution's CCLs? Should they count at all? Or should the institutions also have a separate CCL directly with the CCP, to limit their exposure in the case of a CCP failure? Given the relatively low impact of CCLs that we observe on Hotspot FX, we strongly encourage further research in this area, to help address the many open questions on this topic and to improve understanding of this deeply interesting but hitherto unexplored market mechanism.

Acknowledgments

We thank Franklin Allen, Bruno Biais, Julius Bonart, Jean-Philippe Bouchaud, Yann Braouezec, Damiano Brigo, Rama Cont, Jonathan Donier, J. Doyne Farmer, Ben Hambly, Charles-Albert Lehalle, Albert Menkveld, Stephen Roberts, Cosma Shalizi, Thaleia Zariphopoulou, Ilija Zovko, and seminar participants at the 2012 Institut Louis Bachelier Market Microstructure Conference, the 2015 Financial Risk and Network Theory Conference, the University of Oxford, Imperial College London, the IESEG School of Management, Humboldt-Universität zu Berlin, the Brevan Howard Centre for Financial Analysis, and the Imperial–ETH Workshop for helpful comments and discussions. MDG gratefully acknowledges support from the James S. McDonnell Foundation, the Oxford–Man Institute of Quantitative Finance, and EPSRC (Industrial CASE Award 08001834).

References

T. G. Andersen, T. Bollerslev, F. Diebold, and P. Labys. Great realisations. *Risk Magazine*, 13(3):105–109, 2000.

- Bank for International Settlements. Settlement risk in foreign exchange markets and CLS bank. Technical report, Bank for International Settlements, available at http://www.bis.org/publ/qtrpdf/r_qt0212f.pdf, 2002.
- Bank for International Settlements. Triennial central bank survey: report on global foreign exchange market activity in 2010. Technical report, Bank for International Settlements, available at <http://www.bis.org/publ/rpfx10t.pdf>, 2010.
- Bank for International Settlements. Foreign exchange turnover in April 2013: preliminary global results. Technical report, Bank for International Settlements, available at <http://www.bis.org/publ/rpfx13fx.pdf>, 2013.
- Basel Committee on Banking Supervision. Supervisory framework for measuring and controlling large exposures. Technical report, Bank for International Settlements, available at www.bis.org/publ/bcbs283.pdf, 2014.
- M. Bech. FX volume during the financial crisis and now. *Bank for International Settlements Quarterly Review*, 1:33–43, 2012.
- B. Biais, F. Heider, and M. Hoerova. Clearing, counterparty risk, and aggregate risk. *IMF Economic Review*, 60(2):193–222, 2012.
- D. Brigo, M. Morini, and A. Pallavicini. *Counterparty Credit Risk, Collateral and Funding: With Pricing Cases for All Asset Classes*. Wiley, West Sussex, UK, 2013.
- G. Cesari, J. Aquilina, N. Charpillon, Z. Filipovic, G. Lee, and I. Manda. *Modelling, Pricing, and Hedging Counterparty Credit Exposure: A Technical Guide*. Springer, Berlin, Germany, 2010.
- B. Craig and G. von Peter. Interbank tiering and money center banks. *Journal of Financial Intermediation*, 23(3):322–347, 2014.
- P. Csermely, A. London, L. Y. Wu, and B. Uzzi. Structure and dynamics of core/periphery networks. *Journal of Complex Networks*, 1(2):93–123, 2013.
- P. Erdős and A. Rényi. On the evolution of random graphs. *Publications of the Mathematical Institute of the Hungarian Academy of Sciences*, 5:17–61, 1960.
- D. Fricke and T. Lux. Core–periphery structure in the overnight money market: evidence from the e-mid trading platform. *Computational Economics*, 45(3):359–395, 2012.
- K. Giesecke and S. Weber. Cyclical correlations, credit contagion, and portfolio losses. *Journal of Banking and Finance*, 28(12):3009–3036, 2004.
- Goldman Sachs. Response to enhanced prudential standards and early remediation requirements for covered companies. Technical report, The Goldman Sachs Group, available at <http://www.goldmansachs.com/our-thinking/public-policy/regulatory-reform/comment-letter.pdf>, 2012.
- M. D. Gould, M. A. Porter, S. Williams, M. McDonald, D. J. Fenn, and S. D. Howison. Limit order books. *Quantitative Finance*, 13(11):1709–1742, 2013.
- M. D. Gould, M. A. Porter, and S. D. Howison. Quasi-centralized limit order books. *arXiv:1310.8387*, 2016.

- J. Gregory. *Counterparty Credit Risk: The New Challenge for Global Financial Markets*. Wiley, West Sussex, UK, 2010.
- R. A. Jarrow and F. Yu. Counterparty risk and the pricing of defaultable securities. *The Journal of Finance*, 56(5):1765–1799, 2001.
- P. Jorion and G. Zhang. Credit contagion from counterparty risk. *The Journal of Finance*, 64(5):2053–2087, 2009.
- J.P. Morgan. Response to enhanced prudential standards and early remediation requirements for covered companies. Technical report, J.P. Morgan Chase and Co., available at http://www.federalreserve.gov/SECRS/2012/May/20120518/R-1438/R-1438_043012_107256_494026277948_1.pdf, 2012.
- Knight Capital Group. Retrieved 14 April 2015 from http://www.hotspotfx.com/download/userguide/HSFX/HSFX_UserGuide_wrapper.html, 2015.
- T. V. Koepl. The limits of central counterparty clearing: Collusive moral hazard and market liquidity. *Queen’s Economics Department Working Paper No. 1312*, available at <https://www.econstor.eu/dspace/bitstream/10419/97463/1/756789613.pdf>, 2013.
- C. Lee and M. J. Ready. Inferring trade direction from intraday data. *The Journal of Finance*, 46(2):733–746, 1991.
- L. Y. Liu, A. J. Patton, and K. Sheppard. Does anything beat 5-minute RV? a comparison of realized measures across multiple asset classes. *Journal of Econometrics*, 187(1):293–311, 2015.
- R. M. May, S. A. Levin, and G. Sugihara. Complex systems: Ecology for bankers. *Nature*, 451(7181):893–895, 2008.
- A. J. Menkveld. Crowded trades: an overlooked systemic risk for central clearing counterparties. *Working paper, SSRN eLibrary ID 2422250*, 2015.
- P. Norman. *The Risk Controllers: Central Counterparty Clearing in Globalised Financial Markets*. Wiley, West Sussex, UK, 2011.
- W. Perraudin and P. Vitale. Interdealer trade and information flows in a decentralized foreign exchange market. In J.A. Frankel, G. Galli, and A. Giovannini, editors, *The Microstructure of Foreign Exchange Markets*, pages 73–106. The University of Chicago Press, Chicago, IL, USA, 1996.
- C. Pirrong. Clearing and collateral aandates: A new liquidity trap? *Journal of Applied Corporate Finance*, 24(1):67–73, 2012.
- A. Rehlon and D. Nixon. Central counterparties: what are they, why do they matter and how does the bank supervise them? *The Bank of England Quarterly Bulletin*, 53(2):147–156, 2013.
- R. Roll. A simple implicit measure of the effective bidask spread in an efficient market. *The Journal of finance*, 39(4):1127–1139, 1984.
- K. A. Scott. Counterparty credit limits — U.S. regulations, Basel standards. *New York Law Journal*, 251(2), 2014.

The Basel Committee on Banking Supervision. Supervisory guidance for managing risks associated with the settlement of foreign exchange transactions. Technical report, Bank for International Settlements, available at <http://www.bis.org/publ/bcbs241.pdf>, 2013.

The Counterparty Risk Management Policy Group. Toward greater financial stability: A private sector perspective. Technical report, The Counterparty Risk Management Policy Group, available at <http://www.crmppolicygroup.org/crmppg2/docs/CRMPG-II.pdf>, 2005.

Appendix: Estimating Realized Volatility

In this appendix, we describe our methodology for estimating the realized variance of the quote-price and trade-price series. For a detailed discussion of this methodology and its empirical performance, see Liu et al. [2015].

For concreteness, we describe our methodology for buyer-initiated trades; the corresponding definitions for seller-initiated trades are similar. For a given currency pair on a given trading day, let D denote the total number of buyer-initiated trades that occur, let A_1, A_2, \dots, A_D denote the prices of these trades, and let a_1, a_2, \dots, a_D denote the ask-prices immediately before the arrival of these trades. For a given number K of intervals, and a given number L of sub-samples, let

$$T := D/K$$

denote the sample width and let

$$\tau := T/L$$

denote the sub-sample width. For a given value of j , we calculate the *sell-side trade returns*

$$r_i^A(j) = \log(A_{\lfloor (i+1)T+j\tau \rfloor}) - \log(A_{\lfloor iT+j\tau \rfloor}), \quad i \in \{1, \dots, K-1\}, \quad (13)$$

where $\lfloor x \rfloor$ denotes the largest integer less than or equal to x . We then calculate

$$v_A(j) := \sum_{i=1}^{K-1} (r_i^A(j))^2.$$

We repeat this process for each $j = 0, 1, \dots, L-1$, and finally calculate the *sell-side trade-price quadratic variation*

$$v_A = \frac{1}{L} \sum_{j=0}^{L-1} v_A(j). \quad (14)$$

We calculate the *sell-side quote-price quadratic variation* v_b similarly, using the a_1, a_2, \dots, a_D series.

To identify a suitable range of intervals K to consider, we created volatility signature plots (see Andersen et al. [2000]). We found that these plots were stable for values of K less than about 540 (which corresponds to $T = 1$ minute) for EUR/USD and GBP/USD and for values of K less than about 200 (which corresponds to $T \approx 3$ minutes) for EUR/GBP. For our results in the main text, we use the value $K = 108$, which is within the stable range for all three currency pairs.

Recent Issues

All CFS Working Papers are available at www.ifk-cfs.de.

No.	Authors	Title
580	Nikolaus Hautsch, Michael Noé, and S. Sarah Zhang	<i>The Ambivalent Role of High-Frequency Trading in Turbulent Market Periods</i>
579	Klaus Adam and Henning Weber	<i>Optimal Trend Inflation</i>
578	Günter Coenen, Michael Ehrmann, Gaetano Gaballo, Peter Hoffmann, Anton Nakov, Stefano Nardelli, Eric Persson, and Georg Strasser	<i>Communication of monetary policy in unconventional times</i>
577	Francis X. Diebold and Frank Schorfheide	<i>Real-Time Forecast Evaluation of DSGE Models with Stochastic Volatility</i>
576	Mert Demirer, Francis X. Diebold, Laura Liu, and Kamil Yilmaz	<i>Estimating global bank network connectedness</i>
575	Francis X. Diebold, Laura Liu, and Kamil Yilmaz	<i>Commodity Connectedness</i>
574	Annamaria Lusardi, Olivia S. Mitchell, and Noemi Oggero	<i>Debt and Financial Vulnerability on the Verge of Retirement</i>
573	Donald B. Keim and Olivia S. Mitchell	<i>Simplifying Choices in Defined Contribution Retirement Plan Design: A Case Study</i>
572	Joachim Grammig and Eva-Maria Küchlin	<i>A two-step indirect inference approach to estimate the long-run risk asset pricing model</i>