Counterparty Credit Limits: The Impact of a Risk-Mitigation Measure on Everyday Trading

Martin D. Gould* †[1,2], Nikolaus Hautsch[3,6,7,8], Sam D. Howison[1], and Mason A. Porter[1,4,5]

[1]Mathematical Institute, University of Oxford, Oxford OX2 6GG, UK
[2]CFM–Imperial Institute of Quantitative Finance, Department of Mathematics, Imperial College, London SW7 2AZ, UK
[3]Department of Statistics and Operations Research, University of Vienna, Vienna A-1090, Austria
[4]CABDyN Complexity CCentre, University of Oxford, Oxford OX1 1HP, UK
[5]Department of Mathematics, University of California, Los Angeles, CA 90095, USA
[6]Research Platform "Data Science @ Uni Vienna"
[7]Vienna Graduate School of Finance
[8]Center for Financial Studies, Frankfurt

Abstract

A counterparty credit limit (CCL) is a limit imposed by a financial institution to cap its exposure to a specified counterparty. CCLs help institutions mitigate counterparty risk by enabling selective diversification of their exposures, and they can thereby reduce the probability that an institution defaults after suffering one or more counterparty failures. However, CCLs do not only apply during times of market stress. This raises the question of how CCLs impact the prices that institutions pay for their trades during everyday trading. In this paper, we examine this question both empirically and via a new model of trading with CCLs. We study a high-quality data set for three liquid currency pairs in the foreign-exchange spot market during May and June 2010, and we find that CCLs caused little impact on trade prices during this period. However, our model highlights that, in some situations, CCLs can have a major impact.

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

1 Introduction

The international financial crisis of 2008 underlines the vital importance of understanding counterparty risk. 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 demonstrated how 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

*Corresponding author. Email: gouldm@maths.ox.ac.uk.
†MDG completed part of this work while an academic visitor at the Humboldt-Universität zu Berlin.

and implementing measures to mitigate the risk of default contagions is an extremely important task.

One mitigation measure, which is currently implemented by several multi-institution trading platforms in the foreign-exchange (FX) spot market, is the use of *counterparty credit limits* (CCLs). A CCL is a limit imposed by a financial institution to cap its exposure to a specified counterparty. CCLs are designed to complement existing risk-based capital requirements by protecting financial institutions from large losses resulting from sudden counterparty defaults.

Accompanying this benefit, the application of CCLs also entails an important drawback. With CCLs in place, institutions can access only the subset of trading opportunities that are offered by counterparties with whom they possess sufficient bilateral credit. Therefore, CCLs restrict the set of trading opportunities that institutions can access. In extreme cases, this restriction may be so severe that individual institutions experience a liquidity crisis and fail. In this way, CCLs may themselves influence systemic risk.

Importantly, however, this restriction applies not only during periods of market-wide stress, but also at all other times. One aim of the present paper is to assess the effect of CCLs on the prices that institutions pay for their trades during everyday trading. To study this question, we use an unusually rich data set that describes all trading activity for three liquid currency pairs on Hotspot FX during all of May and June 2010. Hotspot FX is a large electronic trading platform in the FX spot market that enables institutions to apply CCLs. Crucially, the Hotspot FX data enables us to measure how CCLs impact the prices that individual institutions pay for their trades. Because the period from May to June in 2010 was relatively calm, we are able to study how much market participants pay, during 'normal' trading, as a consequence of CCLs. To our knowledge, ours is the first study to investigate this topic.

We introduce the notion of the "skipping cost" of a trade to measure the additional cost that an institution bears from the application of CCLs. In our data set, more than half of the trades have a skipping cost of 0, and the mean skipping cost is less than half a basis point. We do identify a handful of trades with large skipping costs, but 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 on Hotspot FX. These empirical results suggest that CCLs had very little impact on trade prices for the trades that we study.

As our emprical CCL network is fixed (and not known to us), the study of historical data provides no insight into how our results might change if institutions made substantial modifications to their CCLs. We therefore complement our empirical analysis by investigating a second research question: How does network *structure* affect the prices that institutions pay for their trades in the presence of CCLs? We use simulations of an agent-based 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, our model allows us to investigate how varying the density and topology of the connections in a market's CCL network can affect the prices of trades.

Our simulations provide valuable insights into the roles of both density and topology. For example, when the CCL network is dense (in the sense that most institutions can access most trading opportunities), we find that CCLs have very little impact on the prices of trades. However, as the edge density falls, we observe that the skipping costs of trades rise sharply. The network topology also has a noticeable effect: For a given edge density, skipping costs are markedly higher for a network with core–periphery structure than they are for an Erdős–Rényi network (in which one places edges uniformly at random).

We also investigate the impact of CCL network properties on realized trade-price and quote-price volatility. As the edge density falls, the trade-price volatility rises sharply, whereas the quote-price volatility remains almost constant. The volatilities are also both considerably higher for a core–periphery network than for an Erdős–Rényi network.

Together, our empirical and simulation results paint an interesting and complex picture about the impact of CCLs, raising important questions for policy makers and regulators. Both sets of results illustrate that the application of CCLs does not necessarily force many institutions to pay large skipping costs for their trades. Therefore, CCLs may provide market participants with the benefits of selective diversification without causing them to incur large additional costs during the course of their everyday trading. However, our model 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 thus argue that understanding and monitoring how institutions set and adjust their CCLs is a vital step for regulators in assessing how their implementation may impact market stability and, ultimately, whether they may lead to localized liquidity crises that cause individual institutions to fail.

The paper proceeds as follows. In Section 2, we give an introduction to counterparty risk, discuss the strengths and weaknesses of two well-known mechanisms designed to mitigate it, and review a selection of relevant literature. In Section 3, we describe the CCL mechanism in detail, and discuss how CCLs are currently implemented by several large electronic trading platforms in the FX spot market. 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. In the appendices, we give a detailed description of the data that forms the basis of our empirical study, and we describe our methodology for estimating realized volatility.

2 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 lead to a cascade of rapidly propagating institutional failures. Counterparty risk is therefore 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 that arises from derivative contracts (see, e.g., Brigo et al. [2013] for a detailed survey). However, 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.

2.1 Approaches to Mitigating Counterparty Risk

Among several possible approaches to mitigating counterparty risk, two have received particular attention. The first is to novate trade 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 that arise from all contracts agreed between two counterparties. If one counterparty fails, the other is protected via the default-management procedures and the resources of the CCP. During the past decade, several prominent regulatory bodies (see, e.g., The Basel Committee on Banking Supervision [2013] and The Counterparty Risk Management Policy Group [2005]) have argued that CCPs are an effective tool for mitigating counterparty risk. The second approach is to apply 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 that is generated by a CVA to construct or purchase a contingent claim whose payoff is triggered by the default of the given counterparty, such that the resulting net loss is 0.

Despite their clear benefits, the above two approaches to mitigating counterparty risk also suffer from important drawbacks. CCPs require that institutions reserve capital for margin calls and contribute to a default fund. This reduces the amount of capital that institutions have available to conduct trades during the course of everyday trading. Moreover, Pirrong [2012] argued that CCP novation does not reduce the aggregate counterparty risk across all institutions; instead, it concentrates all such risk into the CCP, which thus becomes a single point of failure of 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. Finally, Koeppl [2013] noted that CCPs generate moral hazard by removing the incentive for individual institutions to assess the creditworthiness of their trading counterparties. 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 brings 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, this estimation provides no insight into how to construct a portfolio with the required payoff upon a counterparty default. Indeed, constructing this portfolio is often impossible in practice. Moreover, CVAs are not suitable for assets that are traded on an exchange in which many different institutions access the same centralized set of trading opportunities (as in a limit order book (LOB); see Gould et al. [2013] for a detailed introduction to LOBs), as 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 provides a panacea for counterparty risk. Their failure to provide a conclusive solution is strong motivation for exploring alternative avenues.

2.2 Related Literature on the Role of Connectivity in Risk Analysis

Our work contributes to the rapidly growing literature on possible mechanisms for mitigating counterparty risk. The vast majority of that literature deals with the performance of risk-mitigation measures in the extreme circumstances for which they are designed, and it thereby encompasses the analysis of systemic risk [De Bandt and Hartmann, 2000]. Our work is complementary, in that we examine the effects of risk-mitigation measures on everyday price formation and trading, with a focus on the use of CCLs. Notwithstanding this different emphasis, it is helpful to set the scene for our own work in relation to the existing literature.

As argued by Jarrow and Yu [2001], financial institutions face significant counterparty risks whenever their exposures are concentrated in a small number of counterparties, because the default of any counterparty is likely to cause severe financial distress. Therefore, many financial institutions seek to diversify their counterparty risk exposures by trading with many different counterparties. Moreover, establishing trading relationships with a wide range of other financial institutions reduces the likelihood of experiencing a subsequent liquidity shortage. Several authors have argued that such liquidity shortages can impact systemic risk through *illiquidity contagions* [Anand et al., 2012, Bardoscia et al., 2017, Gai et al., 2011].

Despite these clear benefits, diversification of counterparty risk also has important drawbacks. As noted by Stiglitz [2010] and Roukny et al. [2013], when it comes to counterparty risk, diversification and contagion are two sides of the same coin. All else being equal, the larger the number of different counterparty credit exposures of a given financial institution, the more likely it is to experience a counterparty default. Many authors have studied how the failure of a single institution can propagate through a financial network as a *default contagion* (see, e.g., [Giesecke and Weber, 2004, Jorion and Zhang, 2009, May et al., 2008]).

In an early paper on the topic, Stiglitz [2010] introduced a model in which the failure of a financial institution causes all of its counterparties to fail. In this setting, connectedness among financial institutions leads to default contagion. More recently, other authors have studied more complicated models to examine trade-offs between default contagion (whose drawback increases with the density of counterparty credit exposures) and diversification (whose benefit increases correspondingly) [Bardoscia et al., 2017, Battiston et al., 2012a,b, Tasca et al., 2017].

Battiston et al. [2012a] assumed that when an institution defaults, all of its counterparties suffer a loss that equals their exposure to the defaulting institution. They then reasoned that when a financial institution has a larger number of connections, one should expect it to feel a smaller idiosyncratic shock if one of those counterparties fails, and they concluded that diversification of exposures across many financial institutions has complicated effects on systemic risk. Battiston et al. [2012b] studied a system of coupled stochastic processes to simultaneously examine the impacts of a default contagion and the benefits of diversification. They reported that their model leads to a non-monotonic ("U-shaped") relationship between the amount of connectivity in a financial network and the corresponding probability of a large default cascade. In particular, they stressed that when agents are embedded in a densely-connected network of mutual liabilities, a financial system becomes very unstable and is prone to frequent and large crises. Similarly, Bardoscia et al. [2017] illustrated that processes such as market integration and diversification, which are widely believed to stabilize financial systems, can actually undermine systemic stability by creating cyclical credit-exposure structures that amplify the effects of shocks. Tasca et al. [2017] studied a model in which systemic default is related not only to financial institutions' connections to each other, but also to their exposures to external assets. They concluded that the probability of systemic failure is a non-monotonic function of exposure to external assets, with an interior optimum corresponding to moderate (rather than complete) diversification.

Other authors have focused on how the topology of counterparty exposures impacts default cascades. Roukny et al. [2013] investigated how the size of default cascades varies between Erdős–Rényi networks and networks with a heavy-tailed degree distribution.[1] When considering only the benefits of diversification, the authors concluded that network topology does not

[1] We follow a similar approach in Section 5 to compare skipping costs in Erdős–Rényi networks and core–periphery networks.

impact default cascades. However, when also considering the impact of a default contagion, they reported that network topology strongly impacts their results. In their study, they also illustrated that no single market topology is always better than all others. Luu et al. [2018] examined how network topology affects the dynamics of collateral (and the consequent systemic risk) in the presence of rehypothecation. They observed rather different dynamics for different network topologies. Both Roukny et al. [2013] and Luu et al. [2018] argued that it is important for regulators to be aware of network topology when making policy decisions.

Despite the considerable size of this literature, almost all published work to date on the role of connectivity in risk analysis has focused on the question of how connections between different financial institutions impact default contagions [Gai and Kapadia, 2010]. To our knowledge, no publications have examined the related question of how such connections impact the prices that institutions pay for their trades during the course of everyday trading. This is the primary question that we investigate in the present paper, in the context of CCLs.

3 Counterparty Credit Limits

A key approach — and an alternative to those that we discussed in Section 2.1 — to mitigating counterparty risk is to apply counterparty credit limits (CCLs). Consider a financial market that is populated by a set of institutions, $\Theta = \{\theta_1, \theta_2, \ldots\}$, 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. Such counterparty credit exposures occur in all financial markets in which the agreement and settlement of trades does not occur simultaneously. In the FX spot market, for example, trades agreed on day D are settled on day $D + 2$, so each trade entails exposure to the counterparty during the period between day D and day $D + 2$.

Assigning a CCL to a given counterparty does not require posting collateral; instead, it involves notifying the exchange of the relevant value $c_{(i,j)}$. Institution θ_i cannot enter any trade 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 maximum amount that θ_i and θ_j can trade is therefore equal to $\min\left(c_{(i,j)}, c_{(j,i)}\right)$. We call this quantity the *bilateral CCL between θ_i and θ_j*. These bilateral CCLs determine the subset of trading opportunities that is available to each institution. This subset changes over time according to the relevant institutions' trading activities.

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 $c_{(i,j)} = 0$. Alternatively, if θ_i perceives θ_j to be extremely unlikely to default, then θ_i can assign an unlimited amount of credit to θ_j by setting $c_{(i,j)} = \infty$.

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.[2] In contrast to CVAs, the use of CCLs does not require institutions to estimate the market value of their counterparty risk. Instead, it enables them to specify an upper bound on each of their counterparty exposures. Institutions can thereby use CCLs to mitigate counterparty risk by selective diversification of their exposures.

Several major multi-institution electronic trading platforms in the FX spot market 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. Trades occur via a mechanism similar to a standard limit order book (LOB), except that institutions can only conduct

[2] 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.

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 $c_{(j,i)}$ is exceeded 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 (otherwise centralized) LOB. 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 (i.e., the set of all active orders owned by all market participants). Instead, each institution sees only the active orders that correspond to its own trading opportunities (i.e., that do not violate any of its bilateral CCLs) at time t.[3] 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 subset, filtered according to its CCLs, of all active limit orders.

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

4 Empirical Results

Our empirical investigation uses 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 pairs[4] 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. We give a detailed description of the Hotspot FX data in Appendix A.

4.1 Skipping Costs

We first examine the impact of CCLs on the prices of individual trades. As we discussed in Section 3, 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 occurs 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 that is borne by the institution that submits a market order as a result of CCLs preventing this institution from accessing a better-priced trading opportunity. We call this additional cost the “skipping cost”.

[3]Some QCLOB platforms, such as Reuters and Electronic Broking Services (EBS), offer an additional data feed that, in exchange for a fee, provides snapshots of the global LOB at regular time intervals.

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

For a given currency pair on a given trading day, let p_k denote the price of the k^{th} trade, and let b_k and a_k denote, respectively, the bid-price and ask-price in the global LOB immediately before this trade occurs. The *skipping cost* of a trade is

$$r_k = \begin{cases} p_k - a_k \,, & \text{if the } k^{\text{th}} \text{ trade is a buyer-initiated trade}\,, \\ b_k - p_k \,, & \text{if the } k^{\text{th}} \text{ trade is a seller-initiated trade}\,. \end{cases} \tag{1}$$

In Equation (1), the sign difference between buyer-initiated and seller-initiated trades reflects the fact that every trade has a non-negative skipping cost. In the extreme case with CCLs in which 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.

Because the prices of trades vary across currency pairs and across time, we also normalize each skipping cost by the mid-price $m_k = \frac{1}{2}(a_k + b_k)$ immediately before a trade occurs. Specifically, we calculate the *normalized skipping cost*

$$\tilde{r}_k = \frac{r_k}{m_k}\,, \tag{2}$$

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

In Figure 1, we show the empirical cumulative density functions (ECDFs) of normalized skipping costs $\tilde{r}_k$. More than half of all trades have a (normalized) skipping cost of 0, which implies that they occurred at the best price available in the global LOB at their time of execution. As illustrated by the log-survivor functions,[5] the distribution of normalized skipping costs has a similar shape for all three currency pairs. This similarity extends to about the 99.9$^{\text{th}}$ percentile of the distributions. Beyond this, the currency pair EUR/USD includes a handful of trades with extremely large skipping costs; this does not occur for the other two currency pairs.

In Table 1, we give summary statistics about the 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%). As with Figure 1, these results suggest that the statistical properties of normalized skipping costs are similar for each of the three currency pairs.

Table 1: 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

[5] The survivor function is given by 1 minus the ECDF, and the log-survivor function is its logarithm.

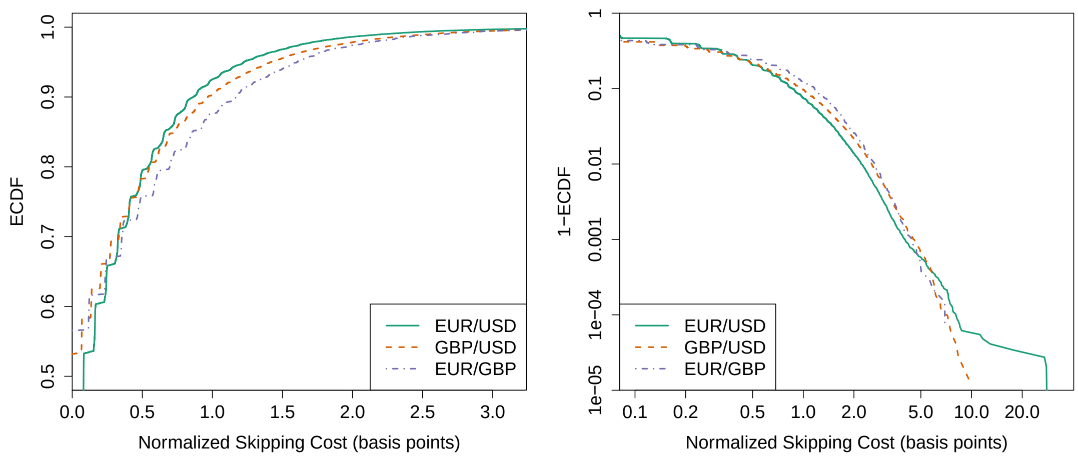

Figure 1: (Left) Empirical cumulative density function (ECDF) and (right) log-survivor function for the normalized skipping costs $\tilde{r}_k$ of (solid green curves) EUR/USD, (dashed orange curves) GBP/USD, and (dash-dotted purple curves) EUR/GBP trades on Hotspot FX during May–June 2010.

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 Appendix A), these skipping costs correspond to a mean additional cost of about £18.00 and \$30.00, respectively, for an institution submitting 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, which corresponds to incurring a total additional cost of about \$3630.00 when submitting a trade of size 1 million euros.

On Hotspot FX, institutions can infer the approximate skipping cost of their trades by comparing their local bid-price 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, as we discussed in Section 3, they observe via their trade-data stream). Given that this is the case, why do some institutions perform trades that have extremely large skipping costs during the course of everyday trading?

We believe 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. 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 everyday trading. In a recent discussion of modern financial markets, Luu et al. [2018] argued that the advent of trading platforms with relatively low barriers to entry (such as Hotspot FX) have blurred the lines between the inter-bank market and less-traditional markets. The significant heterogeneity that we observe in skipping costs is consistent with the idea that a wide and heterogeneous population of financial institutions indeed operate on such platforms, albeit with access to considerably different prices for their trades.

4.2 Price Changes

We now examine price changes between successive trades. Recall that p_k denotes the price of the k^{th} trade for eachcurrency pair on a given trading day. For each k, let $p_{k'}$ denote the price of the *previous* trade in the same direction as the k^{th} trade. Similarly, let $b_{k'}$, $a_{k'}$, and $m_{k'}$ denote, respectively, the bid-, ask-, and mid-prices immediately before the previous trade in the same direction as the k^{th} trade. The *change in trade price* is

$$f_k = \begin{cases} p_k - p_{k'}\,, & \text{if the } k^{\text{th}} \text{ trade is a buyer-initiated trade}\,, \\ p_{k'} - p_k\,, & \text{if the } k^{\text{th}} \text{ trade is a seller-initiated trade}\,. \end{cases} \tag{3}$$

Similar to Equation (2), we also calculate the *normalized change in trade price*,

$$\tilde{f}_k = \frac{f_k}{m_k}\,, \tag{4}$$

which is independent of the size of the underlying exchange rate; it thereby allows easier comparisons across different currency pairs.

Our results in Section 4.1 reveal that skipping costs 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 trading opportunities that are available on the platform. This observation raises the question of how strongly CCLs impact the price changes between successive trades. 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 4.1 suggest is the case), then the trade-price series may include large fluctuations that do not reflect similar changes in an asset's fundamental value. Therefore, the price-formation process on a platform that implements CCLs may be rather different from that on a platform in which all institutions can trade with all others.

To study this question empirically, we introduce the following decomposition of each term in the f_k series into two constituent parts. We define the *change in quote price* between the k^{th} trade and the previous trade in the same direction by

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

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

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

The three price-change series are not independent. For buyer-initiated trades, observe that

$$\begin{aligned} f_k &= p_k - p_{k'} \\ &= (a_k - a_{k'}) + ((p_k - a_k) - (p_{k'} - a_{k'})) \\ &= g_k + h_k\,. \end{aligned} \tag{7}$$

By a similar argument, the same result holds for seller-initiated trades.

Equation (7) enables us to decompose each change in trade price into the corresponding constituent change in quote price and change in skipping cost: The price change between any pair of trades in the same direction is the sum of the price change of the best quotes and the change in skipping costs of the two trades.

We now report 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 changes in trade price. In the left panel of Figure 2, we show a quantile–quantile (QQ) plot of the f_k series versus the g_k series. In this plot, the quantile points cluster tightly along the diagonal, implying 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 2, we show a QQ plot of the f_k series versus the h_k series. In this case, the distribution of changes in skipping costs is concentrated more tightly around 0 than is the distribution of changes in trade prices, suggesting that the changes in skipping cost account for only a small fraction of the total price change between successive trades.

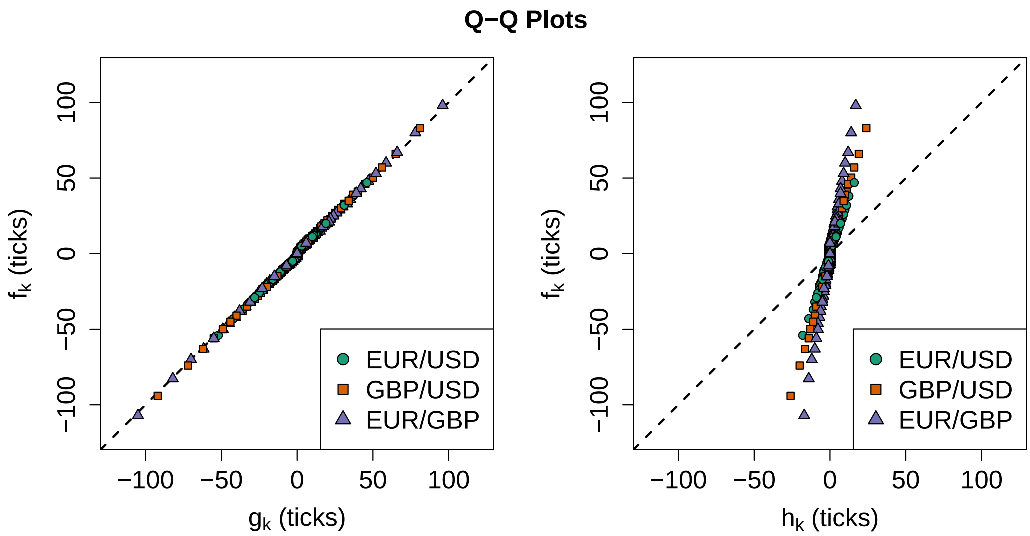

Figure 2: Quantile–quantile (QQ) plots for (left) changes in trade price f_k versus changes in quote price g_k and (right) changes in trade price f_k versus changes in skipping costs h_k for EUR/USD (green circles), GBP/USD (orange squares), and EUR/GBP (purple triangles) trades on Hotspot FX during May–June 2010. In each plot, the points indicate the $0.01, 0.02, \ldots, 0.99$ quantiles of the empirical distributions. The dashed black lines indicate the diagonal.

To assess whether the similarities and differences that we highlight in Figure 2 also hold at the trade-by-trade level (and not only at the level of the unconditional distributions), we make scatter plots of the individual terms of the series. In the left column of Figure 3, we show scatter plots of the f_k series versus 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 data points cluster along the diagonal. In the right column of Figure 3, we show scatter plots of the f_k series versus 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.

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 Pearson correlation ρ between these series (see Table 4.2). For each currency pair, the sample Pearson 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 3) and suggests that changes in trade price are strongly correlated with corresponding changes in the underlying quotes. By contrast, the sample Pearson correlations between the f_k and h_k series are very

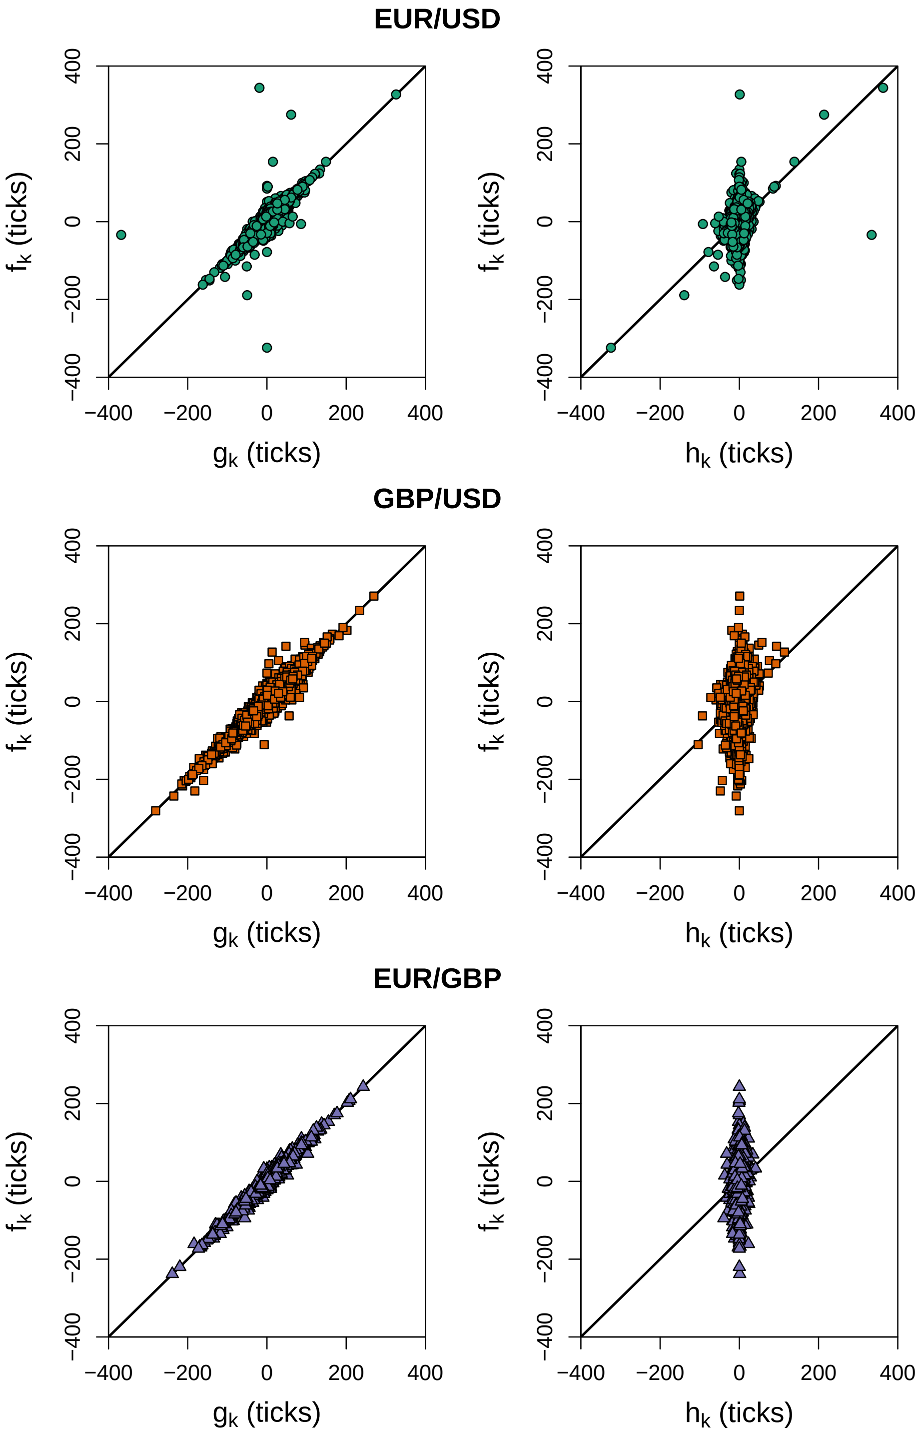

Figure 3: Scatter plots of changes in trade price f_k versus (left column) changes in quote price g_k and (right column) changes in skipping costs h_k for (top row) EUR/USD, (middle row) GBP/USD, and (bottom row) EUR/GBP trades on Hotspot FX during May–June 2010. The solid black lines indicate the diagonal.

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.

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

Table 2: Sample Pearson 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–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.

Taken together, our results in this section suggest the following interpretation of Equation (7). 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 it 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 (additive) 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, during the course of everyday trading, fundamental revaluations in trade prices arise due to corresponding changes in the best quotes. One can regard the f_k series as a noisy observation of the g_k series, where the uncorrelated, additive noise arises from the restriction of institutions' trading activities to their bilateral trading partners. The strength of this effect varies across institutions because of the heterogeneity in their CCLs.

We conclude that CCLs have a small effect on trade-to-trade price changes, and thus on volatility, for the trades that we study. On the one hand, this may arise from the fact that our data describes a relatively calm period for financial markets. On the other hand, it may be because of the underlying network structure of CCLs, which is unobservable to us. Therefore, studying historical data is only one aspect of understanding how CCLs can affect financial markets, because such analysis does not provide insight into how our results may change if institutions make substantial modifications to their CCLs. It is thus important to examine CCL network topology and its effects on a model of trading.

5 A Model of Trade with CCLs

To explore the role of the CCL network's topology, we simulate a trading model with CCLs. 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 that is offering to buy or sell at a mutually agreeable price. A crucial feature is that not all institutions can trade with all others; instead, each institution can trade only subject to its CCLs. These trades occur at prices that depend not only on other institutions' buy and sell prices, but also on the underlying network of bilateral CCLs between them.

As in the model of Luu et al. [2018], we examine a simple trading mechanism — ignoring many features of real markets — to highlight the relationship between the underlying network topology and the consequent trades. Our motivation is to provide a direct link between the institutions' interaction topology (via the network of CCLs) and the corresponding trade prices. More precisely, for two different types of network topologies, we perform numerical simulations to study how the density of edges in the CCL network impacts the trade-price series. Our approach is similar to Roukny et al. [2013], who examined how the size of default cascades varies as a function of the mean number of edges per node, for both Erdős–Rényi networks and networks with a heavy-tailed degree distribution (see Section 2.2).

In order to focus on how restricting the trading opportunities available to institutions affects both the prices of individual trades and the realized volatility of the quote-price and trade-price series, we fix all parameters except those related to the CCL network edge density and topology. Modelling approaches that ignore the heterogeneous impact of CCLs are not appropriate for such an investigation, while ours provides a natural framework for studying this problem. Motivated by our empirical results in Section 4, we study properties (including volatilities) of the trade-price, quote-price, and skipping cost series.

The setting for our model is an infinite-horizon, continuous-time market populated by a set of N institutions, $\Theta = \{\theta_1, \theta_2, \ldots, \theta_N\}$, 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 their view of the likely future value of the asset, as well as differences in their inventory, cash-flow, financing constraints, and so on.[6] 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 using 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 in Section 5.1; for now, we remark only that the values of s_t^i are constrained to never fall below the 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 give 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

$$\mathrm{d}s_t^i = -\kappa \left(s_t^i - s_0\right) \mathrm{d}t, \tag{8}$$

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

$$\mathrm{d}M_t^i = \gamma M_t^i \, \mathrm{d}W_t^{M^i}, \tag{9}$$

[6] We use capital letters to denote institutions' buy and sell valuations to avoid confusion with the (global) bid- and ask-prices, as defined in Section 4.1.

where $\gamma > 0$ is the mid-price volatility (with units of $(\text{time})^{-\frac{1}{2}}$) and $W_t^{M_i}$ are mutually independent Brownian motions.[7]

Equation (8) causes each institution's spread to approach its minimum value s_0 as time progresses after a trade. In the absence of trading, the processes M_t^i are drift-free geometric Brownian motions. Our model minimizes the complications associated with the mixing of price and time scales in the model parameters; the temporal evolution is influenced directly only by the parameter γ (whose inverse defines a time scale) and by trading via the network of CCLs. Although a geometric Brownian motion with no drift has constant mean, its variance increases with time. Without trading, our mid-prices thus spread out progressively and indefinitely over time. As we will 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 discuss this feature in more detail in Section 5.2.

5.2 The CCL Network

We assume that each institution θ_i assigns a CCL to each other institution θ_j. To perform a time-stationary investigation, we assume that each institution's access to trading opportunities does not depend on time, and, specifically, that it does not vary according to the trading history. Therefore, 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.

In a real financial market, a pair of financial institutions can access each other's trading opportunities until they reach their bilateral CCL (see Section 3). Therefore, our modeling framework is a simplification of the way that real CCLs operate in real markets. However, there are three important benefits to this approach in the present context. First, institutions in the FX spot market routinely trade huge volumes of FX each day, yet the modal size of market orders for each of the three currency pairs on Hotspot FX is just 1 million units of the base currency. If a given pair of institutions have a sufficient bilateral CCL to access each other's trading opportunities once, then they are likely to be able to do so again. Second, in the FX spot market, trades agreed on day D are settled on day $D + 2$. Given this relatively short time interval, if an institution θ_i cannot access a trading opportunity that is offered by θ_j, we claim that it is more likely that this is because their bilateral CCL is actually 0, rather than because they have gradually accumulated a very large exposure. Third, this framework makes our model time-stationary. By contrast, tracking the cumulative exposure between a pair of institutions and allowing them to trade only up to their bilateral CCL would yield a model that is non-stationary in time. We opt for this significant gain in model simplicity in what we regard to be an adequately realistic equilibrium-pricing framework, instead of the more complex route of tracking each pair of institutions' cumulative exposures.

We encode CCLs with an undirected, unweighted 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 require this network, which we call a *CCL network*, to be connected; otherwise, the disconnected components can

[7] A possible refinement of the model is to include a common market factor w_T^m in addition to the idiosyncratic noise terms. In that case, $\mathrm{d}M_t^i = \gamma M_t^i \left(\rho_i \, \mathrm{d}W_t^m + \sqrt{1-\rho_i^2}\, \mathrm{d}W_t^{M^i}\right)$, where ρ_i is a (common) Pearson correlation coefficient. We also performed simulations of this more complex model, but we found that this additional complication adds little to our analysis. We therefore restrict our discussion to the simpler model, without a common market factor.

drift apart, as there is no means to link them together. We show some example CCL networks in Figure 4.

Any network with N nodes (and no self-edges or multi-edges) has at most $N(N-1)/2$ edges. A CCL network with N nodes and n edges has an *edge density*

$$d = \frac{2n}{N(N-1)}\,. \tag{10}$$

With our model, we can consider CCL networks with any number of nodes and any configuration of edges. However, throughout this section, 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 that we discuss are *Erdős–Rényi networks* $\mathcal{G}(N,n)$ [Erdős and Rényi, 1960, Newman, 2018], in which we place a specified number of edges 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 Erdős–Rényi networks enables us to investigate the temporal evolution of our model in a simple, stylized framework with no deterministic structure.

To construct Erdős–Rényi networks, we fix the edge density d and then use Equation (10) to calculate the required number of edges n. We place these n edges uniformly at random among the N nodes of the network, and we then check whether the network consists of a single connected component that spans all nodes. If so, we accept it; if not, we reject it and construct an alternative network using the same rules.[8]

For a given choice of d, we construct a sample of 1000 such CCL networks, and we then simulate 1000 independent runs of our model for each of these 1000 CCL networks, where each run uses a different seed for a Mersenne Twister pseudo-random number generator (which controls the temporal evolution of the model). In the left panel of Figure 4, we illustrate a single instantiation of an Erdős–Rényi 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], Rombach et al. [2017]), which have two types of nodes: core nodes and peripheral nodes. Each core node is adjacent to all other core nodes. Each peripheral node is adjacent to exactly one core node, so the degrees of no two core nodes differ by more than 1. In the right panel of Figure 4, we illustrate a core–periphery CCL network with $N = 12$ institutions, with 3 core nodes and 9 peripheral nodes.

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 have a credit line only 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, with a core that consists of large, international banks and a periphery that consists of smaller financial institutions, such as small banks, hedge funds, or mutual funds (see, e.g., Craig and von Peter [2014], Fricke and Lux [2012], Iori et al. [2008]). Therefore, our core–periphery structure represents an approximation of the complex structure of real markets, although we have simplified it to a convenient deterministic form.

To construct core–periphery networks, we first fix the fraction ψ of peripheral 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 construct a single CCL network, and we then simulate 1000 independent runs of the model for it. Each run uses a different seed for the pseudo-random number generator.

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

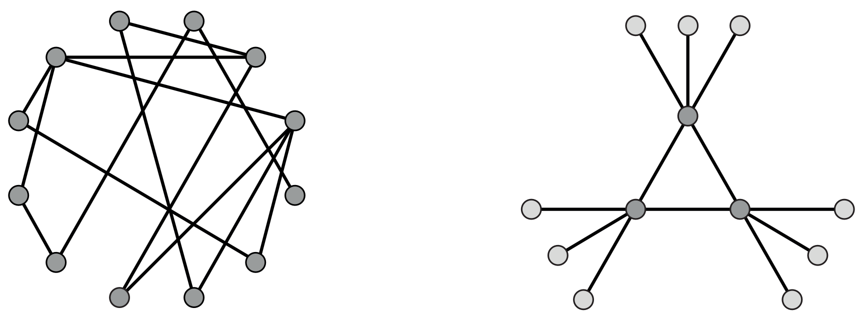

Figure 4: Schematics of CCL networks. (Left) An Erdős–Rényi network with $N = 12$ institutions and $n = 14$ edges and (right) a core–periphery network with $N = 12$ institutions, of which 3 nodes are core nodes (dark grey) and 9 nodes are peripheral nodes (light grey).

Our choice of these two classes of networks resembles those in Luu et al. [2018], who used similar network structures to investigate the dynamics of collateral, and the consequent systemic risk, in the presence of rehypothecation (see Section 2.2). They reported different dynamics when when the underlying network is an Erdős–Rényi network than when it has core–periphery structure. We give a more detailed comparison of our results to those of Luu et al. [2018] in Section 5.7.

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 $B^i_{t^*} = A^j_{t^*}$. We call this price the *trade price*.

In our empirical results in Section 4, recall that we classified 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 use the following simple rule to classify trades. Let

$$\bar{M}_t = \frac{1}{N}\sum_{i=1}^{N} M^i_t \tag{11}$$

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 $p = B^i_{t^*} = A^j_{t^*}$. 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 — which is then matched by this market order — at the trade price. The initiator trades at a relatively unfavorable price; the fewer bilateral CCLs that the initiator has, the further we expect this price 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) 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 $M^i_{t^*}$ and adding $s_0/2$ to $M^j_{t^*}$ (respectively, subtracting $s_0/2$ from $M^j_{t^*}$ and adding $s_0/2$ to $M^i_{t^*}$) and widening each of s^i_t and s^j_t by $s_0/2$. Finally, we reset the values of $B^i_{t^*}$, $A^i_{t^*}$, $B^j_{t^*}$, and $A^j_{t^*}$ according to these new mid-prices and spreads. This separates $B^i_{t^*}$ and $A^j_{t^*}$, which are equal, by $3s_0/2$. All other prices (including $B^j_{t^*}$ and $A^i_{t^*}$) 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 mean 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 then move closer together by s_0, reversing the previous separation. The spreads then revert towards s_0, which stops 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) difference scheme. In general, this choice 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 discrete time-stepping may be so large that it creates more than one trading opportunity. For example, the price moves that occur in a discrete time step may 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 we repeat this procedure until there are no further trading opportunities to consider.

5.5 Parameter Choices and Implementation

Our primary aim is to investigate how CCLs affect skipping costs. We therefore fix the values of γ, κ, and s_0; and we study how our model's output varies for different CCL networks (see Section 5.2).

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 set 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 that we 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 for which $\theta_i \leftrightarrow \theta_j$).

The final parameter in our model is the discretization time step. The dominant term in the discrete temporal 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 in which the discrete time steps regularly create multiple simultaneous trading opportunities. We expect the separation of the mid-prices to be $O(\epsilon M_0/N)$, and we would like this to be at least three times the standard deviation of the noise term. 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 simulation results that we present in Section 5.6, we use $N = 128$ institutions.[9] For each CCL network that we study, we simulate the temporal evolution of our model from $t = 0$ to $t = 10$. We 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 are sufficiently large by examining the numerical stability of our results using a variety of different burn-in periods and total time lengths. Specifically, our results are 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 temporal scale for the reversion of B_t^i and A_t^i) and for all total time periods that are larger than about $t = 2$.

In Figure 5, we show a single simulation of the model in which we use $N = 3$ institutions, the parameter values in Section 5.5, and a CCL network 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. As the figure illustrates, the prices of subsequent trades can deviate considerably from each other. Therefore, our model does a good job of capturing how heterogeneity in institutions' access to trading opportunities (which arise as a direct consequence of their CCLs) can manifest in the trade-price series.

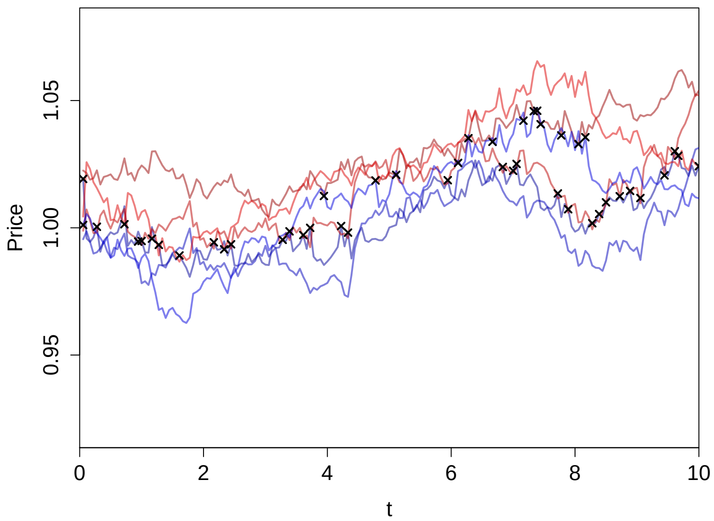

Figure 5: An example simulation of our model on a CCL network with $N = 3$ institutions in which $\theta_1 \leftrightarrow \theta_2$ and $\theta_1 \leftrightarrow \theta_3$, but in which θ_2 and θ_3 cannot trade with each other. The blue curves indicate the institutions' private buy valuations B_t^i, and the red curves indicate the institutions' private sell valuations A_t^i. The black crosses indicate trades.

5.6 Simulation Results

We study buyer-initiated and seller-initiated trades separately via the trade-classification algorithm that we 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 samples, we present our results for all trades together (i.e., we aggregate buyer-initiated and seller-initiated trades).

[9] We also repeated these 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.

In Figure 6, we plot the number of trades that occur for each edge density d. For both Erdős–Rényi networks and core–periphery networks, CCL networks with lower edge densities have fewer trades. The intuition is simple: The lower the edge density, the lower the number of bilateral trading partners in the population; this, in turn, leads to a smaller number of trades occurring within a given time horizon. Figure 6 also illustrates that the mean number of trades that occur for a given CCL network depends on the network topology (and 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 (on average) in an Erdős–Rényi 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 the number of trading partners of each institution has, but also on the identities of those trading partners.

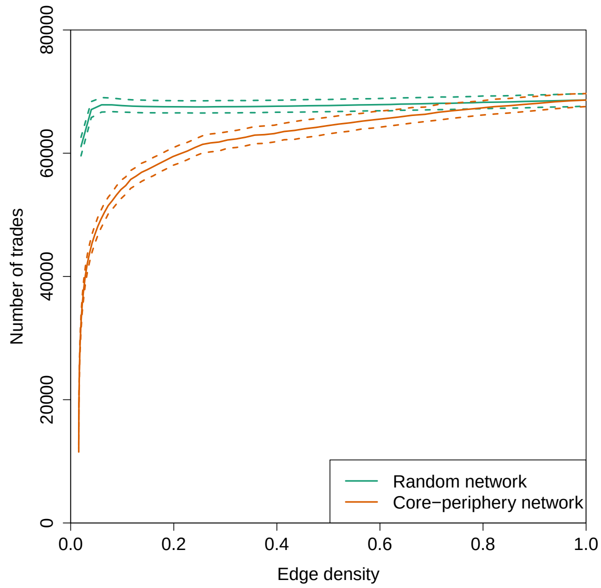

Figure 6: Mean number of trades for the (green) Erdős–Rényi networks and (orange) core–periphery networks. Solid curves indicate the mean across all independent runs of the model; dashed curves indicate one standard deviation away from these means.

In Figure 7, we plot the mean skipping costs of the trades that occur for each edge density d. For both Erdős–Rényi networks and core–periphery networks, CCL networks with lower edge densities have larger skipping costs. Intuitively, this result illustrates that the more restrictive CCLs are to institutions' access to trading opportunities, the larger the skipping costs that institutions pay for their trades.

Figure 7 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 of trades for an Erdős–Rényi network is skipping cost among trades for a core–periphery network.

For both classes of networks, the mean skipping cost decreases rapidly as the edge density increases from 0 to about 0.1. For Erdős–Rényi networks, the mean skipping cost is very close to 0 for all edge densities above about 0.3. In this case, CCLs have a very small impact on individual trade prices. For core–periphery networks, the mean skipping cost remains much larger before eventually decreasing to 0 as the edge density reaches 1 (at which 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 networks.

Because our model allows us to observe all trades that occur between the agents that trade

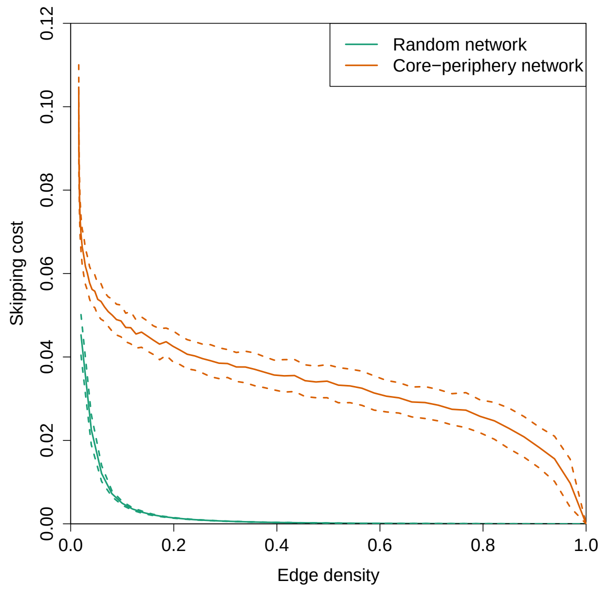

Figure 7: Mean skipping costs of trades for (green) Erdős–Rényi networks and (orange) core–periphery networks. Solid curves indicate the mean across all independent runs of the model; dashed curves indicate one standard deviation.

in it, we are also able to study the realized volatility of both the quote-price and trade-price series. We describe our methodology for measuring the realized volatility in Appendix B. In Figure 8, we plot the trade-price and quote-price volatility for each edge density d.

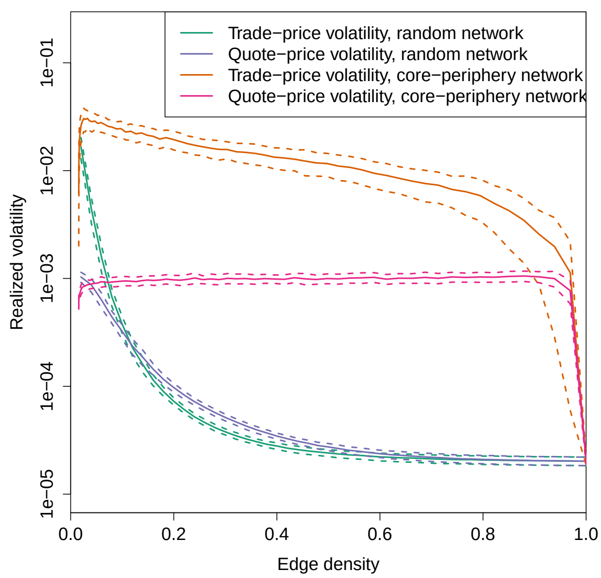

Figure 8: Realized trade-price volatility for Erdős–Rényi (green) and core–periphery (orange) networks, and realized quote-price volatility for the same Erdős–Rényi (purple) and core–periphery (pink) networks. Solid curves indicate the mean across all independent runs of the model; dashed curves indicate one standard deviation.

For the Erdős–Rényi networks, the trade-price volatility exceeds the quote-price volatility when the CCL networks have a very low edge density. As the edge density increases, the trade-price volatility decreases faster than the quote-price volatility. For edge densities larger than about 0.1, the trade-price volatility is approximately equal to the quote-price volatility.

This result is intuitively sensible because quote-price volatility is determined by the maximum among all buy prices (respectively, the minimum among all sell prices) in the market, and it is therefore influenced only by the extreme prices in the population. By contrast, trade-price volatility depends on the prices of all trades that are conducted by all institutions. As the edge density increases, the influence on the number of bilateral CCLs for the institution with the maximum among all buy prices (respectively, the minimum among all sell prices) 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, quote-price volatility is approximately stable across all edge densities, except for values of d that are close to 0 or 1. By contrast, trade-price volatility first increases sharply as d increases slightly above 0, and it subsequently decreases gradually and then decreases sharply as d increases beyond about 0.95. This result illustrates that, for the core–periphery networks, edge density has a much stronger influence on trade-price volatility than it does on quote-price volatility.

In a market with CCLs, one can regard trade-price volatility as the underlying volatility that is observable in the best quotes; there is then an additional contribution due to the CCLs. Therefore, consistent with our results for skipping costs (see Figure 7), Figure 8 suggests that as the edge density of a CCL network decreases, the strength of this additional impact on trade-price volatility from CCLs also increases for both Erdős–Rényi networks and core–periphery networks.

In Figure 9, we plot the log-ratio z of the realized quote-price volatility and the realized trade-price volatility:

$$z := \begin{cases} \log\left(v_B/v_b\right) , & \text{for seller-initiated trades} , \\ \log\left(v_A/v_a\right) , & \text{for buyer-initiated trades} . \end{cases} \tag{12}$$

A positive value of z indicates that trade-price volatility exceeds quote-price volatility, whereas a negative value of z indicates that quote-price volatility exceeds trade-price volatility.

For the Erdős–Rényi networks, the log-ratios are close to 0 (and sometimes even slightly negative) for edge densities that are larger than about 0.1, and they are positive for edge densities that are smaller than about 0.1. For the core–periphery networks, the log-ratios are positive, and they are often large for almost all edge densities. In these cases, the application of CCLs causes the volatility in the trade-price series to exceed the volatility that is 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.

5.7 Discussion of Simulation Results

Consistent with our expectations, we find that the more that institutions' CCLs restrict their access to trading opportunities, the less those institutions trade and the larger the skipping costs that they pay when they do. We also observe that CCLs have a considerable impact on trade-price volatility, but a much smaller impact on quote-price volatility: As the CCLs become progressively more restrictive, the log-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

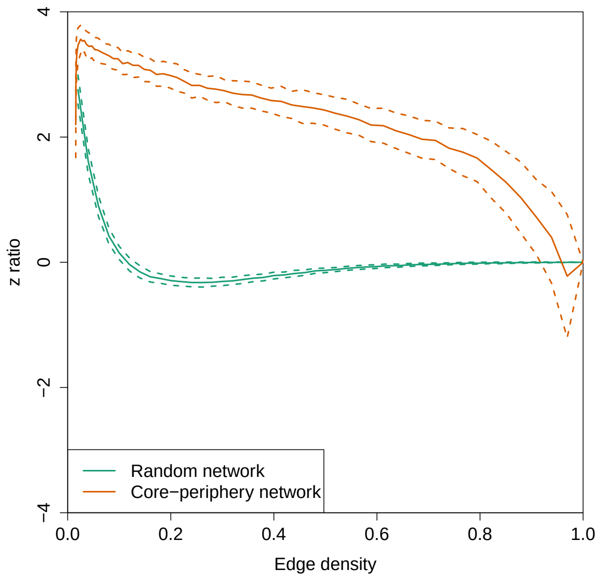

Figure 9: Mean of the log-ratio z of the realized quote-price volatility and the realized trade-price volatility for (green) Erdős–Rényi and (orange) core–periphery networks. Solid curves indicate the mean across all independent runs of the model; dashed curves indicate one standard deviation.

just its edge density, plays an important role in the determining the extent to which CCLs impact trade prices and volatility.[10] The strong variability in log-ratios that we observe in Figure 9 is attributable directly to the impact of CCLs in our model. This result is very interesting from a practical perspective, because it suggests that CCLs can strongly influence the price-formation process via institutions' access to the set of available trading opportunities.

Together, our model simulations suggest that when CCLs cause severe restrictions to institutions' access to trading opportunities, they can significantly impact both trade prices and volatility. Two features are particularly interesting. First, as the edge density of a CCL network decreases, 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 log-ratio decreases. 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 difference in the edge density d impacts skipping costs and trade-price volatility also requires knowledge of the CCL network's topology. Intuitively, this result implies that understanding the possible impact of CCLs in financial markets requires knowledge not only of how many institutions are trading partners, but also of which institutions are trading partners with each other.

It is also important to note some limitations of our approach. First, the institutions in our model do not attempt to implement trading strategies, nor do they adopt any strategic behaviour (such as not conducting a trade at a particularly bad price). Instead, we model the temporal evolution of their buy and sell prices via the simple, stochastic diffusions that we described in Section 5.1. The strategic actions that are undertaken by real market participants can entail the dynamics of real markets to differ considerably from those in our model. Second, we do not attempt to model the complex process by which real market participants choose their personal buy and valuations. In real markets, the probability of experiencing a counterparty

[10] We also investigated the volatility of our empirical price series from Section 4, but the results were uninformative and we do not report on them here.

default can be an important factor that influences a given institution's personal spread (i.e., its choice of s_t^i). By using CCLs to prevent trading with counterparties that it perceives to be unreasonably likely to default, a financial institution may choose to offer other financial institutions (i.e., those with whom it is willing to trade) a smaller spread than would otherwise be the case. In short, in our model, the progressively restrictive application of CCLs progressively increases the skipping costs of trades; however in real markets, the ability to implement CCLs may actually convince some institutions to narrow their spread and thereby offer other institutions better prices than would otherwise be available in the absence of CCLs.

We end by comparing our findings with those of two other models of how networks of relationships between different financial institutions can impact trade. First, we recall the model of Luu et al. [2018], which uses similar network topologies to investigate the dynamics of collateral, and the consequent systemic risk, in the presence of rehypothecation (see Section 2.2). Luu et al. [2018] reported that contagion effects vary much more rapidly as a function of edge density for a core–periphery network than they do for an Erdős–Rényi network. Consistent with these findings, we also find that the impact of CCLs on trade prices varies much more sharply as a function of edge density for a core–periphery network than it does for an Erdős–Rényi network. Luu et al. [2018] argued that "network structures with highly concentrated collateral flows [such as core–periphery networks] are ... characterised by a trade-off between liquidity and systemic risk". In other words, a core–periphery network is more preferable than an Erdős–Rényi network in terms of liquidity, but less preferable than an Erdős–Rényi network in terms of default contagion. We demonstrate that from the perspective of skipping costs, a core–periphery network with a given edge density produces a larger impact on trade prices (on average) than does a corresponding Erdős–Rényi network with the same edge density.

Second, we compare our results with those of Roukny et al. [2013], who studied how a trading network's topology can impact default cascades. Roukny et al. [2013] reported that when considering only load-redistribution (i.e., diversification) effects, network topology does not heavily influence their results. However, they reported that in the presence of a contagion effect, network topology becomes an important factor for default cascades. In our model, a larger edge density leads to a lower mean skipping cost — much like the version of the Roukny et al. [2013] model that considers only load-redistribution effects. However, our simulation results illustrate that the topology of a CCL network can significantly impact both skipping costs (see Figure 7) and the corresponding volatility of the trade-price and quote-price series (see Figures 8 and 9). Moreover, Roukny et al. [2013] argued that "hubs" (i.e., nodes with particular large degrees) promoting and inhibitory roles for market stability, because they both diversify shocks and serve as transmission hubs for default contagions. In our model, we see that a core–periphery network with lower edge density (i.e., with fewer core nodes and more periphery nodes) has considerably higher mean skipping costs than a core–periphery network with higher edge density. Therefore, in the context of our model, hubs always create a promoting effect, because they provide trading opportunities to a wide range of other financial institutions, many of which are otherwise very poorly connected to other institutions.

6 Conclusions and Discussion

We investigated how the application of CCLs impacts the prices of trades during the course of everyday trading. We first examined this issue empirically by studying a data set from a large electronic trading platform in the FX spot market that utilizes CCLs. Although we observed 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 this substantial heterogeneity in skipping costs is a natural consequence of the heterogeneity in the types and sizes of institutions that trade on the platform. By implementing CCLs, Hotspot FX can facilitate trade for a wide variety of different financial institutions while letting them 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, two recent Triennial Central Bank Surveys from Bank for International Settlements 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 [Bank for International Settlements, 2010, 2013]. Our findings are consistent with the hypothesis that a wide variety of different financial institutions, with access to different sets of trading opportunities, interact simultaneously on Hotspot FX.

To complement our empirical analysis, we also introduced a model of a single-asset market with CCLs. In our model, a network of CCLs gives explicit control over the interaction topology between different institutions. By fixing the model parameters and varying only the CCL network topology, we studied how CCLs impact trade in our artificial market. Our main observation is that both the edge density and the network topology are important for determining the skipping costs of trades and the corresponding volatility in the trade-price series. When the restrictions that are imposed by CCLs are particularly severe, they can lead to large skipping costs and a high level of volatility in the trade-price series, without causing a similarly high quote-price volatility.

Several possible extensions to our model may provide further insight into the impact of CCLs. For example, one can modify the temporal evolution of institutions' buy and sell prices to incorporate jumps or stochastic volatility to more closely reflect behaviour in real markets. There are also several possible ways to incorporate strategic considerations into our model. As one example, different institutions can implement different time-update rules for their buy and sell prices to reflect heterogeneity in their trading styles. As another example, each institution can also choose how to update its buy and sell prices according to its CCLs. For example, institutions may 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 illuminate several important questions about the impact of CCLs on everyday trading. To our knowledge, ours is the only empirical or theoretical study to explore the use of this mechanism in a quantitative framework. Our empirical results indicate that CCLs do not strongly impact the prices of the vast majority of trades during everyday trading on Hotspot FX. We therefore argue that one can regard the application of CCLs (and the consequent creation of skipping costs) as a necessary consequence of providing direct market access to a broader 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, skipping costs and trade-price volatility can escalate rapidly. In such situations, it seems plausible that the presence of CCLs may exacerbate systemic risk by severely restricting institutions' access to trading opportunities. Therefore, much like credit valuation adjustments and trade novation via a central counterparty (see Section 2), CCLs do not provide a simple solution to the problem of counterparty risk. However, our empirical results suggest that CCLs can be a sensible approach to this problem under normal market conditions, when CCLs do not overly restrict institutions' access to trading opportunities.

Our model simulations suggest that the impact of CCLs is determined not only by the edge density of a CCL network, but also by its topology. This presents a difficult question for regulators: How can one monitor the state of the CCL network between institutions in real markets in real time? In our opinion, to paint a realistic picture of market dynamics, financial stability policies must consider network effects.

Throughout this paper, our work has examined the question of how CCLs impact trade prices during the course of everyday trading. An important topic for future work is to analyze CCLs during periods of market stress. Specifically, it would be extremely interesting to assess whether institutions modify their CCLs during periods of market stress to reflect the heightened probability of experiencing a counterparty failure. It would also be interesting to examine whether (and when) such modifications cause a significant impact on the statistical properties of the trade-price series, when compared to those of the everyday series that we have studied.

An important open question is whether (and how) CCLs can 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 are still novated by a CCP. This configuration may, in principle, provide institutions with a double-layered protection: Trades are still novated by a CCP, but in the event of a CCP failure, institutions can ensure that they are exposed only to specified counterparties and only up to a pre-specified limit. Before such a configuration could be adopted, however, many important questions about the possible interactions between these two mechanisms need to be addressed. How should trades that are novated by the CCP count towards a given institution's CCLs? Should they count at all? Or should 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 observed 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.

Acknowledgements

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).

Appendix A: The Hotspot FX Data

In this appendix, we provide a detailed description of the data that forms the basis of our empirical analysis. We describe the same data in a more discursive fashion in Gould et al. [2016].

Sample Selection

During our sample period, three major multi-institution trading platforms dominated electronic trading volumes in the FX spot market: Reuters, Electronic Broking Services (EBS), and Hotspot FX.[11] All three of these platforms use similar trading mechanics; in particular, all

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

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 and 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 also 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 include 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 10, 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 3, we provide the corresponding summary statistics. Consistent with the market-wide volume ratios that were 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.

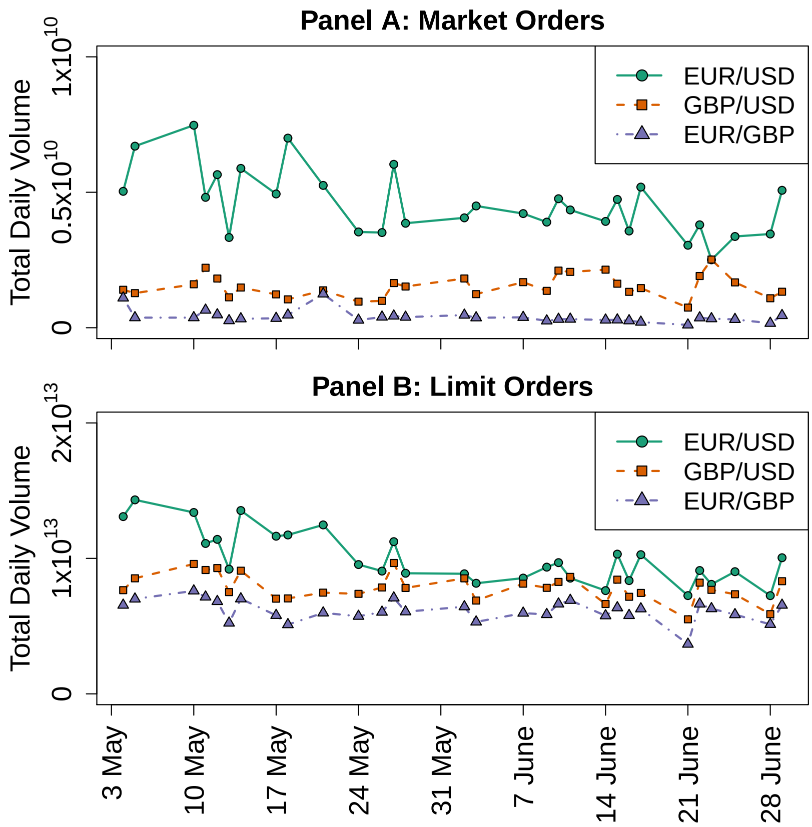

Figure 10: Total daily volumes of (top) market orders and (bottom) limit orders for (green circles) EUR/USD, (orange squares) GBP/USD, and (purple triangles) EUR/GBP activity on Hotspot FX on each day in our sample, measured in units of the counter currency. See Table 3 for the corresponding summary 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 3: Summary statistics for the total daily (A) volume of market orders, (B) number of market orders, (C) volume of limit orders, and (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.

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 or 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 or sell), and trade time. In both files, all times are recorded in milliseconds. For our three currency pairs, the platform's minimum order size is 0.01 units of the base currency, and the platform's tick size (i.e., the smallest permissible price interval between different orders) is 0.00001 units of the counter currency. For further details, 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, so we can 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 doing this 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. Data from platforms with larger tick sizes (such as Reuters and EBS) provide a more coarse-grained view of such price preferences; that makes results more difficult to interpret, particularly among trades for which the CCLs exert a small influence. 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.

For the purposes of our investigation, the Hotspot FX data also has some limitations.

First, the data does not provide any way to identify financial institutions, nor to ascertain which institutions participated in which trades. Therefore, our statistical analysis is limited to studying aggregate behaviour across all institutions, rather than more detailed conditional properties. Second, the tick-data files do not include 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. Third, 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 use 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 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., they "walk up the book"), we record the volume-weighted average price (VWAP) as the price for the whole trade, and we calculate the corresponding skipping cost using this VWAP price.[12]

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 in the present study. Due to the greater heterogeneity among member institutions on Hotspot FX than on Reuters or EBS (see Section 3), it seems reasonable to expect that CCLs have a larger impact on trade prices on Hotspot FX than they do on these other platforms. For example, large banks that trade on Hotspot FX may be unwilling to trade with small counterparties, and they 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 in the interbank market. By studying data from Hotspot FX, we are able to assess the impact of CCLs among a large and heterogeneous population.

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 strictly positive by definition, 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.

Similar to the application of CCLs, bid–ask bounce is a microstructural effect that impacts the trade-price series. Studying all buyer-initiated and seller-initiated trades together may cause bid–ask bounce to obscure the impact of CCLs in the trades that we observe. Throughout this paper, we thus study buyer-initiated and seller-initiated trades separately, in an attempt

[12] Because 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.

to disentangle our results about CCLs from the possible impact of bid–ask bounce.

Appendix B: Measuring Realized Volatility

In this appendix, we describe our methodology for measuring the realized variance of the quote-price and trade-price series in our model. 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 simulation of our model, let X denote the total number of buyer-initiated trades that occur, let $A_1, A_2, \ldots, A_X$ denote the prices of these trades, and let $a_1, a_2, \ldots, a_X$ denote the ask-prices immediately before the arrival of these trades. For a given number K of intervals and a given number L of subsamples, let

$$T := X/K$$

denote the sample width and let

$$\tau := T/L$$

denote the subsample width. For a given lag j, we calculate the *sell-side trade returns*

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

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

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

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

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

We calculate the *sell-side quote-price quadratic variation* v_b similarly from $a_1, a_2, \ldots, a_X$.

To identify a suitable value of K, we created volatility signature plots (see Andersen et al. [2000]) and chose values of K within a plateau. Other values of K in the same plateau produce qualitatively similar results.

References

K. Anand, P. Gai, and M. Marsili. Rollover risk, network structure and systemic financial crises. *Journal of Economic Dynamics and Control*, 36(8):1088–1100, 2012.

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/rpfxf10t.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.

M. Bardoscia, S. Battiston, F. Caccioli, and G. Caldarelli. Pathways towards instability in financial networks. *Nature Communications*, 8:14416, 2017.

S. Battiston, D. D. Gatti, M. Gallegati, B. Greenwald, and J. E. Stiglitz. Default cascades: When does risk diversification increase stability? *Journal of Financial Stability*, 8(3):138–149, 2012a.

S. Battiston, D. D. Gatti, M. Gallegati, B. Greenwald, and J. E. Stiglitz. Liaisons dangereuses: Increasing connectivity, risk sharing, and systemic risk. *Journal of Economic Dynamics and Control*, 36(8):1121–1141, 2012b.

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.

O. De Bandt and P. Hartmann. Systemic risk: A survey. Technical report, European Central Bank, available at `https://www.econstor.eu/bitstream/10419/152469/1/ecbwp0035.pdf`, 2000.

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.

P. Gai and S. Kapadia. Contagion in financial networks. In *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences.* The Royal Society, 2010.

P. Gai, A. Haldane, and S. Kapadia. Complexity, concentration and contagion. *Journal of Monetary Economics*, 58(5):453–470, 2011.

K. Giesecke and S. Weber. Cyclical correlations, credit contagion, and portfolio losses. *Journal of Banking and Finance*, 28(12):3009–3036, 2004.

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.

G. Iori, G. De Masi, O. V. Precup, G. Gabbi, and G. Caldarelli. A network analysis of the Italian overnight money market. *Journal of Economic Dynamics and Control*, 32(1):259–278, 2008.

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.

Knight Capital Group. Retrieved 14 April 2015 from `http://www.hotspotfx.com/download/userguide/HSFX/HSFX_UserGuide_wrapper.html`, 2015.

T. V. Koeppl. 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.

D. T. Luu, M. Napoletano, P. Barucca, and S. Battiston. Callateral unchained: Rehypothecation networks, concentration and systemic effects. *arXiv:1802.02127*, 2018.

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.

M. E. J. Newman. *Networks.* Oxford University Press, 2nd edition, 2018.

P. Norman. *The Risk Controllers: Central Counterparty Clearing in Globalised Financial Markets.* Wiley, West Sussex, UK, 2011.

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.

Puck Rombach, Mason A Porter, James H Fowler, and Peter J Mucha. Core–periphery structure in networks (revisited). *SIAM Review*, 59(3):619–646, 2017.

T. Roukny, H. Bersini, H. Pirotte, G. Caldarelli, and S. Battiston. Default cascades in complex networks: Topology and systemic risk. *Scientific Reports*, 3:2759, 2013.

J. E. Stiglitz. Risk and global economic architecture: Why full financial integration may be undesirable. *American Economic Review*, 100(2):388–92, 2010.

P. Tasca, S. Battiston, and A. Deghi. Portfolio diversification and systemic risk in interbank networks. *Journal of Economic Dynamics and Control*, 82:96–124, 2017.

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.crmpolicygroup.org/crmpg2/docs/CRMPG-II.pdf`, 2005.