BANK OF ENGLAND

Staff Working Paper No. 832

OTC microstructure in a period of stress: a multi-layered network approach

Andreas Joseph, Michalis Vasios, Olga Maizels,
Ujwal Shreyas and John Tanner

October 2019

BANK OF ENGLAND

Staff Working Paper No. 832
OTC microstructure in a period of stress: a multi-layered network approach

Andreas Joseph,[1] Michalis Vasios,[2] Olga Maizels,[3] Ujwal Shreyas[4] and John Tanner[5]

Abstract

Using unique data at transaction and counterparty identity level, we study the microstructure of the Swiss franc FX over-the-counter (OTC) derivatives market during a time of stress that was triggered by the decision of the Swiss National Bank (SNB) to remove the Swiss franc-euro exchange rate floor on 15 January 2015. Building on new methodology based on the topology of the trading network we segment the market into a multi-layered structure. We observe that the SNB announcement had a clear and differentiate impact on the market from this perspective. Clients in a more central position in the network topology were able to enter the market sooner than peripheral counterparties, while the inter-dealer core of the market was largely inactive. Using outstanding positions to proxy demand we observe that clients in greater need of trading were offered unfavourable prices if they found liquidity. Overall, our results point to a shortage of liquidity in the phase of market adjustment and highlight the heterogeneous reactions of different market segments during that time.

Key words: OTC derivatives markets, FX markets, market microstructure, networks, financial stability, market regulation.

JEL classification: G14, G23, G24, C80.

---

(1) Bank of England. Email: Andreas.joseph@bankofengland.co.uk
(2) Norges Bank Investment Management. Email: Michalis.v@gmail.com
(3) Bank of England. Email: Olga.maizels@bankofengland.co.uk
(4) Black Rock Investment Management & Financial Services
(5) Bank of England (retired)

The views expressed in this paper are those of the authors and do not necessarily reflect those of the Bank of England or its committees. The views do also not necessarily reflect those of Black Rock Asset Management & Financial Services or Norges Bank Investment Management. Ujwal Shreyas and Michalis Vasios have worked on this paper while at the Bank of England. Acknowledgement: We would like to thank Paul Bedford, Sujit Kapadia, Angelo Ranaldo, Darrel Duffe Graham Young, Filip Zikes, Rhiannon Sowerbutts and conference and seminar participants at the Bank of England, the US Office of Financial Research, the European Systemic Risk Board, the University of St. Gallen, the 2017 IFABS in Oxford, the Cambridge-INET Institute conference on the microstructure of FX markets, and the 8th IFC Conference at BIS for helpful comments and discussions. An earlier version of the paper was circulated under the title 'Gauging market dynamics using trade repository data: the case of the Swiss franc de-pegging'. The views expressed in this paper are those of the authors, and not necessarily those of the Bank of England or its committees. It is worth noting that while we analyse the impact of the event on the market and its visibility in the data, we are not commenting on the policy decision if the Swiss National Bank itself.

The Bank's working paper series can be found at www.bankofengland.co.uk/working-paper/staff-working-papers

Bank of England, Threadneedle Street, London, EC2R 8AH
Email publications@bankofengland.co.uk

ISSN 1749-9135 (on-line)

1 Introduction

On January, 15 2015 at 9:30 am UK time (GMT), the Swiss National Bank (SNB) surprised markets by announcing it was discontinuing the 1.20 EURCHF floor that had been in place since September 2011. The speed and scale of the subsequent currency move was unprecedented for a major currency making it one of the most stressful 1-day market episodes in the history of FX markets. The exogenous nature and magnitude of this event makes it well-suited for analysing the microstructure of an over-the-counter (OTC) market during a time of stress. How do the various types of market participants respond to a shock? Does the centrality in the trading network matter, and if yes, in what way? Do dealers discontinue intermediation and the provision of liquidity? What lessons can we learn?

To address these important questions, we provide the first micro-level study on the microstructure of the FX OTC derivatives market, and particularly the changes in the provision of liquidity and the trading network, around a market shock. Using unique and proprietary data at transaction and counterparty identity levels and network analytic techniques, we analyse trade-by-trade the impact of the SNB announcement. We find that immediately after the announcement dealers reduced liquidity provision. This coincided with the rapid currency appreciation by 41%. The impact was not homogeneous as the more central and better connected clients were more able to trade with dealers, while investors in the periphery of the trading network were practically absent from the market. Dealers with loss making outstanding positions were less active, while there was limited liquidity and higher prices for clients who wanted to hedge their positions by buying the Swiss franc, both suggesting a supply-driven liquidity crunch. When dealers started engaging more in the inter-dealer and the dealer-to-client market, the buying pressure on the Swiss franc eased, prices reversed and gradually stabilised. Our analysis highlights both the heterogeneous nature of OTC trading networks, with central nodes playing an critical role in liquidity provision and price stability, and the heterogeneous response to markets shocks.

The paper's significance is fourfold. First, analysing the microstructure of OTC derivatives market in a period of stress is important in its own right given the sheer size, the significant role in market efficiency and risk transformation, and the large efforts by regulators to reform these markets in response to the 2008 global financial crisis. Second, a ceteris paribus assessment of the impact of a market event on trading behaviour is hard to obtain. Fortunately, its large magnitude and the fact that it was not anticipated makes the SNB announcement an ideal environment for an event study analysis. Third, the OTC nature of the market makes it difficult to find granular transaction data. The European reporting obligation for derivatives transactions implemented as of 4 July 2012 made these unique data available to the Bank of England, which allows us to reconstruct the real trading network trade-by-trade around the SNB announcement. Fourth, the paper develops a new methodology for assessing trading behaviour from transactional data. Specifically, we use network analytic techniques in an event study context to evaluate the changes in the trading network and the role of the different network segments on liquidity provision and pricing.

Regarding the importance of the OTC derivatives market, OTC derivatives have been at the core of the post-crisis reform agenda. Most notably, in September 2009 G20 leaders agreed to push towards centralised trading and centralised clearing in an attempt to improve the resilience of the affected markets (FSB, 2010). Yet, the FX derivatives market, i.e., the second largest OTC derivatives market with 83 trillion USD of outstanding notional as of the end of 2014 BIS (2015), has not been subject to the new regulatory framework and remains to date largely decentralised, opaque, and reliant on the intermediation by big dealers for matching selling and buying interests (Duffie, 2012). By analysing FX derivatives contracts around a stressful event, our work sheds light on whether this more traditional structure is vulnerable to market shocks and helps informing the debate about future developments in OTC derivative markets. In fact, there is an ongoing discussion amongst regulators and market participants about developing and applying, on a voluntary basis, a set of global principles in the FX market (i.e., the FX Global Code, see Foreign Exchange Working Group, 2018) to promote effective and transparent trading and resilient infrastructure. Our study contributes to this important discussion.

Regarding the methodological contribution, we employ a new analytical framework that utilises the properties of the trading network. We classify market participants in different segments with respect to their centrality and importance to the network. The proposed classification has an intuitive interpretation and provides the basis of our event analysis, in which we examine trade-by-trade the impact of the removal of the Swiss franc floor on the 15th of January 2015 on the different segments or "tiers" of the network. The main innovation is that we do not need to rely on the static dealer-to-client classification to analyse liquidity provision. Instead, we let the network properties determine which entities are more important for "holding the market together" and which entities rely on financial intermediation for accessing OTC markets. Using such an approach in an event study context and at the intra-day level is novel and can be easily applied to analyse other market events.

To conduct our analysis, we analyse five months of EURCHF FX derivatives contracts between November 2014 and March 2015, provided by the Depository Trust & Clearing Corporation (DTCC) the largest European trade repository (TR). The DTCC data include EURCHF forwards and options where at least one of the counterparties is a UK entity and in addition to the standard information about volumes and prices they contain counterparty identities.[1] The time span intentionally starts a few months before and ends few months after the SNB announcement. This setting allows us to reconstruct the trading network and analyse the robustness of our methodology comparing periods before, during and after the event.

[1]Forwards and options account for about 80% of the global FX OTC derivatives market, BIS (2015). The DTCC data cover more than half of the global FX derivatives market (by comparing with statistics from BIS, 2015).

Some clear results emerge from our study. First, far from the schematic dealer-client structure in the standard microstructure models, we find that the FX OTC derivatives market operates under a complex, heterogeneous, multi-layered structure.[2] Using salient topological features of the trading network, we identify three distinct segments: the inner core, the outer-core, and the periphery. The inner core consists of mostly 15 large and well-connected nodes - the dealer-banks - accounting for more than 90% of the market, of which almost half represents interdealer activity. The client's segment is not homogeneous: The outer core largely contains banks, corporates and hedge funds, who trade frequently, are connected with multiple inner core nodes, but they do not interact with other outer-core nodes or with the periphery. Finally, there is a large number of small peripheral nodes consisting of smaller banks, real money investors, and corporates who only have a single connection going in either direction and exclusively rely on the inner core for sourcing liquidity. This structure is different from the one documented in other OTC markets, e.g., the municipal or corporate bond market, where the inter-dealer segment has a core-periphery structure surrounded by a homogeneous client segment.[3]

Second, we document a heterogeneous impact of the SNB event on the different segments of the trading network. In response to the SNB announcement, the provision of liquidity was limited, as evidenced by the low activity by the inner dealer core, and coincides with the rapid currency appreciation. In the first 20 minutes trading was mostly done by the outer client core entering the market. After this initial phase, the inner core become active, coinciding with the price reversal of the Swiss franc and eventual stabilisation. Although the outer-core was able to enter the market in the first hour of the event, the periphery which constitutes the largest fraction of market participants by number was relatively inactive until price stabilised.

[2] See for example the models in Stoll (1978), Amihud and Mendelson (1980), Kyle (1985), and O'Hara and Oldfield (1986) among others.

[3] See the evidence in Hollifield, Neklyudov, and Spatt (2017), Li and Schürhoff (2019), Di Maggio, Kermani, and Song (2017) for the core-periphery structure of the municipal bond, securitisation, and the corporate bond market, respectively.

Third, we find that dealers with loss making outstanding positions were far less active in the first hour of the event with only a single dealer with a net CHF short position entering the market. At the same time, the clients in the outer-core who wanted to hedge their short positions by buying the Swiss franc were less likely to trade in the first phase of the SNB event, when the need for hedging peaked. In fact, the outer-core entities who managed to buy the Swiss franc, they did so at unfavorable prices. These facts point to a supply-driven liquidity shortage, suggesting the important role of the inner-core nodes in our trading network for "holding the market together", particularly in periods of stress. This lends support to the papers by Duffie, Gârleanu, and Pedersen (2005) and Lagos, Rocheteau, and Weill (2011) who argue that illiquidity increases in stressful times when the matching between buyers and sellers is more difficult.

Our results offer practical insights into which theory models better describe the functioning of OTC derivatives markets in normal and stressful times.[4] On the one hand, search models (Duffie, Gârleanu, and Pedersen, 2005; Vayanos and Wang, 2007; Lagos, Rocheteau, and Weill, 2011; Hugonnier, Lester, and Weill, 2014) model trading in decentralised markets in the presence of search costs where prices are negotiated on a bilateral basis and depend on dealers' bargaining power. Recent extensions of these models predict a core-periphery structure, see Neklyudon (2012) and Weller (2014)). On the other hand, network-based models (Babus, 2012; Chang and Zhang, 2018; Wang, 2017; Babus and Kondor, 2018) introduce dealer centrality as an important parameter in the trading process and price formation.

Our findings support the theoretical predictions for a multi-layered (core-periphery) structure. The rapid currency movement and the dry-up of liquidity in the first phase of

[4]Of course, every period of stress has its own characteristics and might affect the trading network in different ways, but similar to other studies analysing single market events (Afonso, Kovner, and Schoar, 2011; Kirilenko, Kyle, Samadi, and Tuzun, 2017; Breedon, Chen, Ranaldo, and Vause, 2018; Menkveld and Yueshen, 2018; Hagströmer and Menkveld, 2019), our analysis should provide important complementary insights on the broader issues of market functioning during stressful times, especially for the under-explored OTC derivatives markets.

the event when dealers were more reluctant to trade, hence suggesting high search costs, offer some support to the search models.[5] However, we also find that investors' centrality to the trading network matters which is consistent with the predictions of the network-based model. One difference is that in the FX OTC derivatives market the inner core consists of a relatively small number of central dealers, who are fully connected between them. Therefore, centrality matters but not for the dealers, but instead for the clients, since it was mainly the central outer-core clients who could access the market in the aftermath of the SNB announcement. This heterogeneity of the clients segment with respect to their centrality is a missing feature of current theoretical models.

Our paper contributes to the empirical literature on the functioning of the OTC derivatives market. Benos, Wetherilt, and Zikes (2013), Loon and Zhong (2016), Fulop and Lescourret (2016), Morrison, Vasios, Wilson, and Zikes (2017), Riggs, Onur, Reiffen, and Zhu (2018) investigate different aspects of CDS markets, such as centralised clearing, transparency, contract standardization, and the transmission of counterparty risk. Podlich and Wedow (2014) asses contagion for dealer and non-dealers using CDS prices. Benos, Payne, and Vasios (2019) and Cenedese, Ranaldo, and Vasios (Forthcoming) look into the market liquidity and the pricing of interest rate swaps (IRS) in the new regulatory environment. We extend the prior studies by examining the microstructure of a previously unexplored market, the FX OTC derivatives market, which has not been subject to any post-crisis reform, such as an increase in transparency, or the introduction of centralised trading and centralised clearing. Additionally, unlike the other studies we employ network analytic techniques in an event study framework.[6]

[5]For example, Lagos, Rocheteau, and Weill (2011) predicts that during a market crash, dealers are less likely to accumulate asset inventories especially when the price correction is expected to be more permanent, which was the case with the Swiss franc event.

[6]Two excellent concurrent papers by Hagströmer and Menkveld (2019) and Breedon, Chen, Ranaldo, and Vause (2018) look at the impact of the Swiss franc removal event on the inter-dealer FX spot market and provide useful insights into the trading by high frequency traders and information revelation as the event evolved. However, their quotes data are anonymised and cover only a small fraction of the highly centralised, fast and electronic spot market.

A more recent strand of literature explores the network properties of OTC markets, with a focus on fixed income products. Hollifield, Neklyudov, and Spatt (2017), and Li and Schürhoff (2019) uncover a core-periphery network structure in the interdealer segment of the securitisation and municipal bond market, respectively. A similar picture is found by Di Maggio, Kermani, and Song (2017) in the corporate bond market. Our paper differs in at least three ways. First, rather than the inter-dealer segment, our focus is on the whole network, including clients who we show exhibit heterogeneity. Second, fixed income markets are essentially different from the FX OTC derivatives market. They are smaller in size, post-trade transparent, and with a large inter-dealer sector.[7] For example, Green, Hollifield, and Schürhoff (2006) report that there are about 600 dealers who intermediate in the US corporate bond market. In contrast, we show that the interdelear sector in the FX OTC derivatives market is very concentrated and homogeneous. Finally, our analysis of the trading network takes place at the intra-day level.

The remainder of this paper is structured as follows. In Section 2 we describe the data. Section 3 provides a granular overview of the Swiss franc (CHF) derivatives market from a network point of view. Particularly, we introduce a data-driven way to partition this network into a multi-layered structure. Section 4 contains the main event analysis, including the heterogeneous response of the trading network. Section 5 concludes and discusses lessons learned.

2 Data

We analyse five months of FX OTC data from between November 2014 and March 2015, provided by DTCC.[8] This trade repository (TR) data is available to us via the Bank of

[7]See Bessembinder, Maxwell, and Venkataraman (2006), Goldstein, Hotchkiss, and Sirri (2006), Edwards, Harris, and Piwowar (2007), and Green, Hollifield, and Schürhoff (2006) for the transparency status of municipal and corporate bond markets. Regarding the size, Di Maggio, Kermani, and Song (2017) report an average daily trading volume of $20 billion in the US corporate bond market compared to about $3.5 trillion daily turnover in USD FX OTC derivatives, BIS (2016).

[8]Depository Trust & Clearing Corporation

England as part of the European reporting obligation for OTC derivatives transactions and covers all trades where one of the counterparties is a UK entity (including subsidiaries of US bank in the UK).[9] They contain reports for forwards and options, with the forward category including both outright forwards and the forward leg of FX swaps. These two types of products account for about 80% of the global FX OTC derivatives market (BIS, 2015). The missing 20% consists of mainly currency swaps, which are reported separately as interest rate derivative products, which are outside the scope of this study.

The DTCC TR data are divided into two types of reports: a) activity reports, which contain trade information on flows, for example new trades, modifications, and valuation and cancellation updates; and b) state reports, which contain all end-of-day outstanding positions between individual counterparties. The state and activity reports are generated each day and are available with a one-day lag. They contain more than 100 fields that include information on trade characteristics and, more importantly, counterparty identities.

We went through several stages of data cleaning, involving filtering, deduplication and group consolidations, outlier detection and the dropping of missing values. The raw data contained about 100 million activity and 300 million state reports, out of which 3.5 million and 10 million contained at least one leg denominated in CHF, respectively. We filtered the activity reports to include only new trades, which left us with approximately 2.5 million reports. In the final step, we removed all duplicate reports, as well as reports from a small number of counterparties whose trade entries appeared incorrect. For state reports we only kept EURCHF trades.[10]

After filtering the data, we were left with a sample of 380,000 activity reports and

[9] The European reporting obligation has been implemented as part of the European Market Infrastructure Regulation (EMIR). The Bank of England's access to trade reports is as per the conditions stated in EMIR under Article 2 of Commission Delegated Regulation (EU) No 151/2013 - Data access by relevant authorities.

[10] Duplication is mainly due to two reasons. First, EMIR is a double-sided reporting regime, so the Bank would see two copies for a single executed trade when both the counterparties are UK entities (and, in this case, where they are both reporting the trade to DTCC). Second, as per the EMIR regulation, the activity reports could contain several copies of the same trade to reflect each of the modification, correction and valuation updates.

around 3 million state reports. To get an idea of the representativeness of the data, we compared them against the most recent BIS OTC derivative statistics BIS (2015). The filtered data contained \$0.9 trillion of Swiss franc forwards and swaps outstanding, and \$0.8 trillion of Swiss franc options outstanding as of 31 December 2014. According to semi-annual BIS data, on 31 December 2014 the notional outstanding of Swiss franc OTC FX derivatives was about \$4.2 trillion.[11] Of this, \$2.1 trillion was accounted for by forwards and swaps, and \$0.9 trillion by options. Although there might be some noise in this comparison due to differences in methodology, it seems that we see a significant portion of the global Swiss franc FX derivatives, especially in options, perhaps a result of London's status as a global centre for FX trading.

The next step in the cleaning process involved the grouping and classification of individual counterparties. Since, trades are reported at the legal entity level, we group single LEIs, which might belong to the same group, manually. Our sample contained about 9,000 active counterparties. To group individual counterparties into consolidated groups, we devised a three-step algorithm, consisting of name matching, intra-group flagging and a manual inspection. The first step created unique identifiers for group names. These were either extracted from a dictionary of names of the largest institutions or from "name stubs" created from removing general and finance-related fill words. The resulting groups were refined using the intra-group trading flag in the data. Note that the intra-group trading flag alone does not serve as a good group identifier in practice as it does not seem to be reported consistently. In a final step, we manually inspected the resulting consolidation groups, which confirmed that the great majority of institutions have been classified correctly in the previous two steps.[12] [13]

[11] This made up 5.5% of the whole USD 75.9 trillion OTC FX derivatives market, which in turn accounts for 12% of overall OTC derivative activity measured by notional size. According to BIS estimates, its semi-annual survey captures 90% of global OTC derivatives market.

[12] There has been progress recently in the consistent assignment and usage of legal entity identifiers (LEI). However, these covered only a fraction of entities in our sample.

[13] For example, having two entities hypothetically called "J.P. Morgan Wealth X Fund" and "JPM WealthX fund (Switzerland)", the name stub/ID would be "jpmorgan" in this case, the word "fund" is

Finally, we manually split the firms into different categories: (a) the G16 dealers, (b) other banks, (c) real money investors/other funds (including pension funds, insurance firms, asset managers, state institutions and unclassified funds), (d) hedge funds, (e) retail FX trading firms, (f) institutional FX trading (non-bank firms offering trading services or prime brokerage to institutional investors), corporates, and (g) others.[14]

Before conducting any in-depth analysis, we provide some first insights into the FX OTC Swiss franc segment. The Swiss franc market is a medium-sized FX market, with a total outstanding notional on 31 December 2014 of $900 billion and $800 billion for forwards and options, respectively. We observe about 3,700 Swiss franc trades per working day, with an average daily traded notional of $118 billion of which 88% represented trading in forwards and the rest in options.

The majority of Swiss franc derivatives trades we see were against the US dollar, accounting for 82% of all Swiss franc forward trades and 64% of all options trades in the market. When looking at the Swiss franc trades against the euro, we observe on average about 800 trades per day with 90% of them being forwards and an average daily trading notional of CHF17.3 billion . The average maturity of forwards is very short. Over half of all forwards have a maturity of no more than a week, and over 96% no more than four months. The maturity for options is slightly longer, although the most common maturity for options is between one week and two months. We can observe some bunching around three, six, nine and twelve month maturities.

suggestive of its type and the suffix "wealthx" may serve as an additional identifier. We would then check if we observe intra-group transactions for these entities and, if yes, group them together to a single identity. We do this under the condition that complimentary information, such as that from different LEIs, names or other intra-group trading patterns do not suggest otherwise. For example a reported intra-group transaction with one of these funds with a company hypothetically called "London Sewage Clearance Ltd." would raise suspicions regarding the correctness of the involved records. Therefore, all such consolidations of counterparties have been checked manually in a final step, especially when involving large and relatively active counterparties.

[14]The 'G16' category is motivated by the "Participating Dealers" in the NY FED's OTC Derivatives Supervisors Group (see https://www.newyorkfed.org/markets/otc_derivatives_supervisors_group.html), plus Standard Chartered: (in alphabetical order) Bank of America-Merrill Lynch, Barclays, BNP Paribas, Citigroup, Crédit Agricole, Credit Suisse, Deutsche Bank, Goldman Sachs, HSBC, JP Morgan, Morgan Stanley, Nomura, Royal Bank of Scotland, Société Générale, Standard Chartered and UBS.

G16 dealers make up a significant majority of trading in both forwards and options (75% and 80%, respectively). The rest of trading in the forwards market is split between banks, buy-side firms, service providers to institutional clients, corporates and others. Trading in the options market is more concentrated, with most non-dealer activity coming from hedge funds. In fact, trading in options makes up 37% of all activity done by hedge funds, by far the highest share of any of our categories. One potential reason for this could be that speculative activity was more likely to be present in the options market, while hedging activity by corporates and various funds was done in the forwards market. There is also a number of unclassified entities, which, however, are typically very small and collectively account for a only around 5% of notional traded.

3 The EURCHF network structure

In this section we introduce our methodology for analysing the network structure of the EURCHF FX OTC market. The analysis focuses on the EURCHF forward market from here onwards, which accounts for about 90% of the EURCHF market in terms of number of trades. We construct two types of networks reflecting the different types of reports under EMIR. There is a daily state network where edges consist of exposures between counterparties (nodes) at the end of the day and the intra-day trading network of executed trades between counterparties during a certain time window. Both networks can be seen as multi-layered either looking at edge or node properties. For instance, links of different maturities can span different sub-networks, while structural properties, i.e., the network topology and node properties therein, allow to separate nodes based on connectivity patterns.

Most of our analysis onwards is based on Figure 1, which is the network from state reports one day before the event. Each node is a counterparty with at least one outstanding position in the EURCHF forward market at the end of that day. Directed links represent aggregate positions between two counterparties, i.e., multiple positions results in a single link. An outgoing (incoming) arrow corresponds to a gross CHF short (long) position. The width of edges and node size is set proportional to gross Swiss franc exposures and the total degree of a node in the network. The layout is based on a force-based algorithm by Yu (2005) which highlights basic operational principles of the market. The market has a hub-centric structure, around a small number of highly connected market participants (highlighted in black) different "layers" within the remaining nodes (red and green). As it will turn out, this classification carries relevance as it allows us to differentiate between more and less active market participants, who might have responded differently to the event.

Figure 1 can be interpreted as a snapshot of trading activity. There is clear structure in the arrangement and sizes of the nodes, with large interconnected trading hubs in the middle connecting a large numbers of clients. This reflects the way OTC markets operate: trading is conducted via the so-call dealers who run the main trading platforms on a principal-to-principal basis and serve client demand. This demand is offset in the inter-dealer market.

We formally classify market participants using the network topology in Figure 1. We define a multi-layered 3-tier structure consisting of an *inner dealer core* (C1; black nodes), an *outer core* of central clients (C2; red nodes) and peripheral clients (P; green nodes). Both segregations between different layers have a functional interpretation.

The inner dealer core is identified using the degree distribution of the state network, i.e. the distribution of the sum of incoming and outgoing connections a node has. This distribution is heavily skewed at all times mirroring the dealer-client picture. We use k-means clustering to identify the two groups from the degree distribution, i.e. $k = 2$. This is a widely used clustering algorithm from unsupervised machine learning, see for example Hastie, Tibshirani, and Friedman (2001); Chakraborty and Joseph (2017). Its basic intuition is to separate observations into k clusters by minimising intra-group distances. The objective function is given by

$$ERR(X) = \frac{1}{n} \sum_{c=1}^{k=2} \sum_{x_i \in k_c} ||x_i - \mu_c||^2 \,, \tag{1}$$

where x_i is a single data input on which clustering is based on, μ_i the centre of cluster c and n is the number of observations, the number of nodes in the network in our case. We use the cumulative degree distribution as input, i.e. $x_i = (d_i, p(d \geq d_i))$, where d_i is the total degree of a node. This approach has the advantage that it can easily account for other node features, which may be derived from the network, e.g., different node centrality measures.[15]

The result from minimising (1) for the degree distribution of the network in Figure 1 is shown on the LHS of Figure 2. To stabilise the clustering we take logarithmic values as is standard when working with tailed distributions. The best-fit degree clusters are connected to each others. High-degree inner core (C1 nodes) are highlighted in red and constitute 1% of market participants (probability on vertical axis). A potential issue with this approach is the stability of the classification over time. However, we do not observe substantial changes in the network structure over time, in other words the hub-centric structure of Figure 1 is very stable. More formally, this is shown on the right panel of Figure 2. In about 90% of days in our sample the identified number of inner core nodes is 14 or 15 with only minor variations. From an operational point-of-view, we also observe that almost all dealer nodes identified in this way serve their respective (mostly out-bound facing) "clouds of clients". These are the single-connected nodes in Figure 1 highlighting the importance of inner core nodes for the provision of liquidity to a large number of peripheral market

[15] We also tested nodes segregations based on fitting a parametric distribution compared to the non-parametric k-means algorithm. This procedure is usually done using a power law or log-normal distribution for skewed data and returns a minimum degree value for which the distribution is valid, in our case defining the inner core. This has two drawbacks in the current situation. First, the tail cut-off degree number can vary substantially across time despite the network structure remaining comparable stable. Second and related, the number of identified C1 nodes is generally small (around ten) making statistical tests to differentiate between different distributions challenging. Given that different distribution carry different interpretations with regard to the underlying network (formation) dynamics, this is not satisfactory.

participants.

The split of client nodes in the different layers uses network topology too. The outer core (red C2 nodes in Figure 1) are defined by the strongly connected subcomponent of the graph, excluding the inner core. This is the subgroup of nodes where each node can be reached from any other node via a directed path across the edges of the network. Despite the abstractness of this definition, it has a clear interpretation within the market. Outer core clients are both selling and buying both currencies, and are often connected to multiple inner core nodes. On the contrary, the remaining periphery nodes (coloured green in Figure 1) have only uni-directional connections to the market to a single inner core dealer in most cases.

Regarding the types of participants found in the three layers, the inner-core consists chiefly, but not exclusively, of the G16 dealer banks. The other two layers of clients contain all the counterparty types defined in the previous section. One difference is that hedge funds are predominantly found in the outer core, while periphery contains more real money investors and corporates.

The three-tier structure of the EURCHF forward market of Figure 1 is stable over time, judged by the number of nodes and structure. However, it is possible that different market participants switch between network layers from one day to the other, which if true will introduce some noise in the proposed classification. We investigate this in Figure 3 which looks at the stability of the different layers over the whole sample period on the daily level. The upper panel shows counterparty fractions in all three layers. The inner core fraction (black line) is stable over time (note that the number of counterparties in the overall market does not vary over time). The proportions of outer core and periphery nodes is stable with some variation near the turn of the year and the CHF floor removal (vertical dotted line). In the lower part of Figure 3, we next look into the number of swaps between outer-core and periphery (green line) and the number of new entries (red line) in either of the two client layers as the fraction of total nodes. Both series suggest only very modest turnover in the outer layers, with a minor spike of new market entrants in the days after the event of about 12%. Both series are stationary, as a unit root can be strongly rejected in both cases. The observations from Figure 3 suggest that our proposed 3-tier network structure is stable over time.

Finally, the proposed network structure representation is also highly indicative of market activity. Across time about 60% of counterparties are classified as belonging to the periphery, while more than 90% of notional is traded in the core (either inner or outer-core), lending weight to this label.

4 The EURCHF floor removal

4.1 Reconstructing the event

We next investigate the EURCHF market dynamics after the announcement of the SNB at 9:30 UK time (GMT) on 15 January 2015 to discontinue the floor to the euro of 1.20 CHF/EUR which had been in place since September 2011. The speed and scale of the subsequent currency move was unprecedented for a major currency. The upper panel of Figure 4 shows the course of the EURCHF spot rate on the EBS trading platform based on individual transactions. The Swiss franc price peaked at 0.85, an appreciation of about 41%, just before 9:50, when it rebounded and stabilised at about 1.05 EURCHF at 10:20. Given the observed price evolution, we divide the event into three periods separated by vertical dashed lines in Figure 4: SNB announcement (9:30), CHF maximum price (9:50), CHF price stabilisation (10:20) and event end (12:00). The first three points correspond to CHF price developments, while the last point is set according to observed client trading activity as will be discussed below. These points in time define three event periods which we use to study the event. (I) CHF appreciation (red), (II) CHF depreciation (blue) and (III) post-event trading at new equilibrium price (green).

We use TR data and reconstruct the trading book on event day trade-by-trade. We start with traded prices for the five-day EURCHF forward market. This is the most liquid part of the Swiss franc OTC derivative market and the one closest to the spot market. This segment accounted for more than half (55%) of all EURCHF forwards traded on that day and for more so in the direct aftermath of the event. The middle part of Figure 4 shows all executed trades on the event day between 9:00 and 14:00. Apart from three outliers of relatively small notional value, the price in the forward market troughed at the same price as in the spot market (0.85 EURCHF) and at the same time (just before 9.50 am).[16]

The dot sizes in Figure 4 are set proportional to the CHF notional of each transaction. These are substantially smaller in periods I and II compared with the pre and post event average, suggesting a lack of liquidity in the direct aftermath of the announcement possibly contributing to the sharp price movement. The lower part of Figure 4 shows the trade-by-trade cumulative inventory of 5-day forwards by network segment. One first observation is that the dealer-to-client market was very thin in terms of notional traded in the first 20 minutes after the event with dealers' net position close to zero. When trading took place it was mainly between dealers and outer-core clients, while the periphery was largely absent until the beginning of period III.

These first insights suggest that the SNB announcement appears to have taken markets by surprise and made investors rush into the market to adjust their portfolios, take new hedges or engage in speculative activity. Hence, the sharp price decline to what seemed to be the new EURCHF fair value was expected to some extent: prices adjusted through the trading process. It is the V-shaped price movement, which is highlighted in the red-blue transition, which makes Figure 4 interesting. If markets were efficient, as economic theory suggests, we should not observe such an exchange rate fluctuation and divergence from

[16] We do observe a small number of trades executed within an hour after the announcement at prices which seem substantially higher than the prevailing price, around 1.15. These might be genuine trades or there might be some issues with the reported execution timestamp. Since these trades have relatively small notional values they do not have a significant impact on the analysis.

fundamentals. The small trade size, the thin trading in the dealer-to-outer-core segment, and the close to zero trading between inner-core and periphery, all point to a liquidity shortage between 9:30 am and 10:20 as a possible explanation.

We test this more formally by comparing liquidity levels before, during and after the event. Given the lack of bid and ask quotes data, we proxy effective spreads using metrics that require only executed prices. The first is the price dispersion measure of Jankowitsch, Nashikkar, and Subrahmanyam (2011), which we define as the average of the relative differences between individual execution prices and the average execution price on day t. More formally,

$$\text{Dispersion}_{i,t} = 10^4 \cdot \sqrt{\frac{1}{N_{i,t}} \sum_{s=1}^{N_{i,t}} \left(\frac{p_{s,i,t} - \bar{p}_{s,i,t}}{p_{i,t}} \right)^2}, \tag{2}$$

where N_t is the total number of trades executed for contract i on day t, $p_{s,i,t}$ is the execution price of transaction s, and $\bar{p}_{s,i,t}$ is the average execution price on contract i and day t. This is based on the premise that traded prices may deviate from the expected value of an asset because of inventory risk for dealers and search cost for investors.

The second is based on the effective spread measure from Roll (1984), which is based on the serial covariance of changes in prices. The measure is defined as:

$$\text{Roll}_{i,t} = 2 \cdot 10^4 \cdot \sqrt{-cov\left(R_{s,i,t}, R_{s-1,i,t}\right)}, \tag{3}$$

where $R_{s,i,t}$ is the intraday return between transaction $s-1$ and s, for trade i on day t. Both metrics are derived from market microstructure models and are commonly used in the context of OTC derivatives markets as proxies of transaction costs (see for example, Goyenko, Holden, and Trzcinka, 2009; Friewald, Jankowitsch, and Subrahmanyam, 2012; Benos, Payne, and Vasios, 2019 among others).[17]

[17]Note that when calculating the liquidity metrics and the realised variance, we remove all trades executed before 10.20 am on 15 January. We do so to control for the rapid Swiss franc appreciation immediately after the event, which, if included in our calculations, would artificially increase our metrics

Table 1 reports the results of an OLS regression of all market variables on a number of date dummies. The post SNB event dummy equals 1 after the event and 0 otherwise. The Wednesday dummy equals 1 on 14 Jan 2015 and 0 otherwise, while the Thursday and Friday dummies are defined in the same fashion. The results suggest that market illiquidity rose sharply on 15 January and remained at an elevated level for the remainder of our sample period. This impairment of market liquidity also evident in the case of the Roll effective spread measure. Trading activity measured by the daily number of trades and notional traded (see specifications 5-8 in Table 1) raised sharply on the event day, as expected, but dropped substantially after the de-peg. It is worth noting that the increase in the number of trades is much stronger than the increase in notional traded, which is reflected on the trade size. Indeed, the average trade sizes for EURCHF forwards dropped by 34% lower on the event day. All these effects are strongly statistically significant.

4.2 Response of network structure to the event

After establishing that there was a shortage of liquidity on the event day, we next examine how this relates to the intraday market dynamics on the event day using the three-tier segmentation from Figure 1. We are particularly interested in seing how the inner-core (dealers) of the trading network, who Lagos, Rocheteau, and Weill (2011) suggest will be less likely to provide liquidity in periods of crisis, responded to the SNB announcement. In addition, we investigate whether the two tiers of clients behaved differently, which if true would suggest that classifying clients using the network topology is a meaningful way for analysing trading behaviour.

We start with a visual investigation of trading dynamics as the event evolved. The upper panel of Figure 5 shows the number of different types of transactions by layer between 9:00 and 14:00 using a 30-min sliding window and 10-min time steps using the three-tier classification from Figure 1, i.e., inter-dealer (C1-C1, black), inner-outer core (C1-C2, red)

of transaction costs (particularly the price dispersion) and volatility.

and inner core-periphery (C1-P, green) trades.[18] Units on the vertical axis are relative to their daily means for each transaction type, which is given in the legend. The background colouring reflects the three periods of the event as defined in Figure 4.

Figure 5 provides some interesting facts. First, the inter-dealer activity (black line) was initially limited (Period I), but gradually increased and stayed at a relatively higher level throughout the CHF depreciation (Period II). More generally, in Period I there were only 4 active dealers trading clients, and only after the exchange rate fell below parity against the euro, more dealers entered the market and started providing liquidity by buying Swiss franc at unfavourable prices, even if the activity was potentially loss-making. Second, the majority of transactions in Periods I and II were between the inner and outer core, where it is worth noting that in Period I the average trade size was very small. Third, there is a clear difference in how the different layers and in particular the two tiers of clients (i.e., outer-core and periphery) reacted during the V-shaped price movement and particularly immediately after the floor removal at 9:30. Most of the immediate reaction came from C1-C2 transactions which constitute about 80% of transactions in periods I (red) and II (blue). On the contrary, most periphery nodes entered the market after the price stabilisation in period III (green).[19]

4.2.1 Trading patterns

We now have initial evidence that the two tiers of clients, defined by the network topology, reacted differently on the event day. We next examine the intraday network structure in more detail using a regression analysis. As discussed in the introduction, the exogenous nature of the event allows us to analyse the changes in the market structure in a period of stress. We therefore focus on the first phase of the event (9:30-9:50) during the rapid cur-

[18]Newly entering counterparties are classified as periphery nodes.

[19]The reversal to the mean of the number of both clients types (C2 and P) at the end of this period motivates the definition of the event's end point in Figure 4.

rency appreciation which was probably most severely affected by the liquidity shortage.[20] The model is the following logit regression:

$$logit(y_{i,p}) \quad = \quad \beta_l X_{i,k} + \gamma Z_i + \epsilon_i. \tag{4}$$

The dependent variable is a dummy variable which indicates if a trade was executed between 9:30-9:50 on the event day. The respective coefficients of interest are β_l for variables X_k, which capture the layer information. The layer information again uses the network of exposures the day before the event and is represented by dummy variables that equal one for inter-dealer ($C1-C1$), inner-outer core ($C1-C2$) or inner core-periphery ($C1-P$) transactions and zero otherwise.[21]. The controls Z are CHF notional, the EURCHF rate and the maturity class of a trade, where we distinguish between trades of maturity of one week (short maturity) and the rest.

The results in Table 2 confirm the key observations in Figure 5. The probability of trading in the first phase of the event is significantly higher for inner-to-outer-core trades than trades that involve periphery nodes. In fact the coefficient of the inner-core-to-periphery trades is negative suggesting that it was difficult for periphery nodes to trade immediately after the event. One explanation could be that outer-core nodes were more exposed to the currency appreciation (i.e., short positions) than periphery nodes. However, this is not the case, as all three layers have roughly equal numbers of counterparties with net short and long positions. Interestingly, the coefficient of inter-dealer trades is negative (but insignificant) reflecting the limited dealer activity. This might be explained by the fact that only one net short, hence with loss-making position, dealer (with 7 transactions) entered the market in the first 20 minutes after the event. This is consistent with anecdotal evidence that dealers were initially reluctant to trade and that the EURCHF floor removal

[20]The analysis have been conducted for all periods for which results will be provided upon request.

[21]About 20 trades we observe for that day are classified $C2-P$. It is not clear if these are genuine or missreported transactions and we drop them.

caused significant distress on dealers' trading platforms.

Table 2 suggests a lack of market liquidity when it was most needed immediately after the SNB announcement and a selective provision of liquidity to clients that are more central to the network and better connected (i.e., the outer core). However, a counter argument could be that the effects are driven by the lack of demand rather than the lack of liquidity supply. To distinguish between the supply and demand channels we test the following hypothesis: Out of the population of clients, the entities with net short Swiss franc outstanding positions were exposed to the currency appreciation the most and as a result they should have demanded liquidity immediately after the SNB announcement in order to close their loss making positions. In line with this, if liquidity was available we should have observed more trading from these entities in the first period of the event.

We test if this is the case using a variation of the baseline specification, where we include a dummy for clients with net short positions on the day before the event: this dummy is an indicator for CHF demand on the event day. The corresponding results are shown in Table 3. Here, we only consider C1-C2 trades as there were very few active periphery nodes between 9:30 and 9:50. All specifications control for the EURCHF rate, the notional and a dummy for a net short position of a dealer. A striking result in Table 3 is that net short clients where less likely to enter the market after the event. Given that common risk management practises would demand to close loss making positions as swiftly as possible, this result points to a lack of provision of liquidity in the first phase of the event. Supportive to the supply channel hypothesis is also the strongly significant and negative coefficient of C1 net short positions: Dealers with loss-making exposures to EURCHF were less active in the market during the currency appreciation.

4.2.2 Pricing patterns

We next move from analysing trading patterns to pricing. The likelihood of trading for counterparties in a certain position can also determined by the level of available prices.

For this, we introduce the *micro price movement* of each trade, defined as the difference of the EURCHF rate of a trade and the average rate of the last h trades before that trade. For trade i, we have

$$\dot{p}_i^h \equiv p_i - \frac{1}{h}\sum_{t=1}^{h} p_{i-t} = p_i - p_{m(h)} \,, \tag{5}$$

where $p_{m(h)}$ stands for the market price defined over the period h. The intuition behind $\dot{p}_i^h$ is that it provides an indication for the direction of the market movement triggered by a single trade by referencing it against earlier executed trades.[22] For the EURCHF rate (measured in CHF/EUR), a positive value of $\dot{p}_i^h$ means that the Swiss franc is cheaper compared to previous transactions, indicating an depreciation. A negative values indicates the opposite. It is not clear what h should be. Ideally it reflects the horizon over which prices incorporate new information via trading which is hard to measure. Effectively the horizon h sets the weight given in (5) of individual transactions against the past. A small h give larger weight to trade i, and vice versa. We set $h = 60$ and discuss this choice below.

The middle panel of Figure 5 shows MPM for all trades between 9:00 and 14:00 GMT on the event day. The colour coding corresponds again to the different layers. In line with the upper panel, outer core transactions constituted the majority of transactions after the SNB announcement (red dots), when MPM dispersion reached the highest level too. Leaving relations to other markets aside (particularly the EURCHF spot market), MPM is indicative of the direction the market is moving. This is shown in the lower part of Figure 5 which plots cumulative sums of consecutive MPM across the whole trading day for EURCHF forwards with up to one week maturity. The grey line represents the sum of micro movements for all trades, which exhibits a similar V-shape like overall price movement after the SNB announcement. The decomposition of this sum into the contributions from each layer of the three-tier market structure is illuminating of the roles individual counterparty types played during the event. These are shown by the black, red and green line for C1-C1,

[22] A number of trades as a reference, opposed to a fixed time window, better accounts for "market velocity", i.e. the prevailing level of market activity at a given time.

C1-C2 and C1-P trades, respectively. The net price movements can be attributed to core activity, while contributions from periphery nodes are negligible.

In equilibrium, price movements within the whole market balance each other. We use this criterion to set $h = 60$. For this value the cumulative MPM of inner and outer core transactions converge to roughly the same EURCHF rate change within period III (black and red lines).[23]

There are qualitative differences between inner and outer core micro price movements as expressed in the respective cumulative sums. The dealer market adjusted to the new exchange rate regime quickly around 9:50, while prices for the outer-core clients, often connects with multiple dealer, continued to appreciate, perhaps a result the limited available liquidity. We analyse the pricing dynamics more formally using the following model for the micro price movement of trade i within the core, i.e. C1-C1 or C1-C2,

$$MPM_i = \beta_p X_{i,k} + \gamma Z_i + \epsilon'_i \tag{6}$$

The coefficients β_p captures the effects of network layers, positions of counterparties and the direction of trades. We use the same controls Z as in Table 3 with the exception of the rate which is already contained in (5). We also include individual dealer fixed effects.[24] Same as before we focus on the first phase of the event, when the Swiss franc appreciated rapidly. The dependent variables is the mean-variance normalised (z-score) of MPM within this period. The results are summarised in Table 4.

The main driving forces in (6) can be attributed to outer core positions and trading direction (columns three and four). Outer core short positions are significantly related with movements stabilising the rabid appreciation of the Swiss franc as show by the positive coefficient. This can be understood as follows. The Swiss franc appreciation leads to loss

[23]We conducted robustness checks for a wide range of values between ten and 150. All presented results hold for values of $h \geq 20$.

[24]Despite clients transacting multiple times, the number of C2 is too large to include individual controls without leading to numerical instability, while the number of dealers is small.

making positions for these counterparties, such that price stabilisation at a cheaper Swiss level is generally in their interest. This finding is represented by the blue line in the lower part of Figure 5 which plot cumulative sums of micro price movements for C1-C2 trades where C2 has a short position.

When we take the direction of trading into account, we find that being short and buying the Swiss franc is associated with a negative sign, which indicates a more expensive Swiss franc compared to the market price defined by previous transactions. In other words, clients with short positions not only were finding it difficult to trade and close their loss-making position, but they did so at more expensive prices. Note that the reverse of specification (3) and (4) is true for CHF long positions of C2 counterparties, which are responsible for the majority of transaction in the market after the event suggesting that they were driving prices upwards. These are also the market participants most associated with cumulative micro price movements, i.e. contributing to the downward movement of outer core trading in the lower panel of Figure 5 (red line).

5 Conclusion

We present the first micro-level study on the microstructure of the FX OTC derivatives market. In particular, we investigate the changes in the provision of liquidity and the reaction of the trading network on 15 January 2015, when the Swiss National Bank surprised markets by announcing the discontinuation of the 1.20 EURCHF floor that had been in place since September 2011. We analyse how did the various types of market participants respond to this shock taking their position in the trading network matter into account. What lessons can we learn?

We look at the above event using unique and proprietary data at transaction and counterparty identity levels. Developing novel methodologies based on the topology of the trading network, we are able to analyse trade-by-trade the impact of the SNB announcement. We start by reconstructing the real network of exposures of the EURCHF FX Forward market using a force-based algorithm by Yu (2005). This highlighted basic working principles of this market, such as its hub-centric structure reflecting the traditional dealer-client view of OTC markets. But there is considerable heterogeneity in the client segment, which motivated the classification of counterparties based on the topology of this network. This led to the definition of a multi-layered structure consisting of the inner dealer core, the outer core of central clients and peripheral clients. Contrary to the classic dealer-client market structure representation, this 3-tier network exhibits a homogeneous large and well-connected inner-core of dealers, surrounded by an outer-core of active clients, often connected to multiple dealers, and "clouds" of hundreds of smaller clients with a single connection to the inner-core. The heterogeneous client segment is a feature not shared by other OTC markets, e.g., the municipal or corporate bond market (see Hollifield, Neklyudov, and Spatt, 2017; Li and Schürhoff, 2019; Di Maggio, Kermani, and Song, 2017).

Based on this perspective, we document a heterogeneous impact of the SNB event on the trading network. In response to the SNB announcement, we observe limited inter-dealer activity, while only the better connected and more central outer-core clients managed to enter the market. The periphery which constitutes the largest fraction of market participants by number was relatively inactive until the EURCHF price stabilised. We observe another heterogeneity with respect to clients' outstanding positions as a proxy for demand in the immediate aftermath of the SNB announcement. It was harder for clients who needed to trade the most (clients with net short outstanding CHF positions) to enter the market in the first phase of the SNB event, when the need for hedging peaked. In fact, the outer-core entities with short CHF positions who managed to buy the Swiss franc, generally did so at unfavorable prices. At the same time, dealers with loss making outstanding positions were far less active in the first hour of the event with only a single dealer with a net CHF short position entering the market. Collectively, these results point to a lack of liquidity supply in the aftermath of the SNB announcement.

Our findings lend support to the papers by Duffie, Gârleanu, and Pedersen (2005) and Lagos, Rocheteau, and Weill (2011) who argue that illiquidity increases in times of stress when the matching between buyers and sellers is more difficult. However, our results also suggest that investors' centrality to the trading network matters which is consistent with the predictions of network-based models. In our case, it is clients centrality that matters, since it was mainly the central outer-core nodes who could access the market in the aftermath of the SNB announcement. This heterogeneity of the clients with respect to their centrality is a missing feature of current theoretical models.

Our analysis has a number of policy implications. The limited provision of liquidity especially for the more peripheral and least connected clients, suggests that more transparency and more centralised trading (e.g., through multi-lateral trading platforms) could play a positive role and make markets more efficient and fairer. Enhanced transparency could also mitigate to some extent dealers advantage at setting prices. Benos, Payne, and Vasios (2019) provide support to this argument for the interest rate swap market.

Another implication is with respect to investors' ability to smoothly and on time unwind large positions during times of market stress. We show that it was actually the clients with loss-making positions, who find it harder to enter the Swiss franc market on the event day. This is something that policy makers should take into account for example when they assess banks' available liquidity to weather stressful events. Our proposed methodology based on network topology could support supervisors to identify the entities that is likely to be more adversely affected during market shocks. Furthermore, our analysis is based on a rich data source which became available following the reform efforts in international financial markets and regulation after the 2008 Financial Crisis. These newly available sources of information in conjunction with insights and methodologies developed in studies like the present one are likely to increase the resilience of financial markets to withstand future shocks.

References

Afonso, G., A. Kovner, and A. Schoar. 2011. Stressed, not frozen: The federal funds market in the financial crisis. *The Journal of Finance*, 66:1109–1139.

Amihud, Y., and H. Mendelson. 1980. Dealership market: Market-making with inventory. *Journal of Financial Economics*, 8:31–53.

Babus, A.. 2012. Endogenous Intermediation in Over-the-Counter Markets. *Working paper.*

Babus, A., and P. Kondor. 2018. Trading and Information Diffusion in Over-the-Counter Markets. *Econometrica*, 86:1727–1769.

Benos, E., R. Payne, and M. Vasios. 2019. Centralized trading, transparency and interest rate swap market liquidity: evidence from the implementation of the Dodd-Frank act. *Journal of Financial and Quantitative Analysis, Forthcoming.*

Benos, E., A. Wetherilt, and F. Zikes. 2013. The structure and dynamics of the UK credit default swap market. *Bank of England Financial Stability paper No 25.*

Bessembinder, H., W. Maxwell, and K. Venkataraman. 2006. Market transparency, liquidity externalities, and institutional trading costs in corporate bonds. *Journal of Financial Economics*, 82:251–288.

BIS. 2015. OTC derivatives statistics end-December 2015. .

BIS. 2016. Triennial Survey of FX and OTC derivatives trading. .

Breedon, F., L. Chen, A. Ranaldo, and N. Vause. 2018. Judgement day: Algorithmic trading around the swiss franc cap removal. .

Cenedese, G., A. Ranaldo, and M. Vasios. Forthcoming. OTC premia. *Journal of Financial Economics.*

Chakraborty, C., and A. Joseph. 2017. Machine learning at central banks. *Bank of England Staff Working Paper No. 674.*

Chang, B., and S. Zhang. 2018. Endogenous market making and network formation. *Working paper.*

Di Maggio, M., A. Kermani, and Z. Song. 2017. The value of trading relations in turbulent times. *Journal of Financial Economics*, 124:266–284.

Duffie, D.. 2012. *Dark markets: Asset pricing and information transmission in over-the-counter markets*, Princeton University Press.

Duffie, D., N. Gârleanu, and L. H. Pedersen. 2005. Over-the-Counter Markets. *Econometrica*, 73:1815–1847.

Edwards, A. K., L. E. Harris, and M. S. Piwowar. 2007. Corporate bond market transaction costs and transparency. *The Journal of Finance*, 62:1421–1451.

Foreign Exchange Working Group. 2018. The FX Global Code. Discussion paper, Global Foreign Exchange Committee.

Friewald, N., R. Jankowitsch, and M. G. Subrahmanyam. 2012. Illiquidity or credit deterioration: A study of liquidity in the US corporate bond market during financial crises. *Journal of Financial Economics*, 105:18–36.

FSB. 2010. Implementing OTC Derivatives Market Reforms. Discussion paper, Financial Stability Board.

Fulop, A., and L. Lescourret. 2016. Transparency Regime Initiatives and Liquidity in the CDS Market. *Working paper.*

Goldstein, M. A., E. S. Hotchkiss, and E. R. Sirri. 2006. Transparency and liquidity: A controlled experiment on corporate bonds. *The Review of Financial Studies*, 20:235–273.

Goyenko, R. Y., C. W. Holden, and C. A. Trzcinka. 2009. Do liquidity measures measure liquidity?. *Journal of Financial Economics*, 92:153–181.

Green, R. C., B. Hollifield, and N. Schürhoff. 2006. Financial intermediation and the costs of trading in an opaque market. *The Review of Financial Studies*, 20:275–314.

Hagströmer, B., and A. J. Menkveld. 2019. Information Revelation in Decentralized Markets. *Journal of Finance, Forthcoming*, pp. 19–2.

Hastie, T., R. Tibshirani, and J. Friedman. 2001. *The Elements of Statistical Learning* . Springer Series in Statistics, Springer New York Inc., New York, NY, USA.

Hollifield, B., A. Neklyudov, and C. Spatt. 2017. Bid-ask spreads, trading networks, and the pricing of securitizations. *The Review of Financial Studies*, 30:3048–3085.

Hugonnier, J., B. Lester, and P.-O. Weill. 2014. Heterogeneity in decentralized asset markets. Discussion paper, National Bureau of Economic Research.

Jankowitsch, R., A. Nashikkar, and M. G. Subrahmanyam. 2011. Price dispersion in OTC markets: A new measure of liquidity. *Journal of Banking & Finance*, 35:343–357.

Kirilenko, A., A. S. Kyle, M. Samadi, and T. Tuzun. 2017. The flash crash: High-frequency trading in an electronic market. *The Journal of Finance*, 72:967–998.

Kyle, A.. 1985. Continuous Auctions and Insider Trading. *Econometrica*, 53:1315–35.

Lagos, R., G. Rocheteau, and P.-O. Weill. 2011. Crises and liquidity in over-the-counter markets. *Journal of Economic Theory*, 146:2169–2205.

Li, D., and N. Schürhoff. 2019. Dealer networks. *The Journal of Finance*, 74:91–144.

Loon, Y. C., and Z. K. Zhong. 2016. Does Dodd-Frank affect OTC transaction costs and liquidity? Evidence from real-time CDS trade reports. *Journal of Financial Economics*, 119:645–672.

Menkveld, A., and B. Yueshen. 2018. The Flash Crash: A Cautionary Tale About Highly Fragmented Markets. *Management Science*.

Morrison, A. D., M. Vasios, M. I. Wilson, and F. Zikes. 2017. Identifying contagion in a banking network. *Bank of England working paper No. 642*.

Neklyudon, A.. 2012. Bid-ask spreads and the over-the-counter interdealer markets: Core and peripheral dealers. *Working paper*.

O'Hara, M., and G. S. Oldfield. 1986. The microeconomics of market making. *Journal of Financial and Quantitative analysis*, 21:361–376.

Podlich, N., and M. Wedow. 2014. Crossborder financial contagion to Germany: How important are OTC dealers?. *International Review of Financial Analysis*, 33:1–9.

Riggs, L., E. Onur, D. Reiffen, and H. Zhu. 2018. Swap Trading after Dodd-Frank: Evidence from Index CDS. Working Paper.

Roll, R.. 1984. A Simple Implicit Measure of the Effective Bid-Ask Spread in an Efficient Market. *Journal of Finance*, 39:1127–39.

Stoll, H.. 1978. The Supply of Dealer Services in Securities Markets. *Journal of Finance*, 33:1133–51.

Vayanos, D., and T. Wang. 2007. Search and endogenous concentration of liquidity in asset markets. *Journal of Economic Theory*, 136:66–104.

Wang, C.. 2017. Core-periphery trading networks. *Working paper*.

Weller, B.. 2014. Intermediation chains. *Working paper*.

Yu, Y. F.. 2005. Efficient and high quality force-directed graph drawing. *The Mathematica Journal*, 10:37–71.

Figures and Tables

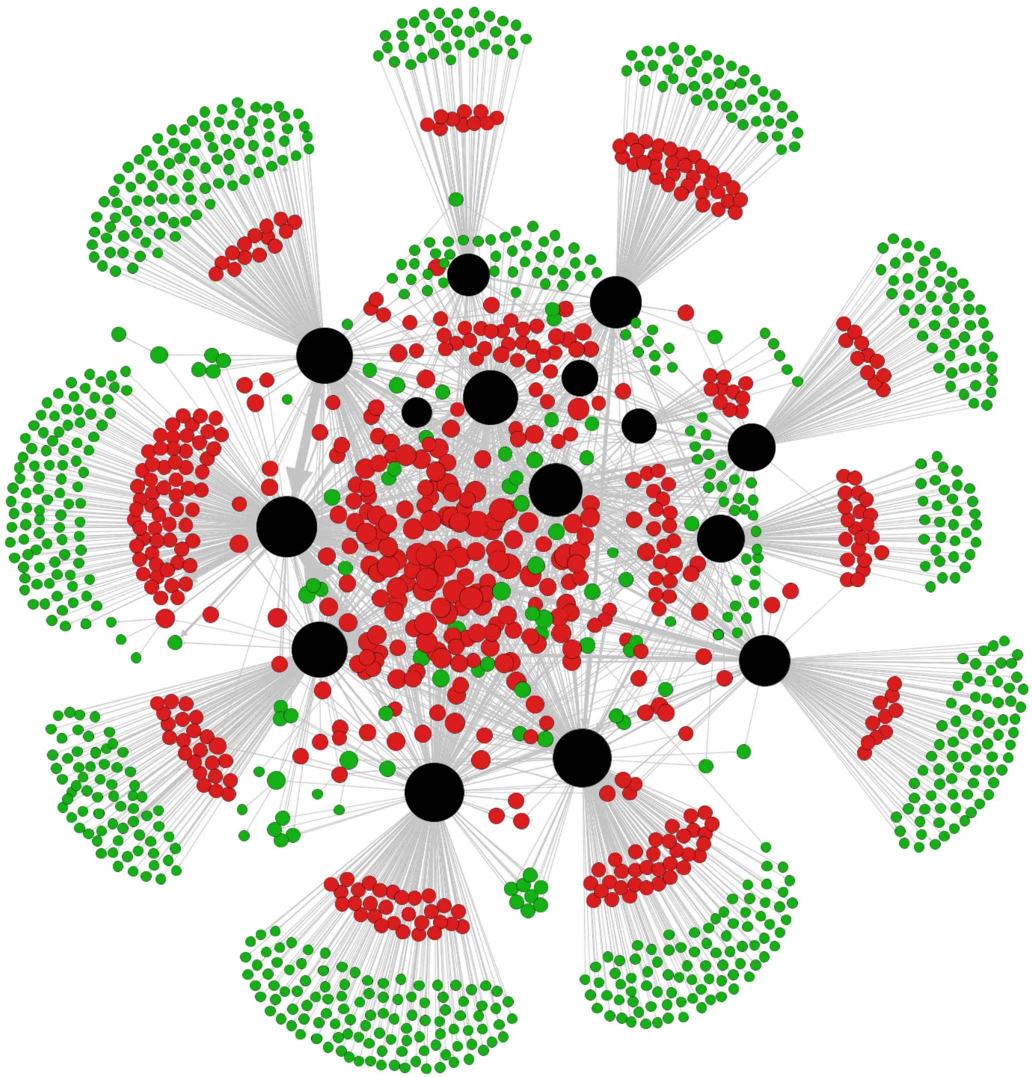

Figure 1: Network of EURCHF forward exposures from EMIR state reports at the end of 14 January 2015 showing all 1433 counterparties of that day. Colour coding indicates the 3-tier structure discussed in the text. The graph layout is based on Yu (2005). Source: DTCC and authors' calculations.

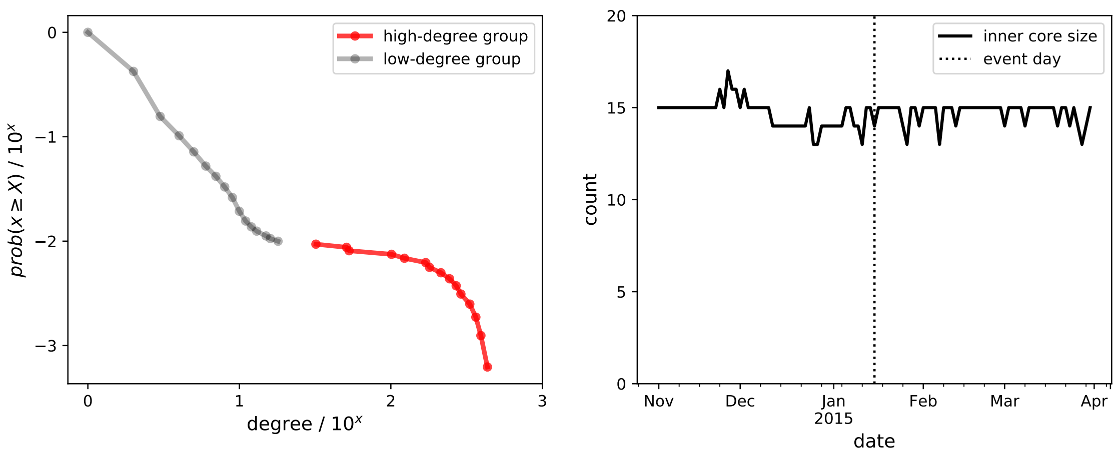

Figure 2: LHS: Cumulative total degree distribution of EURCHF forward network of open exposures on 14. January 2015. Colours indicate k-means separation of distributions into two clusters. RHS: Number of inner core (C1, red cluster on LHS) nodes in network on a daily basis during the full observation period from k-means clustering. Source: DTCC and authors' calculations.

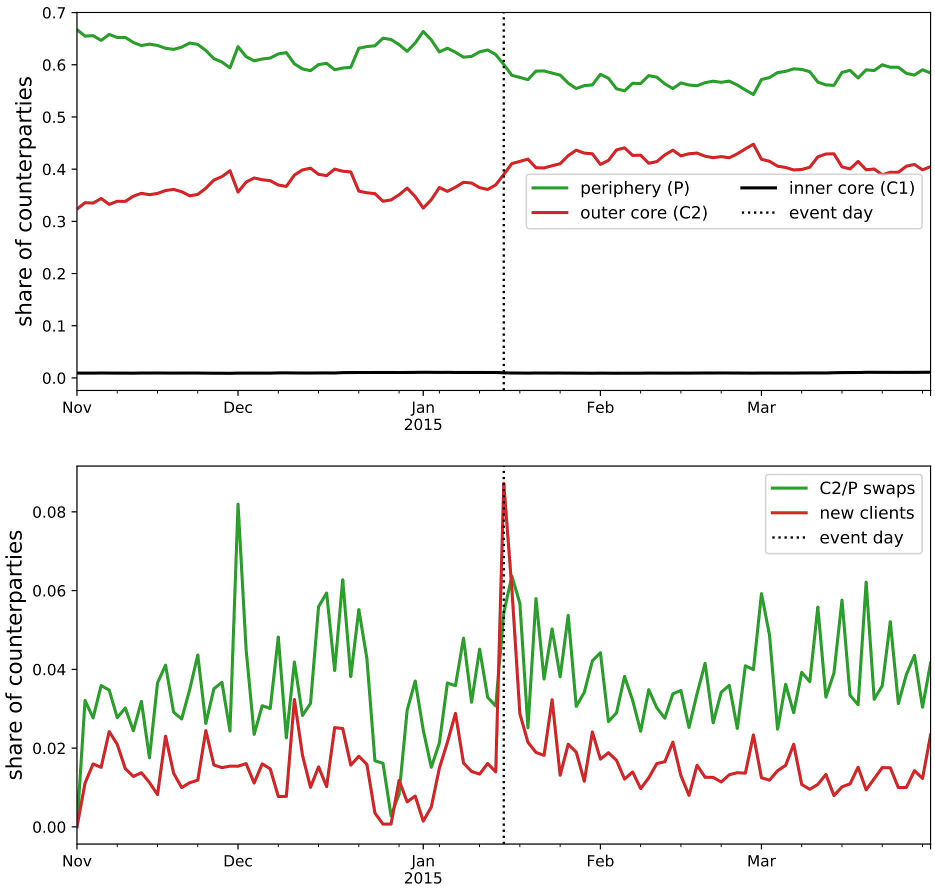

Figure 3: Upper panel: Counterparty fractions in inner dealer core (C1), outer client core (C2) and periphery (P) in the EMIR state report network at each trading day between Nov 2014 and March 2015. Lower panel: Bi-directional flow of clients between C2 and P as fraction of all nodes in the network (green lines) and fraction of new clients entering the market (red line). The event day is marked by the vertical dotted line in each panel. Source: DTCC and authors' calculations.

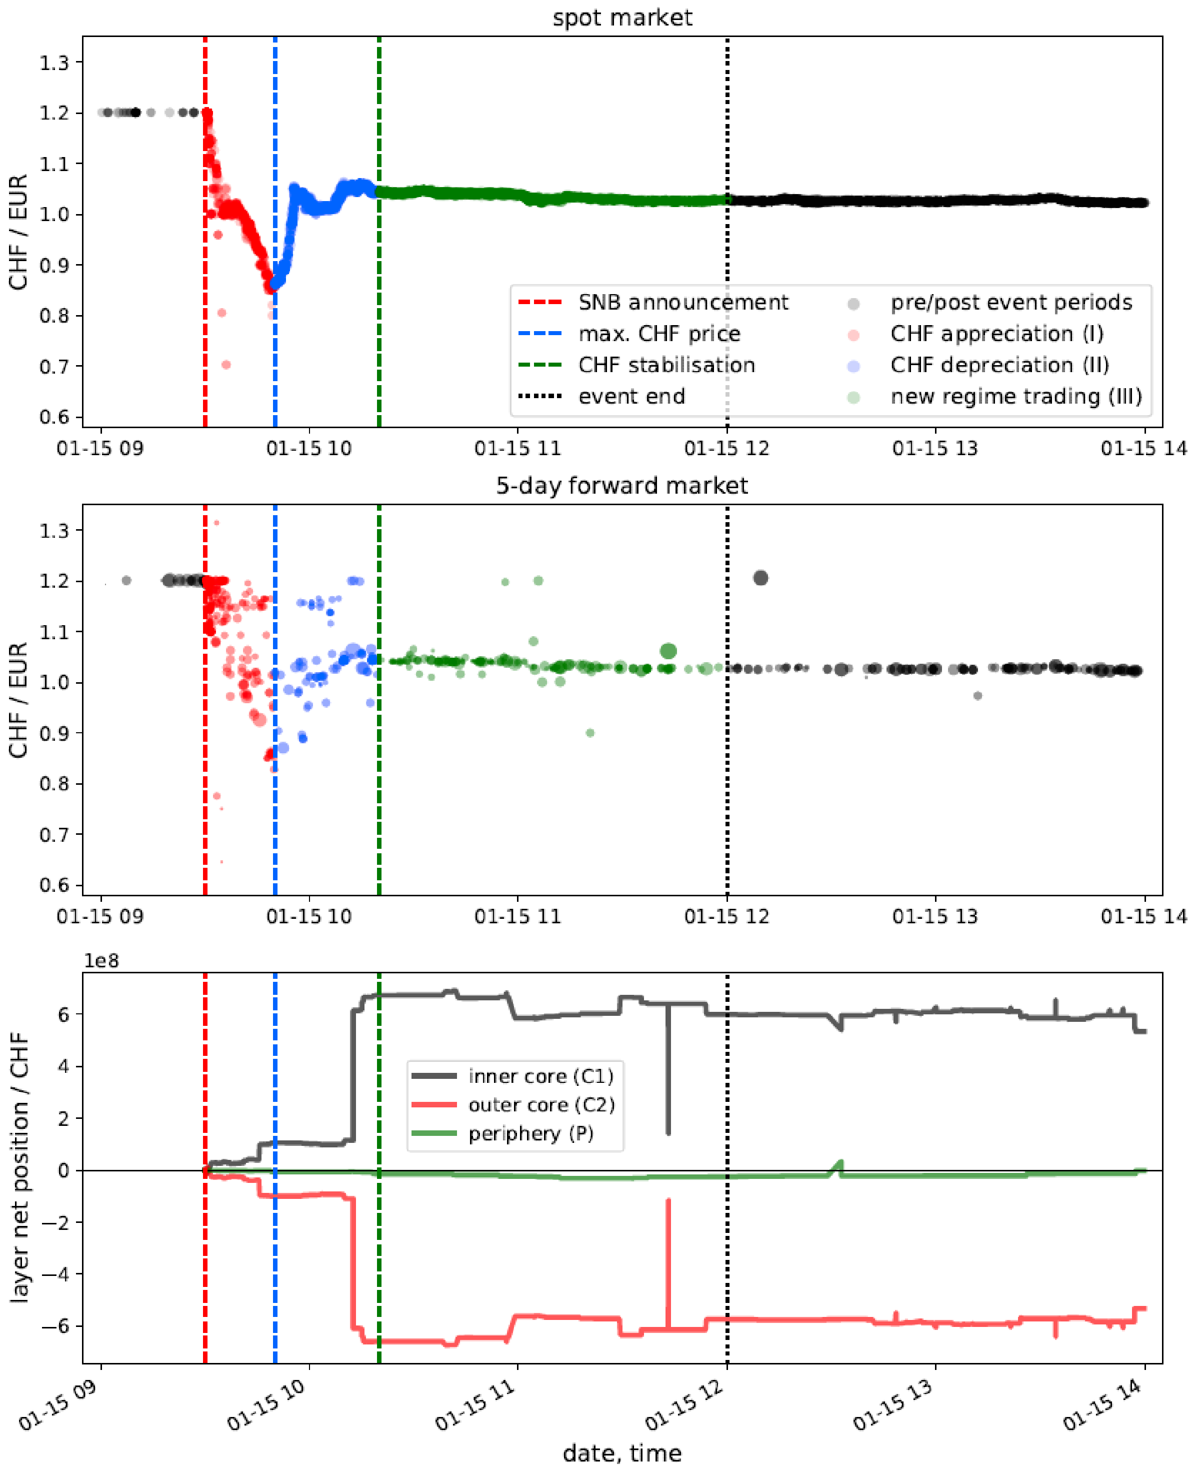

Figure 4: CHF/EUR rates and trading inventories for trading between 9:00 and 14:00 GMT on 15 January 2015. Upper panel: Spot market rates executed on EBS trading platform. Note: We have removed one EBS trade at 9.32 am with a price of just 0.0015. EBS confirmed that the highest value of the Swiss franc traded on the day was 0.85 franc per euro on the day. Vertical dashed lines and colour coding separate events and periods between them, respectively. Middle panel: 5-day forward market rates covered by EMIR TR data. Lower panel: Trading inventories of network layers for 5-day forward market corresponding to middle panel. Sources: EBS, DTCC and Bank calculations.

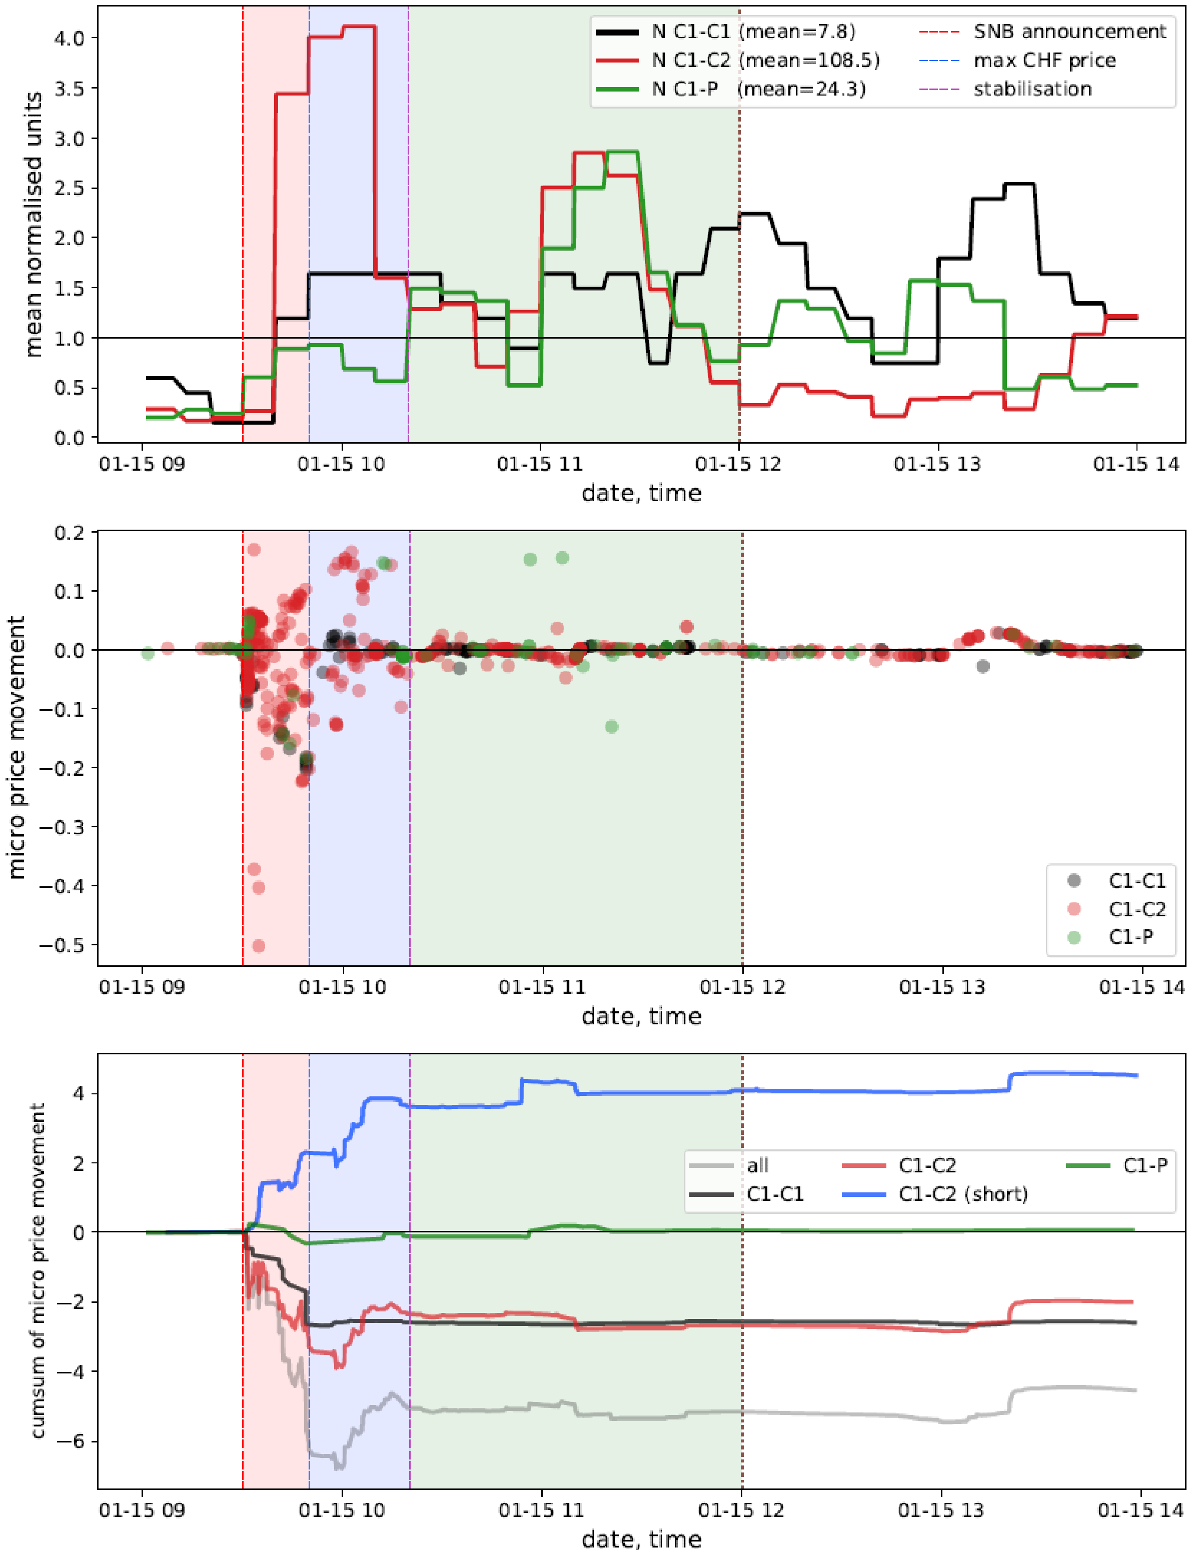

Figure 5: Upper panel: Relative number of transactions by 3-tier classification in 30-min moving-average trading network on 15 January 2015 (10-min step size). Units are relative to the daily mean for each transaction type, which is shown in the legend. Middle panel: Micro price movements (MPM) for each trade coloured by layer type. Lower panel: Cumulative MPM for different layer types. Outer core short CHF positions (blue line) are based on 14 January 2015 exposures. The different event periods are highlighted by different background colours: (I) 9:30-9:50 (red), (II) 9:50-10:20 (blue) and (III) 10:20-12:00 (green). Source: DTCC and authors' calculations.

	Dispersion		Roll		Nbr. trades		Notional $m	
Model	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Maturity	5-day	20-day	5-day	20-day	5-day	20-day	5-day	20-day
Constant	0.642***	0.326***	0.744***	0.941***	396.1***	128.6***	8783.4***	2847.5***
	[0.0200]	[0.0318]	[0.0683]	[0.0813]	[21.86]	[8.253]	[596.7]	[272.4]
Wednesday (14/01)	-0.0291	0.0696**	-0.527***	-0.377***	27.88	-20.59**	2110.1***	843.0***
	[0.0200]	[0.0318]	[0.0683]	[0.0813]	[21.86]	[8.253]	[596.7]	[272.4]
Thursday (15/01)	6.963***	0.913***	14.14***	0.701***	1268.9***	631.4***	10950.4***	1671.8***
	[0.0200]	[0.0318]	[0.0683]	[0.0813]	[21.86]	[8.253]	[596.7]	[272.4]
Friday (16/01)	2.106***	0.917***	8.978***	4.696***	73.88***	393.4***	-1437.3**	868.2***
	[0.0200]	[0.0318]	[0.0683]	[0.0813]	[21.86]	[8.253]	[596.7]	[272.4]
Post SNB event	0.206***	0.218***	2.211***	4.507***	-71.64***	-15.88	-3872.0***	-1185.6***
	[0.0416]	[0.0743]	[0.229]	[0.495]	[23.22]	[9.570]	[641.5]	[313.5]
Adj. R^2	0.919	0.116	0.713	0.434	0.578	0.713	0.344	0.128
N	104	104	104	104	104	104	104	104

Table 1: Regression analysis of the impact of the SNB event on market liquidity and trading activity. The post SNB event dummy equals 1 after the event and 0 otherwise. The Wednesday dummy equals 1 on 14 Jan 2015 and 0 otherwise, while the Thursday and Friday dummies are defined in the same fashion. We report results for EURCHF forwards with maturity up to 5 and 6-20 days. Robust standard errors are shown in square brackets. *, ** and *** denote significance at 10%, 5% and 1% levels, respectively.

Period I (9:30 - 9:50)				
Model	(9)	(10)	(11)	(12)
Constant	-18.673***	-18.528***	-18.870***	-18.326***
	[0.997]	[1.000]	[1.009]	[1.010]
CHR/EUR rate	1.105***	1.088***	1.033***	1.071***
	[0.050]	[0.051]	[0.052]	[0.049]
CHF notional	-0.008**	-0.007*	-0.004	-0.008*
	[0.004]	[0.004]	[0.004]	[0.004]
Short maturity	0.280***	0.281***	0.274***	0.266***
	[0.032]	[0.032]	[0.032]	[0.032]
C1-C1		-0.024		
		[0.018]		
C1-C2			0.072***	
			[0.015]	
C1-P				-0.119***
				[0.025]
pseudo R^2	0.323	0.323	0.334	0.336
N	2758	2758	2758	2758

Table 2: Summary table for logit model in period I with the dependent variables being one if a trade falls within that period. $X-Y$ are dummy variables which take a value of one if a trade is between X and Y, with $\{X,Y\} \in \{C1, C2, P\}$. The CHF notional are logged and normalised (z-scores). Short maturities is a dummy for trades with a maturity of up to one week. Robust standard errors are shown in square brackets. *, ** and *** denote significance at 10%, 5% and 1% levels, respectively.

Period I (9:30 - 9:50)			
Model	(13)	(14)	(15)
Constant	-19.494***	-19.631***	-19.247***
	[1.135]	[1.297]	[1.302]
CHF/EUR rate	1.372***	1.262***	1.244***
	[0.060]	[0.052]	[0.052]
CHF notional	-0.006	0.022***	0.022***
	[0.005]	[0.006]	[0.006]
Short maturity	0.312***	0.134***	0.117***
	[0.041]	[0.042]	[0.041]
C1 (short)		-0.313***	-0.319***
		[0.034]	[0.034]
C2 (short)			-0.045***
			[0.014]
pseudo R^2	0.338	0.415	0.421
N	1859	1859	1859

Table 3: Summary table for logit model with the dependent variables being one if a trade fall period I. All models only consider the C1-C2 part of the market and the C2 position. The CHF notional are logged and normalised (z-scores). Short maturities is a dummy for trades with a maturity of up to one week. Robust standard errors are shown in square brackets. *, ** and *** denote significance at 10%, 5% and 1% levels, respectively.

Period I (9:30 - 9:50)				
Model	(16)	(17)	(18)	(19)
CHF notional	-0.123	-0.123	-0.447***	-0.460***
	[0.104]	[0.104]	[0.112]	[0.113]
Short maturity	-0.981	-0.876**	-1.117	-1.133*
	[1.212]	[0.443]	[1.222]	[0.586]
C1 (short)		-0.524	-0.434	-0.427
		[0.411]	[1.205]	[0.549]
C2 (short)			1.405***	1.633***
			[0.163]	[0.186]
C2 (short, buy)				-0.441**
				[0.197]
dealer-FE	Yes	Yes	Yes	Yes
adj. R^2	0.010	0.010	0.198	0.204
N	322	322	322	322

Table 4: Summary table for micro price regressions during different period (I) for C1-C2 trades only. The dependent variable is the normalised micro price movement (5) of each trade (z-scores) for that period. The CHF notional are logged and normalised (z-scores). Layer dummies as in Tab. 3, buying is a dummy referring to buying the Swiss franc if short on it. Short maturities is a dummy for trades with a maturity of up to one week. All models control for different dealers (C1 nodes). Robust standard errors are shown in square brackets. *, ** and *** denote significance at 10%, 5% and 1% levels, respectively.