

School of Finance



University of St.Gallen

**JUDGMENT DAY: ALGORITHMIC TRADING AROUND THE
SWISS FRANC CAP REMOVAL**

**FRANCIS BREEDON
LOUISA CHEN
ANGELO RANALDO
NICHOLAS VAUSE**

WORKING PAPERS ON FINANCE NO. 2019/12

SWISS INSTITUTE OF BANKING AND FINANCE (S/BF – HSG)

JULY 2019



Judgment Day: Algorithmic Trading around the Swiss Franc Cap Removal¹

By

Francis Breedon,² Louisa Chen,³ Angelo Ranaldo,⁴ Nicholas Vause⁵

July 2019

Abstract

A key issue raised by the growth of algorithmic trading (AT) is how it responds in extreme situations. Using data on foreign exchange (FX) with a precise identification of AT, we find that AT contributed to the deterioration of market quality following the removal of the cap on the Swiss franc on 15 January 2015. Algorithmic traders withdrew market liquidity and generated uninformative volatility, both outside and during periods of perceived central bank intervention. We find that agency algorithms run by banks—rather than proprietary algorithms run by high-frequency traders—were particularly detrimental for market quality.

Keywords: Algorithmic trading, Swiss franc, market liquidity, price efficiency, central bank intervention.

JEL classification: G14, G23

¹ Any views expressed are solely those of the author(s) and so cannot be taken to represent those of the Bank of England or to state Bank of England policy. This paper should therefore not be reported as representing the views of the Bank of England or members of the Monetary Policy Committee, Financial Policy Committee or Prudential Regulation Committee. We are grateful to EBS for providing their anonymised and aggregated market data. EBS has not contributed to any of the analysis in the paper and does not endorse or support any of its conclusions. We are grateful to Alain Chaboud, Erik Hjalmarsson, Andrei Kirilenko, Dermot Murphy, Albert Menkveld, Sophie Moinas and Clara Vega, and to all participants at the 2016 conference on the Microstructure of Foreign Exchange Markets at Cambridge University, the 2016 conference on the Financial Determinants of Foreign Exchange Rates at the Bank of England, the 2018 Stern Microstructure meeting, the 2018 Market Microstructure Conference at the Institut Louis Bachelier, the 2019 workshop on Buy-side Execution in FX Markets at the Banque de France, the 2019 Western Finance Association meetings and seminars at the Swiss National Bank (SNB), the University of Bayreuth and the University of Sussex for their comments. We also thank Patrick Schaffner for excellent research assistance. Our title is inspired by that of a related paper, ‘Rise of the Machines: Algorithmic Trading in the Foreign Exchange Market’ (Chaboud et al. 2014). ‘Judgment Day’ preceded ‘Rise of the Machines’ in the *Terminator* film franchise.

² School of Economics and Finance, Queen Mary University of London. Email: f.breedon@qmul.ac.uk.

³ Department of Accounting and Finance, University of Sussex Business School, University of Sussex. Email: l.x.chen@sussex.ac.uk.

⁴ Swiss Institute of Banking and Finance, University of St. Gallen. Email: angelo.ranaldo@unisg.ch.

⁵ Bank of England. Email: nicholas.vause@bankofengland.co.uk.

1 Introduction

‘Algorithmic traders . . . are a larger presence in various markets than previously, and the willingness of these institutions to support liquidity in stressful conditions is uncertain’.

—Janet L. Yellen

New technologies have dramatically changed financial markets. One of the main innovations in recent years is algorithmic trading (AT), which broadly refers to the direct use of computers to implement trades. AT is now widely used by financial institutions, such as banks and hedge funds, and has important effects on the operation of financial markets. On one hand, it can improve market liquidity by reducing transaction costs and the reliance on financial intermediaries. It can also make security prices more efficient, in the sense that they better reflect fundamental values. It can even reduce risks due to human feelings that lead to, for example, panic reactions and herding behaviour. On the other hand, AT may not be socially desirable because it can increase market power (Hoffmann 2014), adverse selection (Biais et al. 2015), excess volatility and extreme market movements (Foucault et al. 2016)—as occurred in the recent ‘yen flash crash’⁶—and so can potentially threaten financial stability. We focus on this last issue in the current paper.

It is important to realise that investors use algorithms for many different purposes and strategies; therefore, the impact they have on market quality (e.g., liquidity, volatility and price efficiency) depends on why and how AT operates. This diversity contrasts with both the academic and popular literature, which has thus far focused on high-frequency proprietary algorithms. Hasbrouck and Saar (2013) and O’Hara (2015) indicate that a useful distinction is to categorise AT as ‘agency’ or ‘proprietary’. Agency AT is used by buy-side institutions (and the broker-dealer firms that serve them) to minimise execution costs and the price impact of their trading. Proprietary AT, in contrast, aims to profit from the computer trading process itself, using algorithms to identify profitable trades in the market and exercising tight risk control (Bank for International Settlements [BIS] 2011). Typical proprietary AT is undertaken by

⁶ On 3 January 2019, at about 22:40 London time, the yen (JPY) appreciated more than 3% against the dollar in just eight minutes, creating spill-over effects in other foreign exchange (FX) rates (*Financial Times* 2019).

sophisticated firms, such as hedge funds, that rely on low latency and are referred to as ‘high-frequency trading (HFT) firms’. In this paper, we offer some evidence on the differential role of these two broad categories of agency and proprietary AT and analyse how they contribute to foreign exchange (FX) market quality in terms of liquidity and price efficiency.

We analyse the role of AT alongside human trading in the FX market in a period containing the announcement by the Swiss National Bank (SNB) on 15 January 2015 that it had discontinued its policy of capping the value of the Swiss franc (CHF) against the euro (EUR). This ‘Swiss franc event’ represents a natural experiment as one of the largest shocks to financial markets in recent years.⁷ Also, it is the most significant ‘black swan’ event in the period in which AT has been a prominent force in the FX market. A detailed identification of AT and high granular order book data obtained from EBS, which is the leading platform for electronic spot FX trading in many of the major currencies, allows us to shed light on the contribution of AT and human traders to two important dimensions of market quality: liquidity and price efficiency.

A better understanding of whether AT is beneficial or detrimental to market quality in distressed situations is important for at least two reasons. First, being the first paper studying AT and human behaviours in an extreme event destabilising the FX market, our work may support the ongoing reform of trading venues (e.g., the Regulation National Market System [NMS] in the United States [US], and the Markets in Financial Instruments Directive [MiFID] I and II in Europe) and the efforts to regulate over-the-counter (OTC) markets (e.g., the Dodd-Frank Act, European Market Infrastructure Regulation [EMIR] and FX Global Code). The FX market is especially important because it is the world’s largest financial market,⁸ and it is crucial in guaranteeing efficiency and arbitrage conditions in many other markets, including those of bonds, stocks and derivatives. Also, as an OTC market, it is more fragmented and opaque, and it relies on bilateral clearing rather than central clearing (CCP). Moreover,

⁷ For instance, the Dow Jones index dropped by 9% within a few minutes during the 2010 equity ‘Flash Crash’ while the Swiss franc appreciated by more than 40% against the euro over a similar period (FXCM 2015).

⁸ The FX market operated at a daily average trading volume of more than five trillion United States dollars (USD) in 2016 (BIS 2016).

FX spot transactions demand little margin, allowing FX traders to take highly leveraged positions easily (e.g., Galati et al. 2007).

Second, the resilience of an exchange system depends on the market participant types, their shares and their reciprocal influence on each other. As pointed out by Yellen (2017), this is particularly relevant for financial stability because, for instance, the contribution of AT to offer liquidity in calm markets can disappear in distressed situations—that is, when other market participants most need it. If these adverse consequences of AT are predominant, or not offset by human trading or by the prompt interventions of, for example, a central bank, then AT could represent a systemic threat to the whole trading system.

We proceed in three steps. First, we describe our data set, which is representative of the interdealer FX market and at the core of the Swiss franc FX market. Specifically, our data come from the central limit order book for spot FX operated by EBS Service Company Limited, which is part of NEX Markets and a business division of the NEX Group plc. To introduce our analysis, we provide an overview of trading patterns conducted by AT and human traders, as well as the role of the SNB. Second, we perform an in-depth analysis of two important dimensions of market quality: market liquidity and price efficiency. By decomposing order flow and effective spreads by trader type, we highlight the contribution of agency and proprietary AT as well as human traders to liquidity provision and consumption. Third, to address the issue of price efficiency, we examine how AT and human traders contribute to (uninformative) volatility, to the price discovery process and to restoring arbitrage conditions.

Our study delivers three important findings. First, in reaction to the Swiss franc event, we find that AT, especially agency algorithms, significantly worsened market liquidity by reducing their supply and increasing their demand for liquidity—particularly in the direction of Swiss franc appreciation. Conversely, human traders supported liquidity provision. Our estimates of effective spreads show that all spreads exploded on the event day, particularly those of proprietary ATs, and they remained much wider in the post-event period. Second, the supportive role of AT in price efficiency in the pre-event period disappeared during and after the shock. Rather than leading the price discovery process and

closing arbitrage gaps, AT created uninformative volatility. Thus, our analysis delivers a consistent message: in extreme situations such as the Swiss franc event, the contribution of AT to market quality is reduced. Third, though our data set does not allow us to precisely identify SNB trades, our estimates indicate that the central bank's intervention was sizeable and diversified, implying that only part of it was conducted through the EBS platform. Our result showing AT's lack of contribution to sustaining market quality seems to hold true before, during and after central bank interventions occurred.

Of course, as in Tolstoy's comment on unhappy families, each period of market distress is distressed in its own way. Furthermore, the perfect identification of central bank activity and market participants' perception of it, as well as counterfactuals about what would have happened without it, are unachievable. However, like previous papers investigating important shocks, such as the failure of Lehman Brothers in 2008 (Afonso et al. 2011), the 'Flash Crash' in 2010 (Kirilenko et al. 2017) and the Swiss franc event itself (Hagströmer and Menkveld 2019), our analysis should improve our understanding of AT, extreme market events, central bank interventions and broader financial stability issues.

Our paper contributes to three strands of the literature: First, it contributes to prior research on AT. Apart from the important exception of Chaboud et al. (2014), prior research has focused on HFT, which is a specific category of AT.⁹ Without distinguishing agency and proprietary AT categories, Chaboud et al. (2014) show that AT tends to improve price efficiency in 'normal' times and conclude that, 'it cannot be ruled out that AT may, on rare but important occasions, contribute to excess volatility' (Chaboud et al. 2014, p. 2075). We fill this gap by analysing whether AT contributes to market quality in extreme events and find that AT, especially agency AT, was detrimental to market liquidity and both types of AT were unsupportive of price efficiency.

Second, we add to the empirical literature on market crashes and extreme events. By inferring the identity of those engaged in HFT from traders' behaviour (e.g., based on the number of trades per day

⁹ The 'official' definition of 'HFT' is somewhat unclear; the Nasdaq and official sources such as MiFID II effectively use a broad definition of 'HFT' that includes agency algorithms. In addition to characteristics such as sophisticated computer programmes and colocation, the US Securities and Exchange Commission (SEC) definition includes algorithms that initiate trades that rely on speed of execution (see SEC 2010). It is the latter definition we use here to distinguish HFT from other AT.

and net position of inventory), Kirilenko et al. (2017) analyse 15 putative high-frequency traders active on the CME E-Mini S&P 500 Futures market during the 2010 ‘Flash Crash’. Brogaard et al. (2018) study (intraday) short-lived large price movements across a few US stocks (3.5 stocks, on average). Our work differs from these papers in three important aspects: (i) we analyse AT and human traders in an OTC market that, by its nature, may be less resilient to shocks (Duffie 2012); (ii) we access the exact identification of traders’ categories; and (iii) we investigate a genuinely exogenous and market-wide shock.

Third, we contribute to empirical research on FX interventions conducted by central banks, recently surveyed by Fratzscher et al. (2019), by performing a microstructure analysis of (agency and proprietary) algorithmic and human trading during periods of central bank intervention.

2 Literature review and testing hypotheses

The growing literature¹⁰ on AT has tended to focus on HFT, which involves a specific category of proprietary ATs that operate high-speed technologies. The most commonly discussed issue is how HFT affects liquidity and price efficiency. The HFT community’s ability to revise their quotes quickly after the arrival of news reduces the problem of the winner’s curse but creates disincentives for trading for slower traders (Hoffmann 2014; Jovanovic and Menkveld 2016), including increased adverse selection and price impact (Foucault et al. 2016). In reaction to a fundamental shock, such as the SNB’s announcement, it is reasonable to assume that AT’s superior technology to adjust quotes and process information would aggravate the adverse selection affecting human traders. Thus, the first hypothesis is that, in reaction to the SNB shock, AT increased its market share relative to human traders and contributed more to informative trading and price efficiency.

¹⁰ Many excellent papers survey AT and HFT, including those by Biais and Woolley (2012), Chordia et al. (2013), Easley et al. (2013), Gomber et al. (2011), Goldstein et al. (2014), Jones (2013), Kirilenko and Lo (2013), Biais and Foucault (2014), O’Hara (2015) and Menkveld (2016, 2018).

In addition to increasing informational efficiency, the consensus is that HFT market-making enhances market quality by reducing spreads (O’Hara 2015).¹¹ However, it is not clear whether HFTs go with or lean against the wind—that is, whether HFTs amplify price falls (rises) by actively selling (buying), or dampen them by actively buying (selling) (Korajczyk and Murphy 2019; van Kervel and Menkveld 2019; Breckenfelder 2013; Tong 2015). A second hypothesis is that, in reaction to the shock, the share of the liquidity provision originating through AT changed.

While there is evidence that AT can be beneficial to market operation in normal circumstances, it is possible that the opposite may be true in unusual ones. On one hand, AT protocols are generally pre-programmed to operate in ‘normal’ market conditions (Chaboud et al. 2014), but AT may suffer reduced effectiveness or even cease to function in extreme conditions. On the other hand, human traders acting as dealers maintain a greater risk-bearing capacity and apply discretionary judgements (for example, to mutually share risk as a more efficient solution or to conduct better inference of unusual patterns, such as central bank interventions). For instance, Ait-Sahalia and Saglam (2014) and Rosu (2016) model HFT as averse to inventory risk and predict that volatility will lead high-frequency traders to reduce their provision of liquidity. Biais et al. (2015) find a role for HFT in fragmented markets, an issue that produces adverse selection, negative externalities and over-investment in equilibrium.¹² Thus, an alternative hypothesis is that there is a reduction in the activity and contribution to market quality of AT relative to human traders in extreme situations.

By analysing a ‘black swan’ event, our paper is mostly related to the literature on stock markets and future markets in distressed times. Prior research provides mixed evidence of the role of AT (more specifically, HFTs). On one hand, HFTs withdraw from their market-making role during flash crashes (see, for example, Commodity Futures Trading Commission [CFTC]-SEC 2010; Easley et al. 2012; Menkveld and Yueshen 2018) or when market conditions become unfavourable (see, for example,

¹¹ See, for example, Boehmer et al. (2018a, 2018b), Brogaard et al. (2015), Carrion (2013), Conrad et al. (2015), Hasbrouck and Saar (2013), Hendershott et al. (2011), Menkveld (2013) and Malinova et al. (2018).

¹² The role of AT in fragmented markets with multiple exchanges is studied by Pagnotta and Philippon (2018). Other papers analyse the welfare implications from double auctions (e.g., Cespa and Vives, 2015; Du and Zhu, 2015) and asynchronous arrivals (e.g., Budish et al., 2015; Bongaerts and Van Achter, 2016). Bernales (2014) and Rojcek and Ziegler (2016) use numerical methods for dynamic models encompassing the endogenous role of HFT, central limit order books and latency.

Raman et al. 2014; Anand and Venkataraman 2014; Korajczyk and Murphy 2019). On the other hand, HFTs provide liquidity and absorb imbalances created by non-high-frequency traders around large price movements (though only for single stocks, not for broader events) (Brogaard et al. 2018), and HFTs do not change their trading pattern when prices fall during flash crashes (Kirilenko et al. 2017).

In comparison with these papers, one of our key contributions is the analysis of an extreme event in the FX market. This differs from the markets studied in prior AT/HFT research in many important respects, including that (i) leveraged FX traders almost exclusively transact OTC (without clearing via CCPs), (ii) it is dominated by global FX dealers and (iii) central banks can intervene to directly affect FX rates. The first two characteristics may lead to ‘liquidity spirals’ (Brunnermeier and Pedersen 2009; Brunnermeier et al. 2009), causing sudden losses, margin increases and liquidity dry-ups (Mancini et al. 2013), thereby impairing the risk-bearing capacity of international FX dealers (Gabaix and Maggiori 2015) and increasing the risk of coordination failures when speculative positions such as ‘carry trade’ are unwound (Plantin and Shin 2011). The only earlier paper that provides an in-depth analysis of AT in the FX market is that of Chaboud et al. (2014), who show that AT running in ‘normal’ conditions is associated with a reduction in arbitrage opportunities and that AT liquidity provision decreases return autocorrelation.

3 Market structure and data

3.1 Market structure

At the core of the FX market, there is the interdealer segment (Lyons 2001). In it, EBS and Reuters are the two significant global electronic spot-trading platforms in major currency pairs. USD/CHF and EUR/CHF, the focus of this study, primarily trade on EBS (see King et al. 2012), and prices on EBS also constitute the reference values for derivative pricing in these currencies. Moreover, during the period of the Swiss franc event, EBS was the key trading platform for all Swiss franc positions, as the

trading of futures on the Chicago Mercantile Exchange was suspended and OTC trading largely disappeared (Hagströmer and Menkveld 2019).

EBS is an order-driven electronic trading system that unites buyers and sellers of spot FX around the globe on a pre-trade anonymous central limit order book. EBS is accessible to FX dealing banks and, under the auspices of dealing banks (via prime brokerage arrangements), to hedge funds and commodity trading advisors (CTAs). EBS controls the network and each of the terminals on which it is recorded whether a trade is conducted by an ordinary keyboard ('manual' or 'human' trades) or by a direct computer interface ('algorithmic' trades).

As well as this simple distinction by terminal type, EBS requires market participants to identify in which types of trading they engage. This allows EBS to decompose the algorithmic trades into two different categories, referred to as 'bank algorithmic interface' ('bank AI') and 'professional trading community AI' ('PTC AI'), which is how EBS labels the direct computer interfaces.¹³ 'PTC' essentially refers to non-bank entities, such as hedge funds and CTAs, that can directly trade on EBS under the auspices of dealing banks (via prime brokerage arrangements). Trading through this route is all proprietary AT in the sense that it does not include trades directly driven by client demand. HFT algorithms contribute significantly to this type of trading. Bank AI is a mixed category dominated by agency AT. This includes a significant share (EBS estimates approximately 30%) of simple aggregators, which are computer-based trading systems that simply aggregate and process orders received from bank customers and from other types of agency AT, such as hedging and auto liquidation algorithms. This latter category of agency AT has been highlighted as a source of price distortions during other crashes (e.g., McCann and Yan 2015; BIS 2017). While it is likely that bank AI also contains some proprietary AT, it is dominated by agency AT. Hence, we treat it as representative of that category of trading. Figure 1 illustrates the three trade categories and their shares of average transaction volumes in EBS. For the rest of the paper, we discuss bank AI as agency (bank) AT and PTC AI as proprietary (non-bank) AT.¹⁴

¹³ Market data at this level of granularity are not ordinarily sold or distributed by EBS to third parties.

¹⁴ Prior research shows that agency algorithmic execution orders tend to be larger in size and are executed at a lower frequency (BIS 2011), whereas HFT has higher trading volumes and numbers of trades but a smaller trade size than do non-HFT firms or large institutional investors (see, e.g., Aitken et al., 2014; Subrahmanyam and Zhen, 2016; Korajczyk and Murphy, 2019; van Kervel and

Computer-based trading classified as PTC or bank AI accounts for approximately 70% of EBS market's transaction volume. However, the system includes features designed to prevent strategies by which speed or low latency is the sole contributor to success. First, it imposes a minimum quote life (MQL) for the five core currency pairs on EBS, so that once a limit order¹⁵ is submitted, it cannot be cancelled for 250 milliseconds.¹⁶ Second, and more importantly, EBS operates a randomised batching window on all messages that enter its matching engine, referred to as a 'latency floor'. This mechanism generates batching windows of 1–3 milliseconds,¹⁷ during which time messages are processed on a randomised basis. As a result, the first message to arrive may not be the first released, and sub-millisecond differences in latency become less important for trading on EBS.¹⁸

Because EBS is a 'wholesale' trading system, the minimum trade size over our sample period is one million of the base currency,¹⁹ and trade sizes are only allowed in multiples of millions of the base currency.

3.2 Data

Our data consist of both intraday quotes and transactions for EUR/CHF, USD/CHF, EUR/USD, USD/JPY and EUR/JPY from 5 to 23 January 2015. This paper focuses on EUR/CHF and USD/CHF, and additional findings on the remaining currency pairs are summarised in 'Non-CHF foreign exchange rates' in the Online Appendix. We specify 15 January as the Swiss franc event day, 5–14 January as the pre-event period and 16–23 January as the post-event period. Following the convention in the literature

Menkveld, 2019). The descriptive statistics of these features in our data confirm that among the three categories, PTC AI is characterised by the highest average daily trading volume and number of trades but the smallest average daily trade size. See Table OA.2 in the Online Appendix for further details.

¹⁵ We refer to 'limit orders' as those termed 'good-till-cancelled' ('GTC') orders by EBS.

¹⁶ An MQL of 250 milliseconds was applied in the EUR/USD, USD/JPY, USD/CHF, EUR/CHF and EUR/JPY currency pairs on EBS for the dates referenced in this paper.

¹⁷ This is for the dates referenced in this paper.

¹⁸ Note that this is not analogous to the frequent batch auction system described by Budish et al. (2015), which eliminates sniping of stale quotes, but is more like the random order delay system of Harris (2012).

¹⁹ The base currency is the first currency displayed in the symbol of the currency pair. For example, the euro is the base currency of EUR/CHF.

(e.g., Mancini et al. 2013), we focus on data between 8:00 Greenwich Mean Time (GMT) and 17:00 GMT and exclude weekends, when the EBS market’s trading platform is effectively inactive.^{20,21}

The transaction data set records the time stamp to the millisecond of each trade that occurred, along with the transaction price and quantity. Most importantly, the parties providing liquidity (submitting the limit order) and consuming liquidity (submitting the market order²²) have identifiers as ‘human’, ‘bank AI’ or ‘PTC AI’. Thus, each trade can be classified by nine possible combinations of liquidity provider and consumer. In addition, each trade has an indicator of whether the liquidity provider was the buyer or the seller. We line up these millisecond time-stamped transactions with the 100-millisecond quote data, such that for a given transaction, the top 10 anonymised best bid and ask prices at the nearest previous whole 100-millisecond interval are also available. All quotes are firm and, therefore, truly represent the market prices at that instant.²³ Here, the trader type that posts each quote is not available to us, whereas that information is available in the transactions data.

4 Overview of the Swiss franc event

4.1 Policy context

The SNB began intervening in the FX market to cap the value of the Swiss franc against the euro on 6 September 2011. In a press release of the same date, the SNB said, ‘the massive overvaluation of the Swiss franc poses an acute threat to the Swiss economy and carries the risk of a deflationary

²⁰ See Chaboud et al. (2014) for further discussion of trading activity on the EBS market system. In addition, EBS indicated to us that its dealing services are ordinarily open for trading 24 hours a day and 7 days a week, with the exception of a maintenance window that ordinarily occurs from 17:50 EST on Friday until Saturday morning. For the purpose of computing market data products, such as ‘highs’ and ‘lows’, EBS regards trades between 17:00 Friday Eastern Standard Time (EST) and 5:00 Monday Sydney time as not conducted in normal market conditions or market hours. We therefore exclude these trades from the calculations. However, the trades remain in the EBS data.

²¹ In our sample, we drop 5 transactions and 24 quotes on EUR/CHF that took place between 09:32:29 GMT and 09:32:39 GMT on 15 January, where the price was exceptionally low, at 0.0015, with a volume of one million of the base currency for each transaction. EBS confirmed that those transactions were errors made by the traders, and the counterparties settled outside EBS.

²² We refer to ‘market orders’ as those termed ‘immediate or cancel’ (‘IOC’) orders by EBS. Though, unlike conventional market orders, market orders in EBS come with limit prices, beyond which they would not be executed immediately, but instead cancelled.

²³ The historical market data provided by EBS are time-sliced at 100 milliseconds and are therefore a snapshot of the activity during the time period.

development'. Therefore, 'it will no longer tolerate a EUR/CHF exchange rate below the minimum rate of CHF 1.20 . . . and is prepared to buy foreign currency in unlimited quantities' (SNB 2011).

As observable in Figure 2, following the introduction of this policy, the franc generally traded a little below its cap until late 2014, before appreciation pressures intensified towards the end of 2014 due to weakness in the euro-area economy and the safe haven status of the Swiss franc (Rinaldo and Söderlind 2010). In response, the SNB cut its interest rate on sight deposits to minus 0.25%. However, commitment to the exchange rate policy appeared to remain firm. On 18 December 2014, the SNB governor stated that the central bank was 'committed to purchasing unlimited quantities of foreign currency to enforce the minimum exchange rate with the utmost determination' (Jordan 2014).²⁴

Despite that, the capping policy was discontinued at 9:30 GMT²⁵ on 15 January 2015. Releasing an unscheduled press communiqué, the SNB explained the following:

Recently, divergences between monetary policies of the major currency areas have increased significantly. The euro has depreciated significantly against the US dollar and this, in turn, has caused the Swiss franc to weaken against the US dollar. In these circumstances, the SNB concluded that enforcing and maintaining the minimum exchange rate for the Swiss franc against the euro is no longer justified. (2015)

That this news was not generally expected by market participants is reflected in FX options prices leading up to the announcement (e.g., Mirkov et al. 2016; Jermann 2017),²⁶ the financial reporting after the announcement (e.g., Reuters 2015b; Bloomberg 2015) and the delayed reaction to the announcement described below.

²⁴ Similarly, on 12 January 2015, another member of the SNB Governing Board said, 'we are convinced that the minimum exchange rate must remain the cornerstone of our monetary policy' (Reuters, 2015a).

²⁵ In the rest of the paper, all times refer to GMT.

²⁶ Both Mirkov et al. (2016) and Jermann (2017) find low implied break probabilities, though Hertrich and Zimmermann (2017) suggest a larger break probability, approaching 50%, was priced in.

4.2 Trading patterns

Here, we provide an overview of the market reaction to the SNB announcement made on 15 January 2015. In particular, we graphically illustrate the market patterns and key features of algorithmic and human trading around the announcement as a prelude to the more formal analysis in the rest of the paper. We first focus on the announcement time and then on the main intraday periods of the event day.

Panel A of Figure 3 plots the limit order volume to buy ('bid') or sell ('ask') euros for Swiss francs and the mid-price,²⁷ starting a few minutes before the announcement until the end of the event day. This figure shows that the bid volume before the announcement was constantly very large—that is, more than 450 million euros. Although our data from EBS do not include the identities of the institutions submitting the orders, we presume that this wall of bid orders was erected by a very special human trader, the SNB, to enforce the cap. This is based on the following: (i) the regular presence of the SNB before the announcement, (ii) the large size of the order (which is excessive for many market participants), (iii) the limited prices, which were in line with the SNB's cap and (iv) further issues such as the timing, which we discuss later.

Not by coincidence, the order book data suggest that the wall of bids disappeared at 9:29:45 (see Figure 3), reinforcing the idea that the SNB withdrew its limit orders in preparation for the actual announcement of the cap removal, which occurred at 9:30:00. Being an unscheduled announcement, this media release caught market participants by surprise. This may explain why the exchange rate remained essentially unchanged until 9:30:44, at which point it broke through the cap. We take this delayed response as further evidence that the decision to remove the cap was not expected by market participants.²⁸ Once it had broken through the cap, it moved rapidly, and both sides of the order book shrank significantly (indeed, on a few occasions during 9:31, the bid side of the market disappeared completely for a few seconds).

²⁷ The mid-price is the mid-point between the best bid and ask prices.

²⁸ We thank Alain Chaboud for pointing out to us this 'Wile E. Coyote' moment, in which a significant portion of the orders underpinning the value of the euro against the franc were withdrawn, but it was not until about 1 minute later that the price started to fall.

As shown in Figure 3, the post-announcement period on the event day can be divided into three main sub-periods. The first phase spans from the SNB announcement until the re-erection of the bid wall a few seconds after 9:53. This represents the genuine shock phase during which the SNB removed the cap and, as far as we can judge, refrained from intervening in any way.²⁹ The second post-announcement phase starts with the re-appearance of the wall of bid orders of about 4 billion euros after 9:53. Again, we assume this reflects the return of the SNB, as it seems unlikely that other market participants, whether human or algorithmic traders, would have placed such a large volume of orders. However, it is notable that although the scale of bids was very large, the price of those bids was changed regularly so that very few were actually matched against market orders. Because market participants could see the size of the order book, it is very likely that they inferred that the SNB had stepped in, and so revised their beliefs about its willingness to act as ‘the buyer (of euros) of last resort’ and market stabiliser. Similar to Panel A, Panel B of Figure 3 shows the USD/CHF pattern, which has an analogous large volume on the bid side of more than 1 billion US dollars that appeared at about the same time. This phase went on for about one hour, until 10:52. After this point in time, the wall disappeared for a few minutes. Although small, walls intermittently reappeared, suggesting an attempt of the SNB to step back somewhat and let price discovery be conducted by market participants on their own. We consider the sub-period after 10:52 on the event day as the third phase, during which AT and human market participants trade more against each other, perhaps with the belief that the SNB was standing by to intervene when necessary.

Figure 4 focuses on the first phase of the shock, characterised by a sharp appreciation of the Swiss franc. It shows the prices at which different trader types exchanged euros for Swiss francs in the 23 minutes following the SNB announcement, depending on whether their trades were consuming liquidity (market orders, Panel A) or providing it (limit orders, Panel B). The main message is that agency (bank) AT massively consumed and, to a lesser extent, provided liquidity at extreme prices (prices significantly different to those of the immediately preceding trades) on a number of occasions, notably between 9:31 and 9:36. Thus, over 79% of the cumulative appreciation of the franc in the 23 minutes to 9:53 is

²⁹ Both our analysis of trading patterns and discussions with market participants suggest this period contained no intervention.

attributable to agency AT, which accounted for 55% (40%) of the volume of liquidity-consuming (liquidity-providing) trades. That agency (bank) AT both consumed and provided liquidity at extreme prices may reflect the diverse set of traders, including not only different banks but also various clients from whom these trades originated. Human traders were also very active, as they accommodated many extreme-price trades and guaranteed the largest share of the volume of liquidity-providing trades (51%), which was much larger than that of liquidity-consuming trades (21%).

Figure 5 delivers a consistent picture—that is, that computer traders were net purchasers of Swiss francs throughout the day, in particular bank AIs against the euro (Panel A.1) and PTC AIs against the US dollar (Panel B.1), whereas human traders were net purchasers of the base currencies (euros and US dollars). Regarding liquidity provision, Panels B.1 and B.2 of Figure 5 show that human traders were consistently net suppliers of liquidity throughout the day, whereas PTC AI trades consumed it.

Overall, our graphical overview suggests that after the SNB announcement, agency and proprietary AT contributed to the liquidity dry-up (as they were net consumers of liquidity) and the subsequent price disruption (as they were net purchasers of the appreciating currency). Thus, the preliminary results suggest that computers traded ‘with the wind’, buying the franc as it appreciated, whereas humans ‘leaned against the wind’. In the following sections, we examine these issues in more detail.

5 Market liquidity

In this section, we investigate in more detail the contributions of computer and human traders to the market liquidity of EUR/CHF and USD/CHF before, during and after the day of the SNB announcement. We first study the quantity-based (volume) dimension and then the price-based (effective spreads) dimension of market liquidity.

5.1 Liquidity volumes provided and consumed

First, we identify whether the consumer and the provider of liquidity for each trade was a human, an agency (bank) AT or a proprietary (non-bank) PTC AT. We then record the shares³⁰ of total trading volumes in three different periods: a ‘pre-event’ period (5–14 January 2015), the ‘event day’ (15 January 2015) and a ‘post-event’ period (16–23 January 2015). Finally, we record the ‘net liquidity provision’, which is the difference between the share of trades for which a given trader type provided liquidity and the share for which it consumed liquidity. Table 1 reports the main results.

Compared with the pre-event period, the net liquidity provision by computers fell on the event day and in the post-event period, whereas it increased for humans during both periods. The increases in human net liquidity provision were all statistically significant, whereas the decreases for computers were significant for at least one type of AT. In EUR/CHF (Panel A), the net liquidity provision by agency AT significantly declined on the event day, whereas that by proprietary AT significantly declined in the post-event period. In USD/CHF (Panel B), the net liquidity provision by both types of AT significantly declined on the event day, while it was again that of proprietary AT that significantly declined in the post-event period. Economically, the most significant changes were the decline in the share of agency AT liquidity-providing trades on the event day in EUR/CHF and the offsetting increase in the human share. Thus, although agency AT also decreased its share of liquidity-consuming trades on the event day and humans increased their share, this was not enough to change the pattern of net liquidity provision.

5.2 Effective spreads

We now turn our focus to a price-based indicator of liquidity. As shown in Section 5.1, computers reduced their net provision of liquidity on the event day and afterwards, but did they widen the bid–ask

³⁰ Details of the total value and number of trades are shown in the Online Appendix. Overall, all types of traders increased their trading activity, and there was about eight times more CHF trading on the event day than there was on any day in the pre-event period.

spreads on trades for which liquidity was still available? To investigate this question, we calculate a series of effective spreads, s , for each of the trader types:

$$S_{tk} = \frac{q_{tk}(p_{tk} - m_t)}{m_t},$$

where t indexes the time of the trade; k indexes the type of trader providing liquidity (human, bank AI or PTC AI); q is a binary variable equal to +1 for trades in which the liquidity consumer was buying the base currency and -1 for trades in which it was selling it; p is the transaction price; and m is the mid-point between the best bid and ask quotes from any type of trader in the 100-millisecond window in which the trade took place.

To mitigate the effects of some extreme observations, Table 2 shows median effective spreads for EUR/CHF (Panel A) and USD/CHF (Panel B). For the same reason, the tests of equality of the medians reported in the table are non-parametric.³¹ For EUR/CHF, median effective spreads were very similar across trader types in the pre-event period. Although spreads exploded for all trader types on the event day, spreads associated with AT increased much more, and they only contracted a little moving into the post-event period. The results are similar for USD/CHF. In this case, humans offered narrower spreads in the pre-event period. On the event day, the largest spread increase came from proprietary AT, and it remained much wider in the post-event period.

Taking our quantity-based (order flows) and price-based (effective spreads) indicators of market liquidity together, our findings corroborate the alternative hypothesis—that is, that there is a relative reduction in the activity of AT in extreme situations. More specifically, our results show that both agency and proprietary algorithms significantly reduced their net liquidity provision, and their trading activity created much wider spreads. In contrast, human traders significantly increased their net provision of liquidity and did so at narrower spreads than did computers.

³¹ We use K-sample tests to investigate equality across periods and Snedecor and Cochran (1989) tests to investigate equality across trader types.

6 Price efficiency

A second important dimension of market quality, alongside liquidity, is pricing efficiency. In an efficient market, any gaps that might arise between the actual price of an asset and the ‘efficient price’ reflecting its fundamental value tend to be small. They are closed quickly by traders drawing on the available information about the fundamental value.³² Although price efficiency is difficult to measure, it is fair to assume that an efficient market environment is characterised by the following properties: (i) little excess volatility in prices on top of that attributable to changes in the fundamental price, (ii) an effective price discovery process and (iii) no or few deviations from arbitrage conditions. In this section, we examine the contributions of different trader types to these dimensions of efficient pricing around the SNB announcement.

6.1 Contributions to realised volatility

Following O’Hara and Ye (2011), we calculate the contributions of different trader types to the realised variance of returns in our pre-event, event day and post-event periods as follows:

$$V_k = \frac{\sum_{n=1}^N (r_n d_{nk})^2}{\sum_{n=1}^N r_n^2},$$

where k indexes the three types of traders (human, agency [bank] AT and proprietary [non-bank] PTC AT); n indexes the return observations, of which there are N in total; r denotes the returns, which are logarithmic returns derived from successive transaction prices; and d is a dummy variable that equals one if the trader initiating the trade (i.e., consuming liquidity) is of a particular trader type and is zero otherwise. So, for example, if prices changed by 2% during a particular period as a result of two trades, one by a computer that moved the price by 1% and another by a human trader that moved the price by a further 1%, then each type of trader would have contributed 50% of the realised variance in that period.

³² As Hendershott and Menkveld (2014) note, the net provision of limit orders need not be a good measure of liquidity supply, because a market order that leans against price pressure (i.e., goes against the prevailing market trend) can be thought of as contributing to liquidity and reducing volatility.

Panel A.1 of Table 3 shows the results of this breakdown applied to EUR/CHF and USD/CHF during the pre-event, event day and post-event periods. The results for EUR/CHF show a remarkable increase in the variance contribution of agency AT on the event day, of more than 90% of the total variance share. The post-event shares returned to the pre-event ones. The results for USD/CHF show a similar pattern.

Of course, the breakdown in Panel A.1 of Table 3 is effectively a combination of the share of total trading of each type of trader and the per-trade impact on volatility of each type of trade. Changes in variance contributions could therefore simply reflect changes in trading shares. As such, we calculate a per-trade variance impact coefficient, which simply scales variance contributions by the number of trades undertaken by the different trader types. Panel A.2 of Table 3 shows consistent results—that is, a high contribution of agency AT to volatility on the event day and a low contribution of humans and (to a lesser extent) of proprietary AT to volatility.³³

As price jumps can affect realised volatility, we undertake the same analysis for contributions to realised bipower variation (Barndorff-Nielsen and Shephard 2004), which strips out the impact of price jumps. Realised bipower variation is defined as follows (for our case, where we are interested in contributions):

$$BV_k = \sum_{n=2}^N |r_n r_{n-1}| d_{nk} / \sum_{n=2}^N |r_n r_{n-1}|,$$

where r and N are defined in the equation for V_k (above). Panels B.1 and B.2 of Table 3 show the same qualitative results as those using realised variances, suggesting that our findings are not sensitive to price jumps.

Overall, the analysis of realised volatility shows that agency AT was the main driver of volatility in EUR/CHF on the event day, providing empirical support relevant to policy concerns highlighted in other extreme FX market events (e.g., BIS 2017).

³³ We obtain the same qualitative results if we normalise by trade volume rather than by number of trades.

6.2 Price discovery

As the second dimension of price efficiency, we now analyse the price discovery process. As any increase in volatility is valuable if it helps market prices to fully reflect available information, it is important to view the results in Section 6.1 in combination with the analysis of the price discovery process.

As a first step, we perform a simple univariate price-impact regression of the order flow disaggregated by market participants in the spirit of Kyle (1985). More specifically, we regress 5-minute returns on the order flow for each type of trader during the same 5-minute periods. Returns are computed as logarithmic returns (on the base currency) between successive mid-points of the best bid and ask quotes in the final 100 milliseconds of each period. Order flow is computed as the difference between liquidity-consuming purchases of the base currency and liquidity-consuming sales of the base currency by each trader type.

Table 4 reports the results for EUR/CHF (Panel A) and USD/CHF (Panel B). Consistent with the realised variance analysis (see Panel A.2 of Table 3), the order flow originated by agency AT had the largest price impacts on both EUR/CHF and USD/CHF in the post-event period, as well as on USD/CHF on the event day itself. However, it was effectively zero for EUR/CHF on the event day, when agency AT trades generated the most volatility.

We extend the price-discovery analysis using a variant of the vector autoregression (VAR) model developed by Hasbrouck (1991a, 1991b, 2007) and employed by Hendershott et al. (2011).³⁴ This approach is based on estimating the contribution of news and different types of trade to the underlying efficient price. Specifically, we estimate the following model:

$$r_t = \sum_{k=1}^K \alpha_k r_{t-k} + \sum_{k=0}^K \beta_{k,P} x_{t-k,P} + \sum_{k=0}^K \beta_{k,B} x_{t-k,B} + \sum_{k=0}^K \beta_{k,H} x_{t-k,H} + \varepsilon_{r,t},$$

$$x_{t,P} = \sum_{k=1}^K \gamma_k r_{t-k} + \sum_{k=1}^K \delta_{k,P} x_{t-k,P} + \sum_{k=0}^K \delta_{k,B} x_{t-k,B} + \sum_{k=0}^K \delta_{k,H} x_{t-k,H} + \varepsilon_{P,t},$$

³⁴ The model is described in detail in Hasbrouck (2007, pp. 78–85).

$$x_{t,B} = \sum_{k=1}^K \zeta_k r_{t-k} + \sum_{k=1}^K \eta_{k,P} x_{t-k,P} + \sum_{k=1}^K \eta_{k,B} x_{t-k,B} + \sum_{k=0}^K \eta_{k,H} x_{t-k,H} + \varepsilon_{B,t},$$

$$x_{t,H} = \sum_{k=1}^K \lambda_k r_{t-k} + \sum_{k=1}^K \nu_{k,P} x_{t-k,P} + \sum_{k=1}^K \nu_{k,B} x_{t-k,B} + \sum_{k=1}^K \nu_{k,H} x_{t-k,H} + \varepsilon_{H,t},$$

where, as before, r denotes returns on the base currency and x denotes disaggregated order flows for human, agency (bank) AT and proprietary (non-bank) PTC AT calculated over 5-minute periods, indexed by t ,³⁵ and ε_r , ε_P , ε_B and ε_H are the error terms with variances of $\sigma_{\varepsilon,r}^2$, $\sigma_{\varepsilon,P}^2$, $\sigma_{\varepsilon,B}^2$ and $\sigma_{\varepsilon,H}^2$ (i.e., the variances of return shocks and shocks to order flows of each trader type), respectively.

The only structural assumption in this model relates to timings: we assume that proprietary AT—as the fastest trader type in the market—can adjust its order flows to the contemporary order flows of other market participants. Similarly, agency AT—as the next fastest trader type—can adjust its order flows to the contemporary order flows of human traders, but not to those of proprietary AT. Finally, human traders can only adjust their order flows to the previous order flows of other market participants.³⁶ These assumptions are in a similar vein to those of Brogaard et al. (2014) in their study of HFT and non-HFT trading activities.

We estimate this model, selecting $K = 5$ as the optimal number of lags based on AIC criterion, and transform it to a vector moving-average representation by repeatedly substituting for the right-hand-side terms.³⁷ The resulting equation for returns is as follows:

$$r_t = (\varepsilon_{r,t} + \sum_{k=1}^{\infty} a_k \varepsilon_{r,t-k}) + \sum_{k=0}^{\infty} b_{k,P} \varepsilon_{P,t-k} + \sum_{k=0}^{\infty} b_{k,B} \varepsilon_{B,t-k} + \sum_{k=0}^{\infty} b_{k,H} \varepsilon_{H,t-k}.$$

As suggested by Hendershott et al. (2011), the last three terms on the right-hand side may be considered ‘private information’, because they reflect order flows from particular trader types, whereas the first term may be considered ‘public information’. Thus, we can identify separate contributions to efficient

³⁵ A Dicky–Fuller unit root test on the time series of the regression variables confirms the stationarity of the time series in each sub-period, including covariance stationarity.

³⁶ This is the most logical ordering, but the pattern of results across periods remains similar, even with alternative orderings.

³⁷ We do this with ten lags, after which the marginal effect of each substitution becomes small.

pricing from public information and private information pertaining to each of the three types of traders via the following:³⁸

$$\sigma_{\text{w}}^2 = (1 + \sum_{k=1}^{\infty} a_k)^2 \sigma_{\varepsilon,r}^2 + (\sum_{k=0}^{\infty} b_{P,k})^2 \sigma_{\varepsilon,P}^2 + (\sum_{k=0}^{\infty} b_{B,k})^2 \sigma_{\varepsilon,B}^2 + (\sum_{k=0}^{\infty} b_{H,k})^2 \sigma_{\varepsilon,H}^2,$$

where σ_{w}^2 is the variance of the random-walk component, and the terms on the right-hand side represent contributions to this from public information and private information pertaining to PTC AI, bank AI and human traders, respectively.

Table 5 presents the results of this decomposition for EUR/CHF (Panel A) and USD/CHF (Panel B). The results for EUR/CHF show a striking shift in information contributions across the three periods. In the pre-event period, proprietary AT made by far the largest contribution to the efficient price of all types of order flows, with agency AT and human traders contributing very little. On the event day, human traders took over as the most significant contributor, whereas the influence of proprietary AT all but disappeared. Human trading also maintained a significant contribution in the post-event period. The role of agency AT on the event day and in the post-event period was similar to that of human traders, though not as dramatic, stepping up from the pre-event levels. This was dwarfed by the increased contribution of agency AT to the total realised volatility, as highlighted in the previous section.

The pattern for USD/CHF is somewhat less clear, as the informational role of all types of trading was relatively small, pointing to the leading role of the EUR/CHF in discovering the new value of the Swiss franc.

6.3 Arbitrage deviations

Efficient prices should obey arbitrage conditions, such as triangular arbitrage. More specifically, this means that direct quotes for EUR/CHF should not move outside the range defined by implied quotes

³⁸ To be clear, this follows from the assumptions of the model, and not from information in the market data provided by EBS.

derived from USD/CHF and EUR/USD. It is fair to assume that some computer trading in these markets is dedicated to identifying and trading on this triangular arbitrage.

We begin by calculating the frequency and average size of such arbitrage opportunities in the pre-event, event day and post-event periods. Specifically, we examine the best bid and ask quotes in each 100-millisecond window and record the existence of an arbitrage opportunity if a profit could have been made by buying EUR/CHF directly and selling ‘synthesised’ EUR/CHF via USD/CHF and EUR/USD trades. The profits must exceed a minimum of one basis point, and we record the average profitability of all arbitrage opportunities meeting this criterion.

Panel A of Table 6 shows that, by far, the largest and most frequent arbitrage opportunities occurred during the event day itself. Arbitrage opportunities then remained over ten times more frequent in the post-event period compared with the pre-event period. This may suggest that AT had become less active in closing arbitrage deviations during the latter period, possibly in response to the increased volatility or simply a withdrawal from trading following the event. To shed more light on this issue, we follow Chaboud et al. (2014) and estimate a structural vector autoregression (SVAR) model of the relationship between arbitrage opportunities and the trading volumes of the three types of trader (human, agency [bank] AT and proprietary [non-bank] PTC AT):

$$AY_t = \alpha(L)Y_t + \beta X_t + \delta G_t + \epsilon_t,$$

where Y contains four endogenous variables, which are the frequency of arbitrage opportunities and order flows of each trader type relative to the total market order flows. These variables are measured over 5-minute windows. A is a 4×4 matrix of coefficients governing contemporaneous relationships between the endogenous variables. These were estimated using the approach of Rigobon (2003) and two lags of the endogenous variables. X includes six exogenous variables all computed over the preceding 10 minutes: the total trade volumes and return volatilities for each of the three currency pairs in the arbitrage triangle. Finally, G includes nine time-dummy variables, one for each hour of the trading day.

Panel B of Table 6 shows the estimated contemporaneous coefficients that reflect how the trading activity of each type of trader responds to arbitrage opportunities. The coefficients reveal that in the pre-event period, both AT types reacted in the expected way (i.e., to close arbitrage opportunities). In the post-event period, however, no type of trader responded significantly to arbitrage opportunities. Thus, our results suggest a significant reduction in computer resources devoted to triangular arbitrage following the Swiss franc event. Also, the fact that no type of trading made a significant contribution to arbitrage suggests that most mispricings were closed by quote adjustment rather than by active trading.

To sum up, our analysis substantiates the idea that in reaction to extreme events, the contribution of AT to support price efficiency decreases; however, the opposite applies to human traders. This may reflect the possibility that after the announcement, many computer trades, especially agency AT, were driven by liquidity needs and risk exposure reductions rather than by information motives. Such patterns would be consistent with our previous findings (i.e., that computers reduce the liquidity supply and are the net consumers of liquidity). Alternatively, central bank interventions might have affected AT and human trading differently. We analyse the role of the central bank in the next section.

7 SNB intervention

As shown in Section 4.2, the analysis of the limit order book provides hints about the possible central bank interventions. Whether the central bank's presence was perceivable or not is an important issue, because algorithmic and human traders can react differently to it. For instance, in the presence of interventions of a (credible) central bank, human traders may revise their expectations about price reversals, price disruptions and market stability. Being pre-programmed, in contrast, it might be more difficult for algorithmic traders to decipher these issues, especially in unusual market environments. To address this issue, we perform two analyses, starting from a broad perspective and then zooming in on EBS data.

7.1 SNB sight deposits and option hedging

In our first analysis, we use publicly available data to estimate the overall intervention volume of the central bank. Using weekly data on SNB sight deposits (which is the key liability created by FX interventions), we estimate that the total SNB intervention was probably less than 16% of Swiss franc turnover on EBS over our full sample (18% of Swiss franc turnover during the 7 days beginning 15 January).³⁹ As discussed below, we are fairly confident that we can identify SNB limit order intervention on EBS through the techniques used (e.g., the use of a ‘wall’ of orders and some other, more detailed, characteristics). Based on these figures, we estimate that about a quarter of SNB intervention in the event week was through EBS limit orders.⁴⁰ This estimate is consistent with other evidence that suggests SNB favours a diversified approach (Moser 2016), including the substantial use of telephone orders (Fischer 2004), and it is also consistent with anecdotal evidence from market participants.

Our data set does not allow us to trace the impact of the SNB’s other forms of intervention (such as telephone orders) on EBS, but some market participants have highlighted the role of telephone intervention in helping dealers hedge their substantial FX options’ exposure in EUR/CHF. In order to corroborate this hypothesis, we analysed dealers’ option positions using the EMIR trade repository database on option positions.⁴¹

The main idea is that dealer banks supplied protection against Swiss franc appreciation to other market participants (such as Swiss corporates). As a result, dealers entered the event day with a significantly unbalanced (long) exposure in EUR/CHF options, triggering a large delta-hedging requirement after the SNB announcement—an operation commonly undertaken through an agency AT. Focussing on the

³⁹ These weekly estimates are based on the assumption of an even spread of intervention over the week. However, we estimate that the maximum intervention that could have occurred on the event day would be about 29% of EBS turnover. On this basis, our estimates suggest that about 25% of that was direct limit order intervention on EBS (i.e., about 7% of EBS turnover), with the rest being made up of market orders on EBS and trading on other venues (such as a telephone trades or trades on other platforms). Anecdotal evidence suggests the majority would have been telephone trades.

⁴⁰ Any SNB market orders on EBS would have increased liquidity consumption by humans, which is inconsistent with the previous findings that human traders are (net) liquidity suppliers in reaction to the SNB announcement.

⁴¹ This data comes from the regulatory requirement that all OTC and exchange-traded derivatives transactions undertaken by European Union (EU) counterparties since August 2012 should be reported to a trade repository. Relevant parts of this data set are then made available to national regulators (see, for example, Cielinska et al. (2017) for further details).

10 largest dealer banks (that normally represent around 90% of the dealer market, and so, presumably, the same share of bank AI), the trade repository data confirms that, as of 14 January 2015, they had a strongly unbalanced position. These banks had a net long EUR/CHF option position of 84 billion euros notional, and the ratio of long to short trades was about 1.8 to 1 in notional value terms. Using the data on strike prices, implied volatility and maturity available in this database, we calculate the standard Gamma measures. We find that this net position could have generated a hedging demand of about 4.5 billion euros per 1% move in EUR/CHF; the equivalent figure for USD/CHF was less than half this amount at, 2 billion US dollars. Overall, this evidence supports the views of participants that dealers had substantial option hedging requirements on the event day. If they are correct that much SNB telephone intervention was effectively passed directly to option desks, it is likely that this form of intervention actually reduced agency AT liquidity consumption, because option hedging would usually be carried out through this route.

7.2 Event-day results before SNB intervention

In our second analysis, we zoom in on the first intraday sub-period of 23 minutes after the SNB announcement, when it is widely accepted that there was no central bank intervention. This allows us to examine the reaction of AT and human traders to the shock in the absence of central bank intervention. We re-calculate both liquidity provision and realised variance contribution during this phase for EUR/CHF and USD/CHF. The results (see Tables OA.3 and OA.4 in the Online Appendix) are strikingly similar to those of the full event day. For example, for EUR/CHF, humans were significant net liquidity providers (30.5% net liquidity provision) and agency AT was the dominant source (92%) of realised variance. The main difference is that, over this intraday period, the group representing agency AT was the largest net liquidity consumer (-14.6%), a result that is consistent with SNB telephone intervention becoming a source of liquidity to dealer banks, offsetting the need for agency AT later in the day.

7.3 Identifying intervention in the EBS order book

To conclude our analysis, we estimate the limit order trades performed by the central bank, then recalculate the liquidity provision and consumption by the three types of trader by excluding these estimated SNB interventions. As noted above, the data set we have is anonymous, so our identification strategy is based on the SNB intervention approach described in Section 4.2 rather than on direct evidence of the trader's identity. More precisely, we consider a trade to be a SNB intervention when the following three conditions apply: First, the outstanding volume on the bid side of the EBS EUR/CHF order book is 20 million euro or more,⁴² consistent with the analysis provided in Section 4.2. Second, a human trader is on the passive side (maker). Third, 00 or 50 are the last two digits of the transaction price, consistent with past SNB intervention strategy (Fischer 2004).

Given this identification of SNB limit orders, we repeat the analysis underpinning Table 1 excluding these trades. Table 7 shows the adjusted results for liquidity provision and consumption by computers and humans for EUR/CHF. Overall, the main result remains materially unaffected—that is, a reduction in (net) liquidity provision by computers is observed on the event day and in the post-event period.

8 Conclusion

The Swiss franc event is one of the largest shocks to financial markets in recent years and the most significant 'black swan' of the FX market since AT became prominent. Using FX data with a precise identification of AT, we study the reaction to this shock of human traders and AT, which can be further divided into two broad categories: agency (bank) AT and proprietary (non-bank) AT. We also study the role of central bank intervention and how it can affect human traders and AT.

We find that computer trading contributed to the decline in the market quality of the currency pairs directly affected on the event day and afterwards. Agency algorithms, which typically offered

⁴² The daily average bid size during the pre-event period and the event day is 14 million euros.

EUR/CHF liquidity before the shock, reduced their liquidity supply and created uninformative volatility. After the shock, proprietary algorithms subdued the price discovery and (triangular) arbitrage. In contrast, human traders took over as the main contributors to liquidity provision and efficient pricing.

Of course, it is hard to draw general conclusions from one event, not least because not all market shocks are the same. Indeed, an important aspect of this particular event was the role of central bank intervention. Although we present some evidence that the SNB intervention does not alter our main results on computer trading, the overall effects of central bank interventions on human trading are less clear. The intervention may have encouraged human trading, as human traders possibly anticipated or adapted to it. Certainly, the fact that the SNB intervention tactics seem to have been deliberately designed to make their presence highly visible to market participants suggests that the central bank felt that communicating its involvement to the market was an important channel of intervention impact. As well as highlighting this potential co-ordination role of intervention for human trading, our results call attention to other issues relevant to central banks, such as the importance of monitoring outstanding FX option positions to help assess the likely impact of policy announcements and intervention.

In this paper, we study *how* AT reacts to a ‘black swan’ event. Future research should shed light on *why* it does so. For instance, is the AT reduction to market quality due to more stringent capital, trading requirements applied to AT or an inability to discern non-standard dynamics including central bank interventions? Our work should provide several insights for policymakers, market participants and academics. First, the main consensus in the literature, that AT improves liquidity and price efficiency, may be true in normal times, but it is debateable in extreme circumstances. Second, by comparing algorithmic trades across trading types (agency vs proprietary AT) and currency pairs, our paper supports the idea that AT cannot be viewed as a single object. It includes a diverse set of trading strategies, many of which do not fit the stereotype of computer trading (i.e., HFT arbitrage).⁴³ Finally, possible future tasks for policymakers and market regulators may include monitoring computer trading, understanding how algorithms function especially in distressed situations and the conduct of stress tests.

⁴³ This conclusion is similar to that of Boehmer et al. (2018b).

References

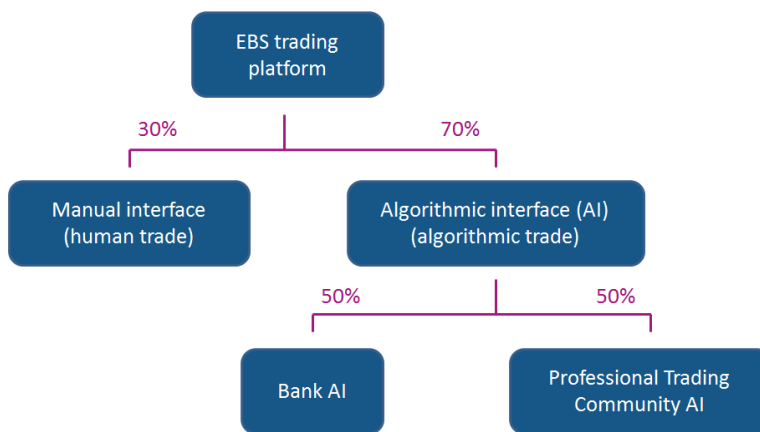
- Afonso G, Kovner A, Schoar A (2011) Stressed, not frozen: The federal funds market in the financial crisis. *J. Finance*. 66:1109–1139.
- Aitken M, Cumming D, Zhan F (2014) Trade size, high-frequency trading, and colocation around the world. *European J. Finance*. 1466–4364 (online).
- Ait-Sahalia Y, Saglam M (2014) High-frequency traders: Taking advantage of speed. Working paper.
- Anand A, Venkataraman K (2014) Should exchanges impose market maker obligations? Working paper.
- Barndorff-Nielsen OE, Shephard N (2004) Power and bipower variation with stochastic volatility and jumps. *J. Financial Econom*. 2:1–37.
- Bernales A (2014) How fast can you trade? High-frequency trading in dynamic limit order markets. Working paper, Banque de France, Paris, France
- Biais Bm Woolley P (2012) High-frequency trading. Manuscript, Toulouse University, Toulouse, France.
- Biais B, Foucault T (2014) HFT and market quality. *Bankers, Markets and Investors*. 128:5–19.
- Biais B, Foucault T, Moinas S (2015) Equilibrium fast trading. *J. Financial Econ*. 116:292–313.
- BIS (2011) High-frequency trading in the foreign exchange market. Report, Study Group established by the Market Committee, Bank for International Settlements, Basel, Switzerland, <https://www.bis.org/publ/mktc05.pdf>.
- BIS (2016) Triennial Central Bank Survey: OTC derivatives positions at end-June 2016. Bank for International Settlements, Basel, Switzerland.
- BIS (2017) The sterling ‘flash event’ of 7 October 2016. Report, Bank for International Settlements Market Committee, Basel, Switzerland.
- Bloomberg (2015) SNB unexpectedly gives up cap on Franc. (January 15).
- Boehmer E, Fong K, Wu J (2018a) Algorithmic trading and market quality: International evidence. Working paper.
- Boehmer E, Li D, Saar G (2018b) The competitive landscape of high-frequency trading firms. *Rev. Financial Stud*. 31:2227–2276.
- Bongaerts D, Van Achter M (2016) High-frequency trading and market stability. Working paper.
- Breckenfelder J (2013) Competition between high-frequency traders and market quality. Working paper.
- Brogaard J, Carrion A, Moyaert T, Riordan R, Shkilko A, Sokolov K (2018) High-frequency trading and extreme price movements. *J. Financial Econ*. 128:253–265.

- Brogaard J, Hagstromer B, Norden L, Riordan R (2015) Trading fast and slow: Co-location and market quality. *Rev. Financial Stud.* 28:3407–3443.
- Brogaard J, Hendershott T, Riordan R (2014) High-frequency trading and price discovery. *Rev. Financial Stud.* 27:2267–2306.
- Brunnermeier MK, Pedersen LH (2009) Market liquidity and fund liquidity. *Rev. Financial Stud.* 22:2201–2238.
- Brunnermeier MK, Nagel, S, Pedersen, LH (2009) Carry trades and currency crashes. *NBER Macroecon. Annu.* 23:313–47.
- Budish E, Cramton P, Shim J (2015) The high-frequency trading arms race: Frequent batch auctions as a market design response. *Q. J. Econ.* 130:1547–1621.
- Carrion A (2013) Very fast money: High-frequency trading on the NASDAQ. *J. Financ. Mark.* 16:680–711.
- Cespa G, Vives X (2015) The beauty contest and short-term trading. *J. Finance.* 70:2099–2154.
- CFTC-SEC Staff (2010) Findings regarding the market events of May 6, 2010. Report, U.S. Commodity Futures Trading Commission, Washington, DC.
- Chaboud A, Chiquoine B, Hjalmarsson E, Vega C (2014) Rise of the machines: Algorithmic trading in the foreign exchange market. *J. Finance.* 69:2045–2084.
- Chordia T, Goyal A, Lehmann B, Saar G (2013) High-frequency trading. *J. Financ. Mark.* 16:637–645.
- Cielinska O, Joseph A, Shreyas U, Tanner J, Vasios M (2017) Gauging market dynamics using trade repository data: The case of the Swiss franc de-pegging. Report, Bank of England Financial Stability Paper No. 41.
- Conrad J, Wahal S, Xiang J (2015) High-frequency quoting, trading, and the efficiency of prices. *J. Financial Econ.* 116:271–291.
- Du S, Zhu H (2015) Welfare and optimal trading frequency in dynamic double auctions. Working paper, National Bureau of Economic Research, no. 20588.
- Duffie D (2012) *Dark Markets: Asset Pricing and Information Transmission in Over-the-Counter Markets* (Princeton University Press, Princeton, NJ).
- Easley D, Lopez de Prado M, O’Hara M (2012) Flow toxicity and liquidity in a high-frequency world. *Rev. Financial Stud.* 25:1457–1493.
- Easley D, Lopez de Prado M, O’Hara M (2013) *High-Frequency Trading – New Realities for Traders, Markets and Regulators* (Risk Books, London, UK).
- Fischer A (2004) Price clustering in the FX market: A disaggregated analysis using central bank interventions. Discussion paper, Centre for Economic Policy and Research, no. 4529.
- Financial Times (2019) Yen flash crash: What happened and why. (January 3), <https://www.ft.com/content/45d3807a-0f5b-11e9-a3aa-118c761d2745>.

- Foucault T, Hombert J, Rosu I (2016) News trading and speed. *J. Finance*. 71:335–382.
- Fratzscher M, Gloede O, Menkhoff, L, Sarno L, Stöhr T (2019) When is foreign exchange intervention effective? Evidence from 33 countries. *Am. Econ. J.* 11(1):132–156.
- FXCM (2015) The SNB flash crash. (March 11). Report, Global Brokerage, Inc., New York, NY.
- Gabaix X, Maggiori M (2015) International liquidity and exchange rate dynamics. *Q. J. Econ.* 130(3):1369–1420.
- Galati G, Heath A, McGuire P. (2007) Evidence of carry trade activity. *BIS Q. Rev.* 4:15–18.
- Goldstein M, Kumar P, Graves F (2014) Computerized and high-frequency trading. *The Financial Review*. 49:177–202.
- Gomber P, Arndt B, Lutat M, Uhle T (2011) High-frequency trading. Working paper, Goethe University, Frankfurt, Germany.
- Hagströmer B, Menkveld A (2019) Information revelation in decentralized markets, *J. Finance*. Forthcoming.
- Harris L (2012) Stop the high-frequency trader arms race. *Financial Times* (December 27).
- Hasbrouck J (1991a) Measuring the information content of stock trades. *J. Finance*. 46:179–207.
- Hasbrouck J (1991b) The summary informativeness of stock trades: An econometric analysis. *Rev. Financial Stud.* 4:571–595.
- Hasbrouck J (2007) *Empirical Market Microstructure* (Oxford University Press, New York, NY).
- Hasbrouck J, Saar G (2013) Low-latency trading. *J. Financ. Mark.* 16:646–679.
- Hendershott T, Menkveld A (2014) Price pressures. *J. Financial Econ.* 114:405–423.
- Hendershott T, Jones C, Menkveld A (2011) Does algorithmic trading improve liquidity? *J. Finance*. 66:1–33.
- Hertrich M, Zimmerman H (2017) On the credibility of the Euro/Swiss Franc floor: A financial market perspective. *J. Money Credit Bank.* 49(2–3):567–578.
- Hoffmann P (2014) A dynamic limit order market with fast and slow traders. *J. Financial Econ.* 113:156–169.
- Jermann U (2017) Financial markets' views about the euro-Swiss franc floor. *J. Money Credit Bank.* 49(2–3):553–565.
- Jones C (2013) What do we know about high-frequency trading? Manuscript, Columbia University, New York, NY.
- Jordan T (2014) Introductory remarks by Thomas Jordan. Speech, https://www.snb.ch/en/mmr/speeches/id/ref_20141218_tjn/source/ref_20141218_tjn.en.pdf.
- Jovanovic B, Menkveld A (2016). Middlemen in limit-order markets. Working paper.

- King M, Osler C, Rime D (2012) Foreign exchange market structure, players and evolution. James J, Marsh I, Sarno L, eds. *Handbook of Exchange Rates* (John Wiley & Sons, New York, NY), 3–44.
- Kirilenko A, Lo A (2013) Moore’s law versus Murphy’s law: Algorithmic trading and its discontents. *J. Econ. Perspect.* 27:51–72.
- Kirilenko A, Kyle A, Samadi M, Tuzun T (2017) The flash crash: The impact of high-frequency trading on an electronic market. *J. Finance.* 72:967–998.
- Korajczyk R, Murphy D (2019). High-frequency market making to large institutional trades. *Rev. Financial Stud.* 32:1034–1067.
- Kyle A (1985) Continuous auctions and insider trading. *Econometrica.* 53:1315–1336.
- Lyons RK (2001) *The Microstructure Approach to Exchange Rates* (MIT Press, Cambridge, MA).
- Malinova K, Park A, Riordan R (2018) Do retail traders suffer from high-frequency traders? Working paper.
- Mancini L, Rinaldo A, Wrampelmeyer J (2013) Liquidity in the foreign exchange market: Measurement, commonality, and risk premiums. *J. Finance.* 68:1805–1841.
- McCann C, Yan M (2015) The recent market turmoil spells trouble for ‘auto-liquidators’ like interactive brokers. Technical report, Securities Litigation & Consulting Group, <http://blog.slcg.com/2015/08/the-recent-market-turmoil-spells.html>.
- Menkveld A (2013) High-frequency trading and the new market makers. *J. Financ. Mark.* 16:712–740.
- Menkveld A (2016) The economics of high-frequency trading: taking stock. *Annu Rev. Financ. Econ.* 8:1–24.
- Menkveld A (2018) High-frequency trading as viewed through an electronic microscope. *Financ. Anal. J.* 74(2).
- Menkveld A, Yueshen B (2018) The flash crash: A cautionary tale about highly fragmented markets. *Manag. Sci.* Forthcoming.
- Mirkov K, Pozdeev I, Söderlind P (2016) Toward removal of the Swiss franc cap: Market expectations and verbal interventions. Working paper, Swiss National Bank, no. 10/16.
- Moser D (2016) On the pulse of the financial markets. Speech at the Geneva Money Market Event, https://www.snb.ch/en/mmr/speeches/id/ref_20161117_mo.
- O’Hara M (2015) High-frequency market microstructure. *J. Financial Econ.* 116:257–270.
- O’Hara M, Ye M (2011) Is market fragmentation harming market quality? *J. Financial Econ.* 100:459–474.
- Pagnotta E, Philippon T (2018) Competing on speed. *Econometrica.* Forthcoming.
- Plantin G, Shin HS (2011) Carry trades, monetary policy and speculative dynamics. Discussion paper, Centre for Economic Policy Research, London, UK.


- Raman V, Robe M, Yadav P (2014) Electronic market makers, trader anonymity and market fragility. Manuscript, Chicago Futures Trading Commission, Chicago, IL.
- Rinaldo A, Söderlind P (2010) Safe haven currencies. *Rev. Finance*. 14:385–407.
- Reuters (2015a) SNB's Danthine says cap on franc remains policy cornerstone. (12 January).
- Reuters (2015b) Swiss franc shock shuts some FX brokers. (16 January).
- Rigobon R (2003) Identification through heteroskedasticity. *Rev. Econ. Stat.* 85:777–792.
- Rojcek J, Ziegler A (2016) High-frequency trading in limit order markets: Equilibrium impact and regulation. Working paper.
- Rosu I (2016) Fast and slow informed trading. Manuscript, HEC, Paris, France.
- Securities Exchange Commission (2010) Concept release on equity market structure. Release no. 34-61358, file no. S7-02-10 (September), 27–41.
- Snedecor GW, Cochran WG (1989) *Statistical Methods*, 8th ed. (Iowa State University Press, Ames, IA).
- Subrahmanyam A, Zhen H (2016) Limit order placement by high-frequency traders. *Borsa Istanbul Rev.* 16(4):185–209.
- Swiss National Bank (2011) Swiss National Bank sets minimum exchange rate at CHF 1.20 per euro. Report, Swiss National Bank, Switzerland, http://www.snb.ch/en/mmr/reference/pre_20110906/source/pre_20110906.en.pdf.
- Swiss National Bank (2015) SNB discontinues minimum exchange rate and lowers interest to -0.75% . Report, Swiss National Bank, Switzerland, https://www.snb.ch/en/mmr/reference/pre_20150115/source/pre_20150115.en.pdf.
- Tong L (2015) A blessing or a curse? The impact of high-frequency trading on institutional investors. Working paper.
- van Kervel V, Menkveld A (2019) High-frequency trading around large institutional orders. *J. Finance*. 74:1091–1137.
- Yellen JL (2017) Financial stability a decade after the onset of the crisis. Speech, Jackson Hole Symposium (August 25), <https://www.federalreserve.gov/newsevents/speech/yellen20170825a.htm>.



Source: EBS.

Figure 1: Indicative breakdown of EBS market trading volumes.


This figure illustrates the three trade categories (i.e., human, bank AI and PTC AI) and their average shares of trading volumes around the time of the event.




Source: Bloomberg.

Figure 2: EUR/CHF exchange rate versus the cap set by SNB.

Panel A: EUR/CHF



Panel B: USD/CHF




Sources: EBS and the authors' calculations.


Figure 3. Price and order book around the SNB announcement.

This figure illustrates the price, bid volume and ask volume around the SNB announcement from 09:25 GMT to 11:00 GMT on 15 January 2015. Panels A and B refer to EUR/CHF and USD/CHF, respectively. Price is the average mid-point between the top 10 best bid and ask prices at a 1-second window. Bid volume is calculated as the total volume over the 10 best bid quotes in a millisecond window, further summed up at a 1-second window. Ask volume is calculated in the same way. Each data point plotted in the charts represents price, bid volume and ask volume within a given second. The SNB announcement was made at 09:30 GMT; the bid 'wall' first appeared at 09:53 GMT and seemed to disappear from 10:53 GMT.

Panel A: By type of trader consuming liquidity (i.e., supplied the market order)



Panel B: By type of trader providing liquidity (i.e., supplied the limit order)




Sources: EBS and the authors' calculations.


Figure 4: EUR/CHF trades in the minutes following the SNB announcement.

This figure illustrates the prices and volume by the three types of trader (i.e., human, bank AI and PTC AI) in the 23 minutes following the SNB announcement, depending on whether their trades were consuming liquidity (market orders, Panel A) or providing liquidity (limit orders, Panel B). Each data point plotted in the charts represents a simple average of the trade prices within a given second. Because of averaging, the prices in the top and bottom panels need not be identical.


Panel A.1: EUR/CHF: Cumulative net purchases of CHF



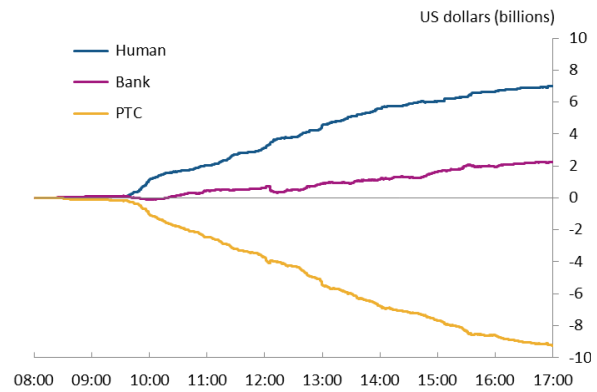
Panel A.2: USD/CHF: Cumulative net purchases of CHF



Panel B.1: EUR/CHF: Cumulative net liquidity provision



Panel B.2: USD/CHF: Cumulative net liquidity provision



Sources: EBS and the authors' calculations.

Figure 5: Market reaction on the day of the SNB announcement.

This figure shows the market reaction on the day of the SNB announcement for EUR/CHF and USD/CHF. Panels A.1 and A.2 show the cumulative net purchase of CHF by the three types of trader (i.e., human, bank AI and PTC AI). Panels B.1 and B.2 show the cumulative net liquidity provision by the three types of trader. ‘Liquidity consumer’ is defined as a trader who submits a market order. ‘Liquidity provider’ is defined as a trader who submits a limit order. The net purchase of CHF for a given type of trader is the difference between the purchasing volume (of the base currency) as a liquidity provider and the selling volume (of the base currency) as a liquidity provider. The net liquidity provision for a given type of trader is the difference between the trading volume as a liquidity provider and trading volume as a liquidity consumer.

Table 1: Liquidity volume by trader type.

This table presents the daily share of trading volume by the three types of trader (i.e., human, bank AI and PTC AI) for EUR/CHF and USD/CHF in the pre-event period (5–14 January), on the event day (15 January) and in the post-event period (16–23 January). Panels A and B present the results for EUR/CHF and USD/CHF, respectively. The daily share of trading volume for a given type of liquidity provider (i.e., submitting the limit order) is first calculated as its volume over the total market volume on a daily basis and is then averaged out within each period. The daily share of trading volume for a given liquidity consumer (i.e., submitting the market order) is calculated in the same manner. The net liquidity provision for a given type of trader is the difference between the daily share of volume as a liquidity provider and daily share of volume as a liquidity consumer. The lower part of Panels A and B shows the result of the t-test on the hypotheses that the daily share of volume is equal between the event day and the pre-event period, and is equal between the post- and pre-event periods. ***, ** and * denote statistical significance at the 1%, 5% and 10% levels, respectively.

Panel A: EUR/CHF

	Liquidity provider			Liquidity consumer			Net liquidity provision		
	Human	Bank AI	PTC AI	Human	Bank AI	PTC AI	Human	Bank AI	PTC AI
<i>Daily share of trading volume (%)</i>									
Pre-event period	38.8	54.1	7.1	19.4	42.0	38.6	19.4	12.1	-31.5
Event day	68.4	25.1	6.5	34.9	26.7	38.4	33.5	-1.6	-31.9
Post-event period	50.1	35.3	14.5	21.5	23.3	55.1	28.6	12.0	-40.6
<i>Statistical tests (t-statistics)</i>									
Event day = pre-event?	6.0 ^{***}	-4.3 ^{***}	-0.3	11.6 ^{***}	-5.0 ^{***}	-0.1	3.1 ^{***}	-2.2 ^{**}	-0.1
Post-event = pre-event?	1.9 ^{**}	-2.5 ^{***}	3.1 ^{***}	1.1	-5.3 ^{***}	4.6 ^{***}	1.7 [*]	0.0	-2.7 ^{***}

Panel B: USD/CHF

	Liquidity provider			Liquidity consumer			Net liquidity provision		
	Human	Bank AI	PTC AI	Human	Bank AI	PTC AI	Human	Bank AI	PTC AI
<i>Daily share of trading volume (%)</i>									
Pre-event period	27.3	31.4	41.3	6.4	18.0	75.6	20.9	13.4	-34.3
Event day	46.5	30.5	23.0	12.3	28.0	59.7	34.2	2.5	-36.7
Post-event period	37.7	40.2	22.1	11.3	26.4	62.4	26.5	13.8	-40.3
<i>Statistical tests (t-statistics)</i>									
Event day = pre-event?	10.2 ^{***}	-0.7	-10.8 ^{***}	9.4 ^{***}	13.1 ^{***}	-18.5 ^{***}	10.3 ^{***}	-8.7 ^{***}	-1.7 ^{**}
Post-event = pre-event?	3.5 ^{***}	3.7 ^{***}	-8.8 ^{***}	3.8 ^{***}	7.4 ^{***}	-7.8 ^{***}	2.1 ^{**}	0.1	-2.2 ^{**}

Sources: EBS and the authors' calculations.

Table 2: Daily effective spreads by type of liquidity provider.

This table presents the effective spreads of EUR/CHF (Panel A) and USD/CHF (Panel B) by the three types of liquidity provider (i.e., human, bank AI and PTC AI) for EUR/CHF and USD/CHF in the pre-event period (5–14 January), on the event day (15 January) and in the post-event period (16–23 January). The effective spread is first calculated trade by trade and is then averaged out within each trading day. The final daily effective spread presented in this table is the average daily effective spread within each period. We use K-sample tests to investigate equality across periods, and Snedecor and Cochran (1989) tests to investigate equality across trader types. ***, ** and * denote statistical significance at the 1%, 5% and 10% levels, respectively.

Panel A: EUR/CHF

	Median spread (<i>basis points</i>)			Statistical tests (<i>p-values</i>)	
	Human	Bank AI	PTC AI	Human = Bank AI?	Human = PTC AI?
Pre-event period	0.21	0.21	0.21	0.76	0.24
Event day	0.92	1.22	1.91	0.02**	0.00***
Post-event period	0.89	0.92	1.72	0.02**	0.00***
Statistical tests (<i>p-values</i>)					
Event day = pre-event?	0.00***	0.00***	0.00***		
Post-event = pre-event?	0.00***	0.00***	0.00***		

Panel B: USD/CHF

	Median spread (<i>basis points</i>)			Statistical tests (<i>p-values</i>)	
	Human	Bank AI	PTC AI	Human = Bank AI?	Human = PTC AI?
Pre-event period	0.24	0.30	0.43	0.00***	0.00***
Event day	0.96	0.94	2.00	0.89	0.00***
Post-event period	0.86	1.05	1.80	0.00***	0.00***
Statistical tests (<i>p-values</i>)					
Event day = pre-event?	0.00***	0.00***	0.00***		
Post-event = pre-event?	0.00***	0.00***	0.00***		

Sources: EBS and the authors' calculations.

Table 3: Contributions to realised variance and realised bipower variation.

This table presents the contributions to realised variance and realised bipower variation for the EUR/CHF and USD/CHF markets by the three types of liquidity consumers (i.e., human, bank AI and PTC AI) in the pre-event period (5–14 January), on the event day (15 January) and in the post-event period (16–23 January). Panel A.1 presents the share of total realised variance, and Panel A.2 is the per-trade share of realised variance. Panel B.1 presents the share of total realised bipower variation, and Panel B.2 shows the per-trade share of total realised bipower variation. The per-trade share of realised variance is normalised such that the average variance per trade across all three periods and all trader types is one. The calculation of all variables in each period use trade-by-trade data within the period. For details about the calculation, see the description in Section 6.1.

Panel A.1: Share of total realised variance

<i>Per cent</i>	Human	Bank AI	PTC AI
<i>EUR/CHF</i>			
Pre-event period	17.4	45.0	37.6
Event day	1.0	90.7	8.3
Post-event period	17.7	46.8	35.6
<i>USD/CHF</i>			
Pre-event period	5.3	17.7	77.1
Event day	6.8	33.7	59.5
Post-event period	8.3	26.4	65.3

Panel A.2: Per-trade share of realised variance

<i>Normalised</i>	Human	Bank AI	PTC AI
<i>EUR/CHF</i>			
Pre-event period	1.14	1.14	0.83
Event day	0.05	3.55	0.15
Post-event period	0.94	1.81	0.64
<i>USD/CHF</i>			
Pre-event period	0.93	0.95	1.02
Event day	0.62	1.22	0.97
Post-event period	0.88	0.98	1.03

Panel B.1: Share of total realised bipower variation

<i>Per cent</i>	Human	Bank AI	PTC AI
<i>EUR/CHF</i>			
Pre-event period	16.4	43.8	39.8
Event day	0.8	91.0	8.2
Post-event period	22.1	29.8	48.1
<i>USD/CHF</i>			
Pre-event period	4.7	21.4	73.9
Event day	9.5	21.2	69.3
Post-event period	9.8	33.2	57.0

Panel B.2: Per-trade share of total realised bipower variation

<i>Normalised</i>	Human	Bank AI	PTC AI
<i>EUR/CHF</i>			
Pre-event period	1.06	1.09	0.86
Event day	0.04	3.41	0.14
Post-event period	0.88	0.87	0.65
<i>USD/CHF</i>			
Pre-event period	0.63	0.87	0.74
Event day	1.24	1.10	1.62
Post-event period	0.92	1.09	0.79

Sources: EBS and the authors' calculations.

Table 4: Price-impact coefficients.

This table presents the estimated price-impact coefficients (β) for the three types of liquidity consumer (i.e., human, bank AI and PTC AI) in the pre-event period (5–14 January), on the event day (15 January) and in the post-event period (16–23 January).

$$r_t = \alpha + \beta_c x_t^c + \epsilon_t,$$

where t is at a 5-minute interval; c is the type of liquidity consumer; r_t is the logarithmic return (on the base currency) between the successive mid-quote of the best bid and ask quotes in the final 100 milliseconds at the 5-minute window, meaning $\ln(\text{midquote}_t/\text{midquote}_{t-1})$; and x_t is the order flow for each type of liquidity consumer during the same 5-minute window, calculated as the difference between liquidity-consuming purchases of the base currency and liquidity-consuming sales of the base currency by each trader type. Panels A and B show the estimated coefficients for EUR/CHF and USD/CHF, respectively. ***, ** and * denote statistical significance at the 1%, 5% and 10% levels, respectively.

Panel A: EUR/CHF

	Human	Bank AI	PTC AI
Pre-event period	0.001	0.000	0.002***
Event day	-0.102	0.000	0.118
Post-event period	0.454***	0.480***	0.151***

Panel B: USD/CHF

	Human	Bank AI	PTC AI
Pre-event period	0.064**	0.073***	0.034***
Event day	0.716	1.421***	-0.371
Post-event period	0.244**	0.516***	-0.023

Sources: EBS and the authors' calculations.

Table 5: Contributions to variance of efficient price.

This table presents the contributions to the variance of efficient price by the three types of liquidity consumer (i.e., human, bank AI and PTC AI) in the pre-event period (5–14 January), on the event day (15 January) and in the post-event period (16–23 January). A number of the most extreme returns on the event day were excluded to avoid allowing them to drive the results. The returns and order flow are defined in Table 4. The procedure of estimating the contributions to variance of efficient price by the three trader types is described in Section 6.2. Panels A and B show the contributions for EUR/CHF and USD/CHF, respectively.

Panel A: EUR/CHF

<i>Per cent</i>	Returns	Order flow		
		Human	Bank AI	PTC AI
Pre-event period	63.8	4.4	0.3	31.5
Event day	11.8	69.2	18.1	0.9
Post-event period	39.9	27.3	19.4	13.4

Panel B: USD/CHF

<i>Per cent</i>	Returns	Order flow		
		Human	Bank AI	PTC AI
Pre-event period	79.5	4.3	11.0	5.2
Event day	19.2	17.9	26.1	36.9
Post-event period	85.3	3.1	11.3	0.2

Sources: EBS and the authors' calculations.

Table 6: Arbitrage opportunities between EUR/CHF, USD/CHF and EUR/USD.

In this table, Panel A presents the average frequency and size of the arbitrage opportunities between EUR/CHF, USD/CHF and EUR/USD at a 1-minute window in the pre-event period (5–14 January), on the event day (15 January) and in the post-event period (16–23 January). An arbitrage opportunity is recorded if the combination of best bid and ask quotes across the three currency pairs in a 100-millisecond window offers a profit in excess of one basis point. Frequency is calculated as the number of arbitrage opportunities over the number of 100-milliseconds in 1 minute (i.e., 600) and is further averaged out within the pre-event period, the event day and the post-event period. Profitability is the average profitability of arbitrage opportunities where they exist and is further averaged out within each period. Panel B presents the estimated contemporary coefficients of order flow by each trader type on the frequency of arbitrage opportunities over a 5-minute window in the pre- and post-event periods. The associated SVAR model and its estimation are described in Section 6.3. ***, ** and * denote statistical significance at the 1%, 5% and 10% levels, respectively.

Panel A: Size and frequency of opportunities

	Frequency <i>Per cent</i>	Profitability <i>Basis points</i>
Pre-event period	0.004	9.9
Event day	0.897	181.2
Post-event period	0.046	4.8

Panel B: Trading on arbitrage opportunities

<i>Coefficients</i>	Pre-event period	Post-event period
Human	0.0022	-0.0010
Bank AI	-0.0061***	-0.0025
PTC AI	-0.0089***	0.0050

Sources: EBS and the authors' calculations.

Table 7: Direct intervention–adjusted EUR/CHF liquidity volumes by trader type.

This table presents a comparative result of Panel A in Table 1 by excluding the estimated SNB trades in the EUR/CHF market. The procedure to estimate SNB trades is described in Section 7.3. ‘Liquidity provider’, ‘liquidity consumer’, ‘net liquidity provision’ and ‘daily share of trading volume’ are defined in Table 1. ***, ** and * denote statistical significance at the 1%, 5% and 10% levels, respectively.

	Liquidity provider			Liquidity consumer			Net liquidity provision		
	Human	Bank AI	PTC AI	Human	Bank AI	PTC AI	Human	Bank AI	PTC AI
<i>Daily share of trading volume (%)</i>									
Pre-event period	38.7	54.2	7.1	19.4	41.9	38.6	19.3	12.3	-31.5
Event day	63.9	28.7	7.5	33.5	25.6	40.9	30.4	3.1	-33.4
Post-event period	50.1	35.4	14.5	21.6	23.3	55.2	28.6	12.1	-40.6
<i>Statistical tests (t-statistics)</i>									
Event day = pre-event?	5.1 ^{***}	-3.8 ^{***}	0.2	10.6 ^{***}	-5.4 ^{***}	0.9	2.4 ^{**}	-1.5	-0.7
Post-event = pre-event?	1.9 ^{**}	-2.5 ^{**}	3.1 ^{***}	1.1	-5.3 ^{***}	4.6 ^{***}	1.7 [*]	-0.0	-2.7 ^{***}

Sources: EBS and the authors’ calculations.