



BANK OF ENGLAND

Staff Working Paper No. 711

Judgement Day: algorithmic trading around the Swiss franc cap removal

Francis Breedon, Louisa Chen, Angelo Ranaldo and
Nicholas Vause

February 2018

Staff Working Papers describe research in progress by the author(s) and are published to elicit comments and to further debate. 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.



BANK OF ENGLAND

Staff Working Paper No. 711

Judgement Day: algorithmic trading around the Swiss franc cap removal

Francis Breedon,⁽¹⁾ Louisa Chen,⁽²⁾ Angelo Ranaldo⁽³⁾ and Nicholas Vause⁽⁴⁾

Abstract

A key issue raised by the rapid growth of computerised algorithmic trading is how it responds in extreme situations. Using data on foreign exchange orders and transactions that includes identification of algorithmic trading, we find that this type of trading contributed to the deterioration of market quality following the removal of the cap on the Swiss franc on 15 January 2015, which was an event that came as a complete surprise to market participants. In particular, we find that algorithmic traders withdrew liquidity and generated uninformative volatility in Swiss franc currency pairs, while human traders did the opposite. However, we find no evidence that algorithmic trading propagated these adverse effects on market quality to other currency pairs.

Key words: Swiss franc, algorithmic trading, liquidity, volatility, price discovery, arbitrage opportunities.

JEL classification: G14, G23.

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

(2) School of Business, Management and Economics, University of Sussex. Email: l.x.chen@sussex.ac.uk. Chen was with the Bank of England when the paper was written.

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

(4) Bank of England. Email: nicholas.vause@bankofengland.co.uk

The views expressed in this paper are those of the authors, and not necessarily those of the Bank of England or its committees. 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 also grateful to Patrick Schaffner for excellent research assistance, and participants at the 2016 conference on Microstructure of Foreign Exchange Markets at Cambridge University, the 2016 conference on Financial Determinants of Foreign Exchange Rates at the Bank of England and a 2017 seminar at the Swiss National Bank for their comments. 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)). 'Judgement Day' followed 'Rise of the Machines' in the Terminator film franchise.

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

Publications and Design Team, Bank of England, Threadneedle Street, London, EC2R 8AH
Telephone +44 (0)20 7601 4030 email publications@bankofengland.co.uk

Contents

1	Introduction	3
2	Literature review	5
3	Market structure and data	6
3.1	Market structure	6
3.2	Data	7
4	Overview of trading patterns around the Swiss franc event	8
4.1	The Swiss franc event	8
4.2	Overview of trading patterns	9
5	Liquidity contributions of computer and human trading	13
5.1	Volumes of liquidity provided and consumed	13
5.2	Effective spreads	15
6	Impact of computer and human trading on volatility	16
6.1	Contributions to realised volatility	17
6.2	Contributions to efficient pricing	18
7	Arbitrage opportunities and market efficiency	21
8	Non-CHF foreign exchange rates	22
9	Conclusion	24
10	References	26

1 Introduction

New technologies have dramatically changed financial markets. One of the main innovations of recent years is computerised 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. It can improve market liquidity by reducing transaction costs (Hendershot *et al.*, 2011) and the reliance on financial intermediaries (Menkveld, 2013). It can also make security prices more efficient, in the sense that they better reflect fundamental values (O’Hara, 2015). It can even reduce risks by lessening the impact of human feelings on overall investor behaviour, such as panic reactions and herding behaviour. On the other hand, AT may have less desirable implications because it can increase market power over slower traders (Hoffmann, 2014), raise adverse selection (Biais *et al.*, 2015), excess volatility or extreme market movements (Foucault *et al.*, 2016) and so potentially harm financial stability. It is this last issue we focus on in this paper.

We analyse the role of AT in foreign exchange (FX) markets in a period containing the 15 January 2015 announcement by the Swiss National Bank that it had discontinued its policy of capping the value of the Swiss franc against the euro. This ‘Swiss franc event’ represents a natural experiment as one of the largest shocks to the FX market in recent years and probably the most significant ‘black swan’ event in the period in which AT has been a prominent force in FX markets.¹ In particular, we study the contribution of AT and human traders to two important dimensions of market quality, namely liquidity and price efficiency. Our analysis is based on a unique dataset with a detailed identification of AT obtained from EBS Market, which is the leading platform for electronic spot FX trading in many of the major currencies.²

A detailed understanding of AT in distressed situations is important for at least two reasons. First, a better comprehension of whether AT is beneficial or detrimental for market quality in extreme situations would help inform the ongoing reform of trading venues, as pursued by Regulation NMS in the United States and MiFID I and II in Europe. Second, the resilience of an exchange system depends on the behaviour of different types of market participant and their reciprocal influence on each other. For instance, a tendency of AT to offer liquidity in calm markets and withdraw it in distressed situations could lead less sophisticated agents to become reliant on high levels of market liquidity only to find it in short supply when they most needed it. If these adverse consequences of AT were predominant or not offset by other traders, then AT could represent a systemic threat to the whole trading system. To shed light on this key issue for financial stability, we analyse whether human traders and AT substitute for or complement each other in supplying and consuming liquidity.

¹ A ‘black swan’ is a metaphor that describes an event that comes as a complete surprise and has a major effect. The term is based on an ancient saying that presumed black swans did not exist, but the saying was rewritten after black swans were discovered in the wild.

² While EBS has supplied the relevant market data, it does not endorse or support any conclusions made in this paper and has not contributed to any of the analysis in it.

We proceed in three steps. First, we describe the EBS Market platform and our sample of data from it. The platform is the central limit order book for spot FX operated by EBS Service Company Limited, which is part of NEX Markets, 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 around the Swiss franc event. Second, we perform an in-depth analysis of market liquidity and price movements by decomposing order flow, effective spreads and intraday volatility by type of trader. This enables us to highlight the contribution of AT and human traders to liquidity provision and consumption, transaction costs and realised volatility. Third, we study the contribution to efficient pricing of AT and human traders. To do this, we analyse the formation of efficient prices by performing a vector autoregression (VAR) as in Hasbrouck (1991(a), 1991(b), 2007), but conditioning on traders' types as in Hendershott *et al.* (2011). We substantiate the analysis of price efficiency by looking at arbitrage opportunities (Chaboud *et al.*, 2014) and variance ratios (O'Hara and Ye, 2011).

Our study delivers two important findings. First, in reaction to the Swiss franc event, we find that AT tended to consume liquidity and reinforce the price disruption. Opposite and offsetting patterns apply for human traders, who supported market quality by providing liquidity and aiding price discovery. Second, we find that this market quality degradation coming from AT was concentrated in the shocked FX rate (EUR/CHF) and, to a lesser extent, USD/CHF. Non-CHF currency pairs (USD/JPY, EUR/JPY and EUR/USD in our sample) were essentially unaffected.³ This suggests that AT models were somewhat compartmentalised, which, along with human trading, helped to sustain market quality beyond the CHF currency pairs.

Of course, like Tolstoy's comment on unhappy families, each period of market distress is distressed in its own way. This limits our ability to draw general conclusions. However, as in previous papers investigating important shocks such as the failure of Lehman Brothers in 2008 (Afonso *et al.*, 2011) or the 'Flash Crash' in 2010 (Kirilenko *et al.*, 2017), our analysis of a single event should provide indicative evidence on these broader issues.

Our paper contributes to the growing literature on algorithmic trading, which we survey in the next section. Except for the thorough analysis by Chaboud *et al.* (2014), prior research on FX algorithmic trading is scant. While that work shows AT improves price efficiency in 'normal' market conditions, we find that AT reactions to an extreme event were detrimental to market liquidity and price resilience, whereas human traders sustained market quality during the event. FX markets have a different structure to other markets where AT is prevalent, with more trading conducted bilaterally and on multilateral platforms, and less on centralised exchanges. Hence, it is important to build the literature on FX markets.

The paper proceeds as follows. After briefly describing related literature (Section 2) and our data (Section 3), we give an overview of trading patterns around the Swiss franc event (Section 4). We then undertake a more formal analysis of liquidity (Section 5), volatility

³ The currency codes are: CHF for Swiss francs, EUR for euros, JPY for Japanese yen and USD for US dollars.

(Section 6) and market efficiency (Section 7) for Swiss franc currency pairs, followed by analysis of non-CHF FX rates (Section 8).

2 Literature review

The literature on AT has grown substantially in the last few years.⁴ The theoretical literature focuses on how high-frequency trading (HFT), which is the most commonly discussed form of AT, affects liquidity and price efficiency. The HFT community's ability to revise their quotes quickly after news arrives reduces the winner's curse problem but creates disincentives for trading by slower traders (Hoffmann, 2014; and Jovanovich and Menkveld, 2016), including by increasing adverse selection and price impact (Foucault *et al.*, 2016). Aït-Sahalia and Saglam (2014) and Rosu (2016) model HFT that is 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 that produces adverse selection, negative externalities and over-investment in equilibrium.⁵

The focus of prior empirical research has been the stock market. As described by O'Hara (2015), the consensus is that HFT market making enhances market quality by reducing spreads and raising informational efficiency.⁶ However, it is not clear whether HFTs go with or lean against the wind, that is amplify price falls (rises) by actively selling (buying) or dampening them by actively buying (selling) (see Korajczyk and Murphy, 2016; van Kervel and Menkveld, 2016; Breckenfelder, 2013; and Tong, 2015).

By analysing a 'black swan' event, our paper is related to the stock market literature, which provides mixed evidence on the role of HFT in distressed markets. On the one hand, HFTs are found to withdraw from their market-making role during 'flash crashes' (see *e.g.* CFTC-SEC, 2010; Easley *et al.*, 2012; Kirilenko *et al.*, 2017; and Menkveld and Yueshen, 2015) or when market conditions become unfavourable (see *e.g.* Raman *et al.*, 2014; Anand and Venkataraman, 2015; and Korajczyk and Murphy, 2015). On the other hand, HFT provides liquidity and absorbs imbalances created by non-high frequency traders around large price movements (Brogaard *et al.*, 2017) and the short-interval return volatility of most stocks varies inversely with a market-wide measure of correlated HFT strategies (Boehmer, Li, and Saar, 2016).

The only earlier paper that provides an in-depth analysis of AT in the FX market is Chaboud *et al.* (2014). They show that algorithmic trading is associated with a reduction in arbitrage opportunities while AT liquidity provision decreases return autocorrelation.

⁴ There are many excellent surveys on AT and HFT, including recent papers by Biais and Woolley (2011), 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), SEC (2010), O'Hara (2015) and Menkveld (2016, 2017).

⁵ The role of AT in fragmented markets with multiple exchanges is studied in Pagnotta and Philippon (2015). 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, general limit order book and latency.

⁶ See, for example, Boehmer *et al.* (2015), Brogaard *et al.* (2015), Carrion (2017), Conrad *et al.* (2015), Hasbrouck and Saar (2013), Hendershott *et al.* (2011), Menkveld (2013) and Malinova *et al.* (2013).

3 Market structure and data

3.1 Market structure

Two significant global electronic spot trading platforms in major currency pairs are EBS and Reuters. USD/CHF and EUR/CHF, which are the focus of this study, trade primarily on EBS (see King, *et al.*, 2011). Prices on EBS also constitute the reference 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 trading of futures on the Chicago Mercantile Exchange was suspended and over-the-counter trading had largely disappeared (Hagströmer and Menkveld, 2016).

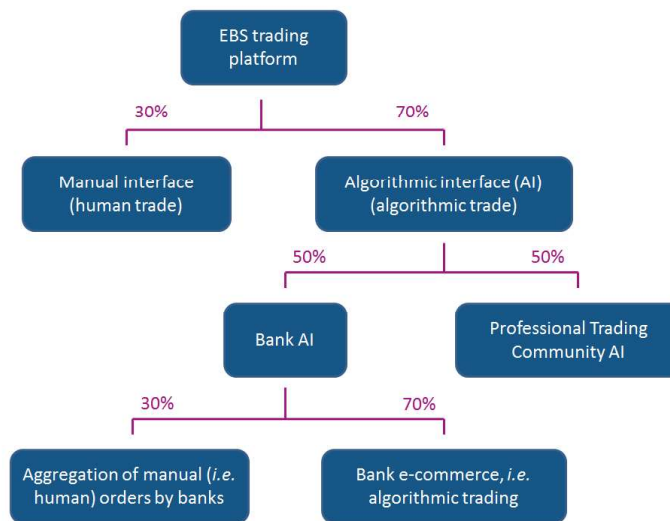
EBS Market is an order-driven electronic trading system, which unites buyers and sellers of spot FX across the globe on a pre-trade anonymous central limit order book. EBS is accessible to foreign exchange dealing banks and, under the auspices of dealing banks (via prime brokerage arrangements), to hedge funds and commodity trading advisors (CTAs). EBS controls the network and each of the terminals on which trading is conducted, and records whether a trade is placed 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 what types of trading they engage in. This allows EBS to decompose the algorithmic trades into two different categories, referred to as 'bank AI' and 'PTC AI', where PTC stands for Professional Trading Community and AI for algorithmic interface, which is how EBS labels the direct computer interfaces. Market data at this level of granularity is not ordinarily sold or distributed by EBS to third parties. PTC essentially refers to principal trading firms, hedge funds and commodity trading advisors, which can trade directly on EBS under the auspices of dealing banks (via prime brokerage arrangements). Trading through this route is all AT in the sense that trades are initiated by computers. Bank AI is harder to categorise since it includes a significant share (EBS estimates approximately 30%) of aggregators that are computer-based trading systems that simply process orders received from bank customers and then execute them algorithmically. The source of these trades may be described as 'agency algorithms' or 'smart execution algorithms', as discussed by Gomber *et al.* (2011), rather than the proprietary algorithms usually highlighted in the AT literature. Note that the former category also includes 'auto-liquidation' algorithms that respond to margin calls, which have been highlighted as a source of price distortions during crashes (*e.g.* McCann and Yan, 2015). Note that PTC firms or banks may also submit manual orders. The three trade categories and their shares of average transaction volumes in EBS Market are represented in Figure 1.

Computer-based trading classified as PTC or bank AI accounts for approximately 70% of EBS Market transaction volume. However, the system includes features designed to prevent strategies where 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 Market, so that once a good-

till-cancelled order is submitted, it cannot be cancelled for 250 milliseconds.⁷ Second, and more importantly, EBS Market 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 milliseconds⁸, in which 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 Market. Note that this is not analogous to the frequent batch auction system described by Budish, Crampton and Shim (2015), which eliminates sniping of stale quotes, but is more like the random order delay system of Harris (2012).

Figure 1: Indicative breakdown of EBS trading volumes



Source: EBS.

As EBS Market is a ‘wholesale’ trading system, the minimum trade size over our sample period is one million of the base currency⁹, and trade sizes are only allowed in a multiple 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 during 5-23 January 2015. 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. Throughout the subsequent analysis, we focus on data between 8.00 GMT and 17.00 GMT,

⁷ A MQL of 250 milliseconds was applied in the EUR/USD, USD/JPY, USD/CHF, EUR/CHF and EUR/JPY currency pairs on EBS Market for the dates referenced in this paper.

⁸ For the dates referenced in this paper.

⁹ The base currency is the first currency displayed in the symbol of the currency pair. For example, the euro is the base currency of EURCHF.

which represent the effective trading hours of the day. We also exclude weekends, when the EBS trading platform is inactive.^{10,11}

The transaction data set records the time stamp to the nearest millisecond of each trade that occurred, along with the actual transaction price, the amount transacted, and the direction of the trade. Importantly for our study, the nature of the party providing liquidity (submitting the good-till-cancelled (GTC) order) and consuming liquidity (submitting the immediate-or-cancel (“IOC”) order) is categorised by EBS as human, bank AI or PTC AI.¹² Thus, each trade may be classified according to nine different 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 quote data, which are available at 100 millisecond intervals, so 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.¹³ Unlike in the transactions data, the type of trader that posts each quote is not available to us.

4 Overview of trading patterns around the Swiss franc event

4.1 The Swiss franc event

The Swiss National Bank (SNB) began intervening in FX markets 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 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).

Following the introduction of this policy, the franc generally traded a little below its cap (Figure 2). It pushed against it for a period in mid-2012, before appreciation pressures again intensified towards the end of 2014 due to weakness in the euro-area economy. 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). Similarly,

¹⁰ See Chaboud *et al.* (2014) for further discussion of trading activity on the EBS system. In addition, EBS indicated to us that their dealing services are ordinarily open for trading 24 hours a day, 7 days a week, with the exception of a maintenance window that ordinarily occurs from 5:50pm New York time on a Friday until Saturday morning. For the purpose of computing market data products, such as ‘highs’ and ‘lows’, EBS regards trades between 5pm Friday New York time and 5am Monday Sydney time as not conducted in normal market conditions or market hours and excludes these trades from the calculations. However, the trades remain in EBS’ data.

¹¹ In our sample, we drop five transactions and 24 quotes on EUR/CHF that took place between 09:32:29 GMT and 09:32:39 GMT on January 15, where the price was exceptionally low at 0.0015, with volume of one million of the base currency for each transaction. EBS confirmed that those transactions turned out to be errors made by the traders, and the counterparties have settled solutions outside EBS.

¹² GTC and IOC orders respectively align closely with the concepts of limit and market orders in the broader market microstructure literature.

¹³ The historical market data provided by EBS is time-sliced at 100 milliseconds and is therefore a snapshot of the activity during previous time-period. Consequently the quote and paid/given trade data provided are a summary and not a full life-cycle of every event.

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

Figure 2: EUR/CHF exchange rate versus the cap set by the Swiss National Bank



Source: Bloomberg

Despite that, the policy of capping the value of the franc against the euro was discontinued at 9:30 GMT on 15 January 2015. A press release gave the following explanation: “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”. This news came as a complete surprise to market participants, as was reflected both in FX options prices leading up to the announcement (Mirkov *et al.*, 2016) and financial reporting after it (*e.g.* Reuters, 2015b; Bloomberg, 2015).


4.2 Overview of trading patterns

In this section we provide an overview of the market reaction to the SNB announcement on 15 January 2015. In particular, we illustrate graphically some of the key features of human and algorithmic trading around the announcement as a prelude to the more formal analysis in the rest of the paper. We first focus on the seconds around the announcement, then the minutes, and finally the hours.

Figure 3 plots the volume of GTC orders to buy (‘bid’) or sell (‘ask’) euros for Swiss francs as well as the mid-price, which is the mid-point between the best bid and ask prices, during the 90 seconds following the announcement. Not shown, for reasons of sensitivity, is a sharp fall in outstanding orders to sell francs in the seconds approaching 9:30. While our data from EBS does not include the identities of the institutions submitting the orders, we presume this was driven by the SNB in preparation for its 9:30 announcement. The chart does show that it was not until about 44 seconds after the announcement that the market reacted significantly. At this time, the mid-price started to reflect a very rapid appreciation of the franc, as both sides

of the order book collapsed in size.¹⁴ Indeed, for a few seconds during the minute of 9.31, there were no orders to buy euros in exchange for Swiss francs at any price. We take the delayed response to the announcement as further evidence that it was not anticipated.

Figure 3: EUR/CHF price and orders in the seconds following the SNB announcement



Sources: EBS and authors' calculations.


Figure 4 shows the prices at which different types of trader exchanged euros for Swiss francs in the 30 minutes following the SNB announcement depending on whether their trades were consuming liquidity (top panel) or providing it (bottom panel). Trades that consume liquidity result from IOC orders, while those that provide it result from GTC orders. The top panel shows that bank AIs consumed liquidity at extreme prices (prices significantly different to those of immediately preceding trades) on a number of occasions, notably between 9.31 and 9.36. Thus, over 75% of the cumulative appreciation of the franc in the 20 minutes to 9.50 was attributable to bank AIs, which accounted for 61% of the volume of liquidity-consuming trades. Indeed, we show below that bank AIs accounted for an even larger share of the realised variance of the EUR/CHF rate at this time. The lower panel shows that bank AIs also provided liquidity for some of the extreme-price trades. That bank AIs both consumed and provided liquidity at extreme prices may reflect the diverse set of traders from whom these trades may originate. This includes not only the different banks but also their various clients. In addition, a roughly equal number of extreme-price trades were accommodated by human traders. Indeed, human traders accounted for a significantly higher share of liquidity-providing trades (50%) than they did for liquidity-consuming trades (19%) during the 20 minutes to 9.50 when the Swiss franc appreciated sharply.¹⁵

¹⁴ 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 seconds later that the price started to fall.


¹⁵ Figure 4 excludes trades that we infer (as discussed later) may have involved the SNB, though we found few such trades in the first twenty minutes after the announcement.

Figure 4: EUR/CHF trades in the minutes following the SNB announcement⁽¹⁾

By type of trader consuming liquidity (i.e. supplied the IOC order)



By type of trader providing liquidity (i.e. supplied the GTC order)




⁽¹⁾ Each data point plotted in the charts represents a simple average of trade prices within a given second. Due to averaging, the prices in the top and bottom panels need not be identical.
Sources: EBS and authors' calculations.


Figure 5 gives an overview of the reaction of both the EUR/CHF and USD/CHF markets to the SNB announcement over the whole trading day of 15 January 2015. The first row shows that the Swiss franc appreciated extremely sharply against both 'base' currencies in the first 20 minutes following the announcement, but that sizeable portions of these gains were reversed in the subsequent hour. After that the two spot rates were much more stable, with the Swiss franc worth about 10% more than at the start of the day. The second row shows that algorithmic traders were net purchasers of Swiss francs over the day, particular bank AIs against the euro and PTC AIs against the US dollar, while human traders were net purchasers of the base currencies. Thus, computers traded 'with the wind', buying the franc as it appreciated, while humans 'leaned against the wind'. Note, however, that human traders did not make net purchases of the base currencies in the key 20-minute period immediately after the announcement. Finally, the third row shows that human traders were consistently net suppliers of liquidity over the day, while PTC AI trades consumed it. However, as we shall see below, net liquidity consumption by PTC AIs is not unusual in these two currency-pairs.

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


EUR/CHF: spot exchange rate




USD/CHF: spot exchange rate




EUR/CHF: cumulative net purchases of CHF




USD/CHF: cumulative net purchases of CHF



EUR/CHF: cumulative net liquidity provision



USD/CHF: cumulative net liquidity provision



Sources: EBS and authors' calculations.

Overall, our graphical overview suggests that algorithmic traders contributed both to the liquidity dry-up (as they were net consumers of liquidity) and the price move (as they were net purchasers of the appreciating currency) after the SNB announcement. In the following sections, we examine these issues in more detail.

5 Liquidity contributions of computer and human trading

In this section we investigate the contributions of computer and human traders to the liquidity of EUR/CHF and USD/CHF markets before, during and after the day of the SNB announcement. We study both quantity-based and price-based indicators of liquidity.

5.1 Volumes of liquidity provided and consumed

First, we focus on volumes of liquidity. As in Section 4, we identify whether the consumer and provider of liquidity for each trade was a human (H), a bank AI (B) or a PTC AI (P). We then record the shares of total trade volumes in three different periods for which providers and consumers of liquidity were H, B or P. The three periods are a ‘pre-event’ period (5-14 January 2015, excluding weekends), the ‘event day’ (15 January 2015) and a ‘post-event’ period (16-23 January 2015, excluding weekends). Finally, we record ‘net liquidity provision’, which is the difference between the share of trades for which a given type of trader provided liquidity and the share for which it consumed liquidity. The results are reported in Table 1.

Table 1: Liquidity volumes by trader type

<i>EUR/CHF</i>									
	Liquidity provider			Liquidity consumer			Net liquidity provision ⁽¹⁾		
	Human	Bank AI	PTC AI	Human	Bank AI	PTC AI	Human	Bank AI	PTC AI
<i>Share of trade 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)⁽²⁾</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***
<i>USD/CHF</i>									
	Liquidity provider			Liquidity consumer			Net liquidity provision ⁽¹⁾		
	Human	Bank AI	PTC AI	Human	Bank AI	PTC AI	Human	Bank AI	PTC AI
<i>Share of trade volume (%)</i>									
Pre-event period	3.8	15.4	69.1	16.6	31.5	40.3	-12.8	-16.1	28.9
Event day	8.9	25.6	60.3	34.1	34.6	26.2	-25.1	-9.0	34.1
Post-event period	7.7	24.9	56.3	27.6	39.0	22.4	-19.8	-14.1	34.0
<i>Statistical tests (t-statistics)⁽²⁾</i>									
Event day = pre-event?	17.6***	10.0***	-8.8***	20.3***	2.4**	-11.6***	20.8***	-6.6***	-4.4***
Post-event = pre-event?	8.7***	5.9***	-9.3***	8.1***	3.7***	-9.0***	5.6***	-0.87	-2.9***

⁽¹⁾ Share of volume as liquidity provider minus share of volume as liquidity consumer.

⁽²⁾ *** / ** / * denote statistical significance at the 1% / 5% / 10% level.

Sources: EBS and authors' calculations.

The table shows that net liquidity provision by AIs fell on the event day and in the post-event period compared with the pre-event period, while it increased for humans. The increases in human net liquidity provision were all statistically significant, while the decreases for algorithmic traders were significant for at least one type of AI. In EUR/CHF, net liquidity

provision by bank AIs declined significantly on the event day, while it was that of PTC AIs that declined significantly in the post-event period. In USD/CHF, net liquidity provision by both types of AI declined significantly on the event day and that of PTC AIs declined significantly in the post-event period. Economically, the most significant changes were the decline in the share of bank AI liquidity-providing trades on the event day in EUR/CHF and the offsetting increase in the human share. Thus, although bank AIs also increased their share of liquidity-consuming trades on the event day, and humans reduced their share, this was not enough to change the pattern in net liquidity provision.

One possible explanation for the importance of human traders in supporting liquidity on the event day and in the post-event period is that the SNB was an active trader in this category. If that were the case, it would be hard to draw general conclusions from our analysis about the role of human traders in extreme events. Hence, we have estimated how the SNB may have traded during this period. We need to make estimates, based on assumptions, as our data from EBS is anonymous so we cannot identify from this any trades of the SNB or any other individual institution. We then repeat the analysis underpinning Table 1 excluding these estimated trades to see if the results are materially affected.

Our estimates of SNB trading activity are based on three assumptions. First, that SNB activity over our full sample period cannot have exceeded about 16% of the turnover on EBS. This figure is derived from public data on SNB sight deposits, which is the key liability created by FX interventions. It is an upper bound because other public documents suggest that the SNB favours a diversified approach to FX interventions (Moser, 2016), including substantial use of telephone orders (Fischer, 2004). Hence, only a share of any such activity may take place directly through EBS. The other assumptions are based on Fischer (2004), who notes that SNB interventions cluster around prices ending in 00 or 50, and that they tend to come in bursts of many transactions within short periods. Hence, we assume that bid orders ending in 00 or 50 originated from the SNB if there was a trade at the same price within 100 milliseconds with a human liquidity provider that was buying EUR. Our strategy does not estimate IOC orders placed by the SNB, so we may overestimate net liquidity provision by the SNB.

Table 2 shows adjusted results for liquidity provision and consumption by computers and humans excluding estimated SNB trades. It only shows results for EUR/CHF, as we made relatively more adjustments for this currency pair. These results are very similar to those in Table 1. In particular, we still find that human traders supported liquidity on the event day and in the post-event period, offsetting a reduction in net liquidity provision by computers. Hence, we proceed using the full data set in the rest of the paper.

Table 2: Adjusted EUR/CHF liquidity volumes by trader type

	Liquidity provider			Liquidity consumer			Net liquidity provision ⁽¹⁾		
	Human	Bank AI	PTC AI	Human	Bank AI	PTC AI	Human	Bank AI	PTC AI
<i>Share of trade 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 test (t-statistics) ⁽²⁾</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 ^{***}

⁽¹⁾ Share of volume as liquidity provider minus share of volume as liquidity consumer.

⁽²⁾ *** / ** / * denote statistical significance at the 1% / 5% / 10% level.

Sources: EBS and authors' calculations.

5.2 Effective spreads

We now turn our focus to a price-based indicator of liquidity. We saw in Section 5.1 that computers reduced their net volume of liquidity provision on the event day and afterwards, but did they widen bid-ask spreads on trades for which they did still provide liquidity? To investigate this question, we calculate a series of effective spreads, s , for each of our 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, *i.e.* supplying the GTC order (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.

Table 3 shows median effective spreads for EUR/CHF and USD/CHF. We report median values as our calculated spreads include some extreme observations, particularly on the event day. For the same reason, the tests of equality of the medians reported in the table are non-parametric tests.¹⁶ For EUR/CHF, median effective spreads were very similar across trader types in the pre-event period. They increased very sharply for all trader types on the event day, but significantly more so for computer trades than human trades, 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, but all spreads again increased very sharply on the event day, particularly those of PTC AIs, and they remained much wider in the post-event period than in the pre-event period.

¹⁶ In particular, we use K-sample tests to investigate equality across periods and Snedecor and Cochran (1989) tests to investigate equality across trader types.

Table 3: Effective spreads by type of liquidity provider*EUR/CHF*

	Median spread (<i>basis points</i>)			Statistical tests ⁽¹⁾ (<i>p-values</i>)	
	Human	Bank AI	PTC AI	Human = Bank AI?	Human = PTC AI?
Pre-event period	0.21	0.21	0.21	0.45	0.41
Event day	0.92	1.22	1.91	0.03**	0.00***
Post-event period	0.91	0.91	1.70	0.12	0.00***
Statistical tests ⁽¹⁾ (χ^2 statistics)					
Event day = pre-event?	303.1***	413.3***	222.9***		
Post-event = pre-event?	830.8***	880.3***	826.6***		

USD/CHF

	Median spread (<i>basis points</i>)			Statistical tests ⁽¹⁾ (<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	0.38	0.00***
Post-event period	0.86	1.05	1.80	0.00***	0.00***
Statistical tests ⁽¹⁾ (χ^2 statistics)					
Event day = pre-event?	71.7***	73.3***	97.1***		
Post-event = pre-event?	429.6***	666.7***	864.2***		

⁽¹⁾ The tests are described in footnote 18. *** / ** / * denote statistical significance at the 1% / 5% / 10% level.

Sources: EBS and authors' calculations.

Taking our quantity-based and price-based indicators of market liquidity together, we conclude that those classified as AI traders significantly reduced their net volume of liquidity provision and, where they did still provide liquidity, this was only at much wider spreads. In contrast, human traders significantly increased their net volume of liquidity provision relative to AIs and did so at narrower spreads than algorithmic traders.

6 Impact of computer and human trading on volatility

A second important dimension of market quality, alongside liquidity, is pricing efficiency. In theory, in an efficient market any gaps that might arise between the actual price of an asset and its 'efficient price', which reflects its fundamental value, tend to be small and closed quickly by traders drawing on the available information about the fundamental value.¹⁷ As a result there would be little excess volatility in prices, *i.e.* on top of that attributable to changes in the efficient price. In this section, we will examine the contributions of different trader types to efficient pricing of the EUR/CHF and USD/CHF exchange rates around the SNB announcement. However, we begin relatively simply by looking at contributions of different trader types to the realised variance of returns.

¹⁷ As Hendershott and Menkveld (2014) note, the net provision of limit orders need not be a good measure of liquidity supply, since a market order that leans against price pressure (goes against the prevailing market trend) can be thought of as a contribution to liquidity and reducing volatility.

6.1 Contributions to realised volatility

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

$$V_k = \sum_{n=1}^{N_t} (r_n d_{nk})^2 / \sum_{n=1}^{N_t} r_n^2$$

where k indexes our different types of trader (human, bank AI or PTC AI) and 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 type and is otherwise zero. 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. Results of this breakdown applied to EUR/CHF and USD/CHF during our three periods are shown in the left-hand panel of Table 4.

Table 4: Contributions to realised variance

<i>Share of total variance</i>				<i>Per-trade share of variance</i>			
<i>Per cent</i>	Human	Bank AI	PTC AI	<i>Normalised⁽¹⁾</i>	Human	Bank AI	PTC AI
<i>EUR/CHF</i>				<i>EUR/CHF</i>			
Pre-event period	17.4	45.0	37.6	Pre-event period	1.14	1.14	0.83
Event day	1.0	90.7	8.3	Event day	0.05	3.55	0.15
Post-event period	17.7	46.8	35.6	Post-event period	0.94	1.81	0.64
<i>USD/CHF</i>				<i>USD/CHF</i>			
Pre-event period	5.3	17.7	77.1	Pre-event period	0.93	0.95	1.02
Event day	6.8	33.7	59.5	Event day	0.62	1.22	0.97
Post-event period	8.3	26.4	65.3	Post-event period	0.88	0.98	1.03

⁽¹⁾ Such that the average variance per trade across all three periods and all trader types is one.
Sources: EBS and authors’ calculations.

The results for EUR/CHF show a remarkable increase in the variance contribution of bank AIs on the event day. This jumped to over 90% of the total, before the variance shares of all trader types returned close to pre-event levels in the post-event period. The results for USD/CHF show a similar pattern, but the event-day increase in the bank AI share is much smaller and the variance shares of the different trader types do not fully revert to pre-event levels in the post-event period. As bank AIs include both proprietary algorithmic trading by banks (estimated to be about 70% of total bank AI trading) and computer systems that aggregate and transact client orders (estimated to be about 30% of total bank AI trading), either of these could have been responsible for the event-day jumps in variance contributions. To the extent that PTC AIs run similar proprietary trading algorithms to those of banks, the fact we see a different pattern for bank AI and PTC AI suggests that aggregators may have played an important role.

The breakdown in the left-hand panel of Table 4 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. We therefore calculate a per-trade variance impact coefficient, which simply scales variance contributions by the number of trades undertaken by the different trader types.

The right-hand panel of Table 4 shows the results of this analysis. The remarkably high contribution of bank AIs to EUR/CHF volatility on the event day is still present on a per-trade basis, as is an increase in its contribution to USD/CHF volatility. This panel also shows that the contribution to volatility of human trades declined on a per-trade basis for both EUR/CHF and USD/CHF, particularly on the event day itself. We obtain the same qualitative results if we normalise by trade volume rather than number of trades.

As a complement to our per-trade variance results, Table 5 shows estimated price-impact coefficients (Kyle, 1985). These were derived by regressing five-minute returns on the net order flow for each type of trader during the same five-minute periods. Specifically, returns were 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. Net order flow was 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 5: Price impact coefficients⁽¹⁾

	<i>EUR/CHF</i>				<i>USD/CHF</i>		
	Human	Bank AI	PTC AI		Human	Bank AI	PTC AI
Pre-event period	0.001	0.000	0.002 ^{***}	Pre-event period	0.064 ^{**}	0.073 ^{***}	0.034 ^{***}
Event day	-0.102	0.000	0.118	Event day	0.716	1.421 ^{***}	-0.371
Post-event period	0.454 ^{***}	0.480 ^{***}	0.151 ^{***}	Post-event period	0.244 ^{**}	0.516 ^{***}	-0.023

(1) *** / ** / * denote statistical significance at the 1% / 5% / 10% level.

Sources: EBS and authors' calculations.

The results share some similarities with the right-hand panel of Table 4. In particular, bank AI order flow 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, the price impact of bank AI order flow was effectively zero in EUR/CHF on the event day, while bank AI trades generated the most volatility. This suggests that bank AIs contributed a lot of uninformative ‘noise’ to the EUR/CHF market on the event day. We investigate the information contributions of our three trader types in more detail in the next section.

6.2 Contributions to efficient pricing

Since it would be possible to argue that any increase in volatility driven by AI or human trading was valuable if it helped market prices to fully reflect available information, it would be desirable to view the results in Section 6.1 in combination with estimated contributions of different trader types to the formation of efficient prices. We estimate the latter using a variant

of the vector autoregression (VAR) model developed in Hasbrouck (1991(a), 1991(b), 2007) and employed in Hendershott *et al.* (2011).¹⁸ Specifically, we estimate the following model:

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

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

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

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

where r denotes returns on the base currency and x denotes its order flows, *i.e.* net liquidity-consuming purchases, both are calculated over five-minute periods, indexed by t . More specifically, returns are logarithmic returns based on the mid-point of the best bid and ask quotes in the final 100 milliseconds of the current and previous periods, and order flows are computed separately for the different types of trader.

The only structural assumptions in this model relate to timings. Thus, we assume that PTC AIs – as the fastest traders in the market – can adjust their net orders to the contemporary order flow of other market participants. Similarly, bank AIs – as the next fastest trader type – can adjust their net orders to the contemporary order flow of human traders, but not to that of PTC AIs. Finally, human traders can only adjust their net orders to the previous order flows of other market participants.¹⁹ These assumptions are in a similar vein to those of Brogaard *et al.* (2014) in their study of HFT and non-HFT trading activity.

We estimate this model, selecting $K = 5$ as the optimal number of lags, and transform it to a vector moving-average representation by repeatedly substituting for the right-hand side terms.²⁰ The resulting equation for returns is:

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

As suggested by Hendershott *et al.* (2011), the last three terms may be considered ‘private information’ as they reflect order flows from particular trader types, while the first term may be considered ‘public information’. Thus, we can identify separate contributions to efficient pricing from public information and private information pertaining to each of our three trader types via:²¹

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

¹⁸ The model is described in some detail on pages 78-85 of Hasbrouck (2007).

¹⁹ This is the most logical ordering, but even with alternative orderings the pattern of results across periods remains similar.

²⁰ We did this ten times, by when the marginal effect of each substitution had become small.

²¹ To be clear, this follows from the assumptions of the model, and not from information in the market data provided by EBS.

where σ_r^2 is the overall variance of returns, and the terms on the right-hand side respectively represent contributions to this from public information and private information pertaining to PTC AI, bank AI and human traders. $\sigma_{\varepsilon,r}^2$, $\sigma_{\varepsilon,p}^2$, $\sigma_{\varepsilon,B}^2$, $\sigma_{\varepsilon,H}^2$ denote the variances of return shocks and shocks to order flows of each trader type. Table 6 presents the results of this decomposition.

Table 6: Estimated contributions to variance of efficient returns

<i>EUR/CHF</i>					<i>USD/CHF</i>				
<i>Per cent</i>	Returns	Order flow			<i>Per cent</i>	Returns	Order flow		
		Human	Bank AI	PTC AI			Human	Bank AI	PTC AI
Pre-event period	63.8	4.4	0.3	31.5	Pre-event period	79.5	4.3	11.0	5.2
Event day ⁽¹⁾	11.8	69.2	18.1	0.9	Event day ⁽¹⁾	19.2	17.9	26.1	36.9
Post-event period	39.9	27.3	19.4	13.4	Post-event period	85.3	3.1	11.3	0.2

⁽¹⁾ A small number of the most extreme returns on the event day were excluded to avoid these driving the results.
Sources: EBS and authors' calculations.

The results for EUR/CHF show a striking shift in information contributions across our three periods. In the pre-event period, PTC AIs are estimated to have made by far the largest contribution to the variance of efficient returns of all types of order flows, with bank AIs and human traders contributing very little. On the event day, human trading took over as the most significant contributor, while the influence of PTC AIs all but disappeared. Human trading also maintained a significant contribution in the post-event period. The role of bank AIs on the event day and in the post-event period was similar to that of human traders, though not as dramatic, stepping up from pre-event levels. This was dwarfed by the increased contribution of bank AI to total realised volatility highlighted in Section 6.1.

The pattern for USD/CHF is somewhat less clear, as the informational role of all types of trading is estimated to be relatively small in both the pre-event and post-event periods. When the contribution of public information collapsed on the event day, however, order flows did temporarily become much more informative, particularly those from AI trades.

This section suggests that human traders played an important role in the discovery of efficient prices, notably in the EUR/CHF market on the event day. Here, they substituted for computers, notably PTC AIs, which still contributed less to price discovery than human traders in the post-event period. Combining these results with previous ones, showing that AI traders made large contributions to realised volatility, we conclude that these traders added significant noise to FX rates following the SNB announcement. This may reflect the possibility that many computer trades after the announcement were driven by liquidity needs rather than information. Such 'fire-selling' would be consistent computers being net consumers of liquidity, as we saw in Section 5.

7 Arbitrage opportunities and market efficiency

Studies have found that computer trading algorithms sometimes help to iron out market imperfections such as arbitrage opportunities. In this section, we present measures of market efficiency relating to triangular arbitrage to see if there were changes in the efficiency of Swiss franc currency markets as computer traders withdrew liquidity following the SNB announcement on 15 January 2015.

The Swiss franc is one of the few currencies continuously quoted directly against both the euro and the US dollar. Some computer trading in these markets may therefore be engaged in triangular arbitrage. This involves searching for and trading on instances of direct quotes for EUR/CHF that have moved out of line with implied quotes derived from USD/CHF and EUR/USD.

We begin by calculating the frequency and average size of such arbitrage opportunities on the day of the SNB announcement and in the periods preceding and following it. 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 or *vice versa*. The profits have to exceed a *de minimus* one basis point, and we record the average profitability of all arbitrage opportunities meeting this criterion. The results are shown in the left-hand panel of Table 7.

Table 7: Arbitrage opportunities between EUR/CHF, USD/CHF and EUR/USD

<i>Size and frequency of opportunities</i>			<i>Trading on arbitrage opportunities⁽¹⁾</i>		
	Frequency ⁽²⁾ <i>Per cent</i>	Profitability ⁽³⁾ <i>Basis points</i>	<i>Coefficients</i>	Pre-event period	Post-event period
Pre-event period	0.004	9.9	Human	-0.0022	0.0010
Event day	0.897	181.2	Bank AI	0.0061***	0.0025
Post-event period	0.046	4.8	PTC AI	0.0089***	-0.0050

(1) *** / ** / * denotes statistical significance at the 1% / 5% / 10% level.

(2) Percentage of 100 millisecond periods in which the combination of best bid and ask quotes across the three currency-pairs offers a profit in excess of one basis point.

(3) Average profitability of arbitrage opportunities where they exist.

Sources: EBS and authors’ calculations.

Not surprisingly, 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 suggests that algorithmic trading may have become less active in this latter period, possibly in response to the increased volatility of the two CHF rates in the arbitrage triangle.

To investigate more thoroughly how the role of algorithmic trading in arbitrage in the post-event period compares with that of the pre-event period, we estimate a structural vector autoregression (SVAR) model of the relationship between arbitrage opportunities and the

trading volumes of different types of trader. This analysis closely follows that of Chaboud *et al.* (2014). In particular, we estimate the following model:

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

where Y contains four endogenous variables, the first of which measures the frequency of arbitrage opportunities, while the remaining ones measure the order flow of each trader type relative to total market order flow. These variables are measured over five-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 (1993). Two lags of the endogenous variables are also included in the model, as are six exogenous variables, X . These are total trade volumes and return volatilities for each of the three currency pairs in the arbitrage triangle, all computed over the preceding ten minutes. Finally, we include nine time dummy variables, G , one for each hour of the trading day.

The right-hand panel of Table 7 shows the estimated contemporaneous coefficients that reflect how the trading activity of each type of trader responds to arbitrage opportunities. It reveals statistically significant positive responses of human, bank AI and PTC AI trading in the pre-event period, with both bank AI and PTC AI trading responding more than that of human traders. In the post-event period, however, human traders pursued triangular arbitrage opportunities more than bank AIs and PTC AIs, even though the relationship was weaker and no longer statistically significant. Thus, our results suggest a significant reduction in computer resources devoted to triangular arbitrage following the Swiss franc event, with human trading becoming the most important source of arbitrage (though not significantly so). The fact that no type of trading made a significant contribution to arbitrage suggests that most mis-pricings were closed by quote adjustment rather than active trading.

8 Non-CHF foreign exchange rates

Although we have some evidence that AI trading adversely affected market liquidity and price formation in exchange rates featuring the Swiss franc after the SNB announcement, an important question for financial stability is whether these effects spread to exchange rates more widely. If so, to what extent did AI trading undermine the quality of these FX markets? In order shed light on this issue, we now focus on three other currency pairs, EUR/USD, USD/JPY and EUR/JPY. We chose these three cross-rates as EUR/USD and USD/JPY are the two most traded currency pairs in all FX markets, and while EUR/USD is associated with EUR/CHF and USD/CHF through triangular arbitrage, USD/JPY is not. We added EUR/JPY to study arbitrage in another triangle of currencies.

First we repeat the analysis of liquidity provision and consumption of Section 5.1 for these three currency pairs. Here, we focus on net liquidity provision, which is shown in Table 8.

Table 8: Net liquidity provision⁽¹⁾

	EUR/USD			USD/JPY			EUR/JPY		
	Human	Bank AI	PTC AI	Human	Bank AI	PTC AI	Human	Bank AI	PTC AI
<i>Share of trade volume (%)</i>									
Pre-event period	20.9	5.0	-25.9	20.1	7.7	-27.8	35.3	18.5	-53.8
Event day	17.7	-2.1	-15.6	18.2	-2.4	-15.8	31.9	18.4	-50.4
Post-event period	18.5	4.1	-22.6	19.6	7.7	-27.3	36.6	20.4	-57.0
<i>Statistical tests (t-statistics)⁽²⁾</i>									
Event day = pre-event?	-3.3 ^{***}	-6.7 ^{***}	9.2 ^{***}	-1.6	-8.0 ^{***}	7.4 ^{***}	-2.1 [*]	-0.1	1.8
Post-event = pre-event?	-1.3	-0.6	1.4 [*]	-0.3	0.0	0.1	0.3	0.6	-0.8

⁽¹⁾ Share of volume as liquidity provider minus share of volume as liquidity consumer.

⁽²⁾ *** / ** / * denote statistical significance at the 1% / 5% / 10% level.

Sources: EBS and authors' calculations.

Table 9: Effective spreads by type of liquidity provider*EUR/USD*

	Median spread (<i>basis points</i>)			Statistical tests ⁽¹⁾ (<i>p-values</i>)	
	Human	Bank AI	PTC AI	Human = Bank AI?	Human = PTC AI?
Pre-event period	0.17	0.21	0.34	0.00 ^{***}	0.00 ^{***}
Event day	0.22	0.30	0.51	0.00 ^{***}	0.00 ^{***}
Post-event period	0.20	0.25	0.42	0.00 ^{***}	0.00 ^{***}
<i>Statistical tests⁽¹⁾ (χ^2 statistics)</i>					
Event day = pre-event?	42.6 ^{***}	62.3 ^{***}	84.4 ^{***}		
Post-event = pre-event?	51.5 ^{***}	101.8 ^{***}	160.7 ^{***}		

USD/JPY

	Median spread (<i>basis points</i>)			Statistical tests ⁽¹⁾ (<i>p-values</i>)	
	Human	Bank AI	PTC AI	Human = Bank AI?	Human = PTC AI?
Pre-event period	0.14	0.24	0.35	0.00 ^{***}	0.00 ^{***}
Event day	0.18	0.27	0.47	0.00 ^{***}	0.00 ^{***}
Post-event period	0.14	0.26	0.41	0.00 ^{***}	0.00 ^{***}
<i>Statistical tests⁽¹⁾ (χ^2 statistics)</i>					
Event day = pre-event?	18.7 ^{***}	11.8 ^{***}	45.6 ^{***}		
Post-event = pre-event?	0.6	13.2 ^{***}	57.5 ^{***}		

EUR/JPY

	Median spread (<i>basis points</i>)			Statistical tests ⁽¹⁾ (<i>p-values</i>)	
	Human	Bank AI	PTC AI	Human = Bank AI?	Human = PTC AI?
Pre-event period	0.22	0.36	0.45	0.00 ^{***}	0.00 ^{***}
Event day	0.46	0.56	0.92	0.91	0.00 ^{***}
Post-event period	0.28	0.48	0.64	0.00 ^{***}	0.00 ^{***}
<i>Statistical tests⁽¹⁾ (χ^2 statistics)</i>					
Event day = pre-event?	28.3 ^{***}	19.0 ^{***}	73.9 ^{***}		
Post-event = pre-event?	14.8 ^{***}	20.2 ^{***}	60.0 ^{***}		

⁽¹⁾ The tests are described in footnote 18. *** / ** / * denote statistical significance at the 1% / 5% / 10% level.

Sources: EBS and authors' calculations.

The table shows no significant change in the net supply of liquidity in non-CHF currency pairs in the post-event period compared with the pre-event period. On the event day, there was some decline in net liquidity provision in EUR/USD and USD/JPY by bank AIs, but this was compensated for by an increase from PTC AIs.

Second, we look at effective spreads, as in Section 5.2. Table 9 shows how effective spreads varied on and after the day of the SNB announcement by type of liquidity provider for our three non-CHF currency pairs. Across these currency pairs, spreads on human trades were always narrower than spreads on computer trades, but all spreads widened on the event day and reverted towards pre-event values in the post-event period, such that the relative positions of spreads with human and computer traders did not change.

Finally, we repeat our analysis of the frequency and size of arbitrage opportunities in Section 8.1 for the USD/JPY-EUR/JPY-EUR/USD triangle. If the Swiss franc event had a widespread impact on computer trading in all FX cross-rates we might expect more arbitrage opportunities to have appeared in this triangle. However, as Table 10 shows, although there was some increase in arbitrage opportunities on the event day, which was statistically significant, they were almost identical in the pre-event and post-event periods.

Table 10: Arbitrage opportunities between USD/JPY, EUR/JPY and EUR/USD

	Frequency ⁽¹⁾ <i>Per cent</i>	Profitability ⁽²⁾ <i>Basis points</i>
Pre-event period	0.012	0.2
Event day	0.494	1.9
Post-event period	0.014	0.2

⁽¹⁾ Percentage of 100 millisecond periods in which the combination of best bid and ask quotes across the three currency-pairs offers a profit in excess of one basis point.

⁽²⁾ Average profitability of arbitrage opportunities where they exist.

Sources: EBS and authors' calculations.

9 Conclusion

The Swiss franc event is probably the most significant shock to FX markets since computerised algorithmic trading has been prominent. Studying the reaction to this shock, we find that algorithmic trading contributed to the decline of EUR/CHF and USD/CHF market quality on the event day and afterwards as they withdrew liquidity and generated uninformative volatility. Human traders took over as the main contributors to efficient pricing, while algorithms tended to amplify price movements by following trends. Trades by bank algorithms, in particular, contributed substantially to non-informative EUR/CHF volatility. This may indicate that the role of trade aggregators, which has been highlighted in other extreme market events (see for example BIS, 2017), was important. Both PTC and bank algorithms, which had traded on

triangular arbitrage opportunities before the event, ceased doing so afterwards. However, adverse effects on non-CHF currency pairs were limited, suggesting that algorithmic trading did not propagate contagion in the FX market, at least on this occasion.

Of course, it is hard to draw general conclusions from one event, not least because we have only studied *how* algorithmic trading reacted and not *why* it did so. Was it due to more stringent capital or trading requirements applied to algorithmic trading or because of the behaviour of the algorithms themselves? If the latter, were there types of algorithm that reacted more and types that reacted less or even partially offset the behaviour of the others? If that were the case, the mix of algorithms operating at the time of future financial market shocks could affect the scale of any amplification that occurs. Indeed, that mix could be affected by the adaptation of algorithms as they experience periods of market stress like the one studied in this paper. Nevertheless, our results contrast with evidence that algorithmic trading in aggregate improves liquidity and price discovery in normal times. This suggests there is some value in maintaining a diversity of trader types to help keep markets resilient through different trading conditions.

10 References

- Afonso, G., Kovner, A. and Schoar, A. (2011), ‘Stressed, not frozen: the federal funds market in the financial crisis’, *Journal of Finance*, vol. 66, pp. 1109–1139.
- Ait-Sahalia, Y. and Saglam, M. (2014), ‘High-frequency traders: taking advantage of speed’, *working paper*.
- Anand, A., and Venkataraman, K. (2014), ‘Should exchanges impose market maker obligations?’, *working paper*.
- Baldauf, M. and Mollner, J. (2015), ‘High-frequency trading and market performance’, *working paper*.
- Bernales, A. (2014), ‘How fast can you trade? High-frequency trading in dynamic limit order markets’, *Banque de France working paper*.
- Biais, B. and Foucault, T. (2014), ‘HFT and market quality. Bankers, markets and investors’, *manuscript*.
- Biais, B., Foucault, T. and Moinas, S. (2015), ‘Equilibrium fast trading’, *Journal of Financial Economics*, vol. 116, pp. 292–313.
- Biais B. and Woolley, p. (2011), ‘High-frequency trading’, *Toulouse University manuscript*.
- BIS (2017), ‘The sterling ‘flash event’ of 7 October 2016’, *Bank for International Settlements Market Committee report*.
- Bloomberg (2015), ‘SNB unexpectedly gives up cap on Franc’, 15 January.
- Boehmer, E., Fong, K. and Wu, J. (2015), ‘International evidence on algorithmic trading’, *working paper*.
- Boehmer, E., Li, D. and Saar, G. (2017), ‘The competitive landscape of high-frequency trading firms’, *working paper*.
- Bongaerts, D. and 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. and Sokolov, K. (2017), ‘High-frequency trading and extreme price movements’, *working paper*.
- Brogaard, J., Hagstromer, B., Norden, L. and Riordan, R. (2015), ‘Trading fast and slow: co-location and market quality’, *Review of Financial Studies*, vol. 28, pp. 3407–3443.
- Brogaard, J., Hendershott, T. and Riordan, R. (2014), ‘High-frequency trading and price discovery’, *Review of Financial Studies*, vol. 27, pp. 2267–2306.

- Budish, E., Cramton, P. and Shim, J. (2015), 'The high-frequency trading arms race: frequent batch auctions as a market design response', *Quarterly Journal of Economics*, vol. 130, pp. 1547–1621.
- Carrion, A. (2017), 'Very fast money: high-frequency trading on NASDAQ', *Journal of Financial Markets*, forthcoming.
- Cespa, G. and Vives, X. (2015), 'The beauty contest and short-term trading', *Journal of Finance*, vol. 70, pp. 2099-2154.
- Chaboud, A., Chiquoine, B., Hjalmarsson, E. and Vega, C. (2014), 'Rise of the machines: algorithmic trading in the foreign exchange market', *Journal of Finance*, vol. 69, pp. 2045-2084.
- CFTC-SEC Staff Report (2010), 'Findings regarding the market events of May 6, 2010'.
- Chordia T., Goyal, A., Lehmann B. and Saar G. (2013), 'High-frequency trading', *Journal of Financial Markets*, vol. 16, pp. 637-645.
- Conrad, J., Wahal S. and Xiang, J. (2015), 'High-frequency quoting, trading, and the efficiency of prices', *Journal of Financial Economics*, vol. 116, pp. 271-291.
- Du, S. and Zhu, H. (2015), 'Welfare and optimal trading frequency in dynamic double auctions', *National Bureau of Economic Research working paper*, no. 20588.
- Easley, D., Lopez de Prado, M. and O'Hara, M. (2012), 'Flow toxicity and liquidity in a high-frequency world', *Review of Financial Studies*, vol. 25, 1457–1493.
- Easley, D., Lopez de Prado, M. and O'Hara, M. (2013), 'High-frequency trading – new realities for traders, markets and regulators', *Risk Books*, London.
- Fischer A. (2004), 'Price clustering in the FX market: a disaggregated analysis using central bank interventions', *Centre for Economic Policy and Research discussion paper*, no. 4529.
- Foucault, T., Hombert, J. and Rosu, I. (2016), 'News trading and speed', *Journal of Finance*, vol. 71, 335-382.
- Goldstein, M., Kumar, P. and Graves, F. (2014), 'Computerized and high-frequency trading', *The Financial Review*, vol. 49, 177-202.
- Gomber, P., Arndt, B., Lutat, M. and Uhle, T. (2011), 'High-frequency trading', *Goethe University working paper*.
- Hagströmer, B. and Menkveld, A. (2016) 'A network map of information percolation', *working paper*.
- Harris, L. (2012) 'Stop the high-frequency trader arms race', *Financial Times*, December 27.
- Hasbrouck, J. (1991a), 'Measuring the information content of stock trades', *Journal of Finance*, vol. 46, pp. 179–207.

- Hasbrouck, J. (1991b), 'The summary informativeness of stock trades: an econometric analysis', *Review of Financial Studies*, vol. 4, 571–595.
- Hasbrouck, J. (2007), 'Empirical market microstructure', *Oxford University Press*, New York.
- Hasbrouck, J. and Saar, G. (2013), 'Low-latency trading', *Journal of Financial Markets*, vol. 16, pp. 646-679.
- Hendershott, T., Jones, C. and Menkveld, A. (2011), 'Does algorithmic trading improve liquidity?' *Journal of Finance*, vol. 66, 1-33.
- Hendershott, T. and Menkveld, A. (2014), 'Price pressures', *Journal of Financial Economics*, vol. 114, 405-423.
- Hoffmann, P. (2014), 'A dynamic limit order market with fast and slow traders', *Journal of Financial Economics*, vol. 113, pp. 156-169.
- Jones, C. (2013), 'What do we know about high-frequency trading?', *Columbia University manuscript*.
- Jovanovic, B. and Menkveld, A. (2016), 'Middlemen in limit-order markets', *working paper*.
- Kirilenko, A., Kyle, A., Samadi, M. and Tuzun, T. (2017), 'The flash crash: the impact of high-frequency trading on an electronic market', *Journal of Finance*, vol. 72, 967-998.
- Kirilenko, A. and Lo, A. (2013) 'Moore's Law versus Murphy's Law: algorithmic trading and its discontents', *Journal of Economic Perspectives*, vol. 27, pp. 51-72.
- Korajczyk, R. and Murphy D. (2016) 'High-frequency market making to large institutional trades', *working paper*.
- Jordan, T. (2014), 'Introductory remarks by Thomas Jordan', speech available [here](#).
- King, M., Osler, C. and Rime, D. (2011), 'Foreign exchange market structure, players and evolution', *Norges Bank working paper*.
- Kyle, A. (1985), 'Continuous auctions and insider trading', *Econometrica*, vol. 53, pp. 1315–1336.
- Lo, A. and MacKinlay, A. (1988) 'Stock market prices do not follow random walks: evidence from a simple specification test', *Review of Financial Studies*, vol. 1, pp. 41–66.
- Malinova, K., Park, A. and Riordan, R. (2013), 'Do retail traders suffer from high-frequency traders?', *working paper*.
- McCann, C. and Yan, M. (2015), 'The recent market turmoil spells trouble for "auto-liquidators" like interactive brokers', Securities Litigation & Consulting Group technical report, available [here](#).
- Menkveld, A. (2013), 'High-frequency trading and the new market makers', *Journal of Financial Markets*, vol. 16, pp. 712-740.

- Menkveld, A. (2016), ‘The economics of high-frequency trading: taking stock’, *Annual Review of Financial Economics*, vol. 8, pp. 1-24.
- Menkveld, A. (2017), ‘High-frequency trading as viewed through an electronic microscope’, *Financial Analysts Journal*, forthcoming.
- Menkveld, A. and Yueshen, B. (2015), ‘The flash crash: a cautionary tale about highly fragmented markets’, *working paper*.
- Mirkov, K., Pozdeev, I. and Söderlind, P. (2016), ‘Toward removal of the Swiss franc cap: market expectations and verbal interventions’, *Swiss National Bank working paper*, no. 10/16.
- Moser, D. (2016), ‘On the pulse of the financial markets’, speech at Geneva Money Market Event, available [here](#).
- Pagnotta, E. and Philippon, T. (2015), ‘Competing on speed’, *working paper*.
- O’Hara, M. (2015), ‘High-frequency market microstructure’, *Journal of Financial Economics*, vol. 116, pp. 257–270.
- O’Hara, M. and Ye, M. (2011), ‘Is market fragmentation harming market quality?’, *Journal of Financial Economics*, vol. 100, pp. 459–474.
- Raman, V., Robe, M. and Yadav, P. (2014), ‘Electronic market makers, trader anonymity and market fragility’, *Chicago Futures Trading Commission manuscript*.
- 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’, *Review of Economics and Statistics*, vol. 85, pp. 777-792.
- Rojcek, J. and Ziegler, A. (2016), ‘High-frequency trading in limit order markets: equilibrium impact and regulation’, *working paper*.
- Rosu, I. (2016), ‘Fast and slow informed trading’, *HEC (Paris) manuscript*.
- Securities Exchange Commission (2010), ‘Concept release on equity market structure’, release no. 34-61358, file no. S7-02-10.
- Swiss National Bank (2011), ‘Swiss National Bank sets minimum exchange rate at CHF 1.20 per euro’, available [here](#).
- Swiss National Bank (2015), ‘SNB discontinues minimum exchange rate and lowers interest to –0.75%’, available [here](#).
- Tong, L. (2015), ‘A blessing or a curse? The impact of high-frequency trading on institutional investors’, *working paper*.
- van Kervel, V. and Menkveld, A. (2016), ‘High-frequency trading around large institutional orders’, *working paper*.