Toxic Arbitrage*

Thierry Foucault† Roman Kozhan‡ Wing Wah Tham§

February, 2015

Abstract

Short lived arbitrage opportunities can arise when the prices of asset pairs do not adjust to information at the same speed. These opportunities are toxic because they expose investors to the risk of trading with arbitrageurs at stale quotes. Hence, more frequent toxic arbitrage opportunities and a faster arbitrageurs' response to these opportunities can impair liquidity. We provide supporting evidence using data on triangular arbitrage in currency markets. In our sample, a 1% increase in the likelihood that a toxic arbitrage terminates with an arbitrageur's trade (rather than a quote update) is associated with a 4% increase in bid-ask spreads. Our findings suggest that fast arbitrageurs' response to toxic arbitrage opportunities enhances pricing efficiency while raising trading costs for other market participants.

Keywords: Arbitrage; Adverse Selection; Liquidity; High Frequency Trading.

JEL Classification: D50, F31, G10

*We thank Mark Van Achter, Hank Bessembinder, Geir Bjønnes, Michael Brennan, Alain Chaboud, Pierre Collin-Dufresnes, Jean-Edouard Colliard, Matthijs Fleischer, Arie Gozluklu, Terry Hendershott, Frank de Jong, Andrew Karolyi, Pete Kyle, Grace Xing Hu, Olga Lebedeva, Bruce Lehmann, Katya Malinova, Albert Menkveld, Michael Moore, Pamela Moulton, Maureen O'Hara, Marco Pagano, Andreas Park, Angelo Ranaldo, Vikas Raman, Dagfinn Rime, Fabrice Riva, Gideon Saar, Daniel Schmidt, Elvira Sojli, Clara Vega, Kumar Venkataraman, Chen Yao, and Mao Ye. We are also grateful to seminar and conference participants at Bath University, Cornell University, EIEF and Consob, the Norwegian School of Business, Vrije University, Manchester University, KU Leuven, the Autorité des Marchés Financiers, the 2015 American Finance Association meetings, the 9^{th} Central Bank Workshop on the Market Microstructure of Financial Markets, the 2013 BIRS workshop on modeling high frequency trading activity, the 6^{th} Erasmus Liquidity Conference, the Conference on Liquidity and Arbitrage Trading in Geneva, and the 2014 CityU Finance Conference in Hong Kong.

†Department of Finance, HEC, Paris, 1 rue de la Libération, 78351 Jouy en Josas; tel: +33 13967 9569; e-mail: foucault@hec.fr

‡Warwick Business School, University of Warwick, Scarman Road, Coventry, CV4 7AL, UK; tel: +44 24 7652 2114; e-mail: Roman.Kozhan@wbs.ac.uk

§Econometric Institute, Erasmus School of Economics, Erasmus University Rotterdam, Burg. Oudlaan 50, PO Box 1738, 3000DR, Rotterdam, the Netherlands, tel: +31 10408 1424; e-mail: tham@ese.eur.nl

1 Introduction

Arbitrageurs play a central role in financial markets. When the Law of One Price (LOP) breaks down, they step in, buying the cheap asset and selling the expensive one. Thereby, arbitrageurs enforce the LOP and make markets more price efficient. In theory, arbitrage opportunities should disappear instantaneously. In reality, they do not because arbitrage is not frictionless. As Duffie (2010) points out: *"The arrival of new capital to an investment opportunity can be delayed by fractions of a second in some markets, for example an electronic limit order-book market for equities, or by months in other markets [...]."*

Various frictions (e.g., short-selling costs, funding constraints, idiosyncratic risks, etc.) explain why some arbitrage opportunities persist (see Gromb and Vayanos (2010)). For very short-lived arbitrage opportunities–those lasting *fractions of a second*– attention costs and technological constraints are the main impediments to a seamlessly Law of One Price. These barriers are falling as high frequency arbitrageurs invest massively to detect and react ever faster to arbitrage opportunities. Returns on high speed arbitrage are substantial because arbitrage opportunities are very frequent at time scale of milliseconds, due to market fragmentation and the proliferation of derivatives assets (ETFs, futures, etc.). For instance, a report from the Tabb group estimates high frequency arbitrageurs' profits at 21 billion for year 2009 alone (see Sussman et al. (2009)). This evolution has triggered debates about the social value of high speed arbitrage, and in particular about whether arbitrage strategies "*benefit or harm the interests of long-term investors and market quality [...]*" (U.S. Securities and Exchange Commission (2010), Section B, p.51).

Arbitrageurs can be beneficial or harmful for other investors, depending on the cause of arbitrage opportunities. When these opportunities are due to transient demand or supply shocks ("price pressures"), arbitrageurs implicitly act as liquidity providers in exploiting them (see, for instance, Holden (1995), Gromb and Vayanos (2002), and Gromb and Vayanos (2010)). In this case, trades between arbitrageurs and their counterparties are mutually beneficial.[1] However, short lived arbitrage opportunities are also due to asynchronous adjustments in asset prices following information arrival. Arbitrageurs' profits in these trades are obtained at the expense of liquidity suppliers with stale quotes.[2] Thus, asynchronous price adjustments to information

[1] For instance, Gromb and Vayanos (2002) write (on p.362): *"In our model, arbitrage activity benefits all investors. This is because through their trading, arbitrageurs bring prices closer to fundamentals and supply liquidity to the market."*

[2] This problem is not new. In the 90s, professional day traders (so-called SOES bandits) were picking off

in asset pairs generate "toxic" arbitrage opportunities, i.e., opportunities in which liquidity suppliers are at risk of being adversely selected.[3] Through this channel, high speed arbitrageurs can harm market liquidity because liquidity suppliers charge larger bid-ask spreads to cover the risk of trading at stale quotes (Copeland and Galai (1983)).

Our contribution is to provide evidence about this channel. To our knowledge, our paper is first to do so. This is important for at least two reasons. First, arbitrage is a central notion in finance. Thus, understanding how it affects liquidity is of broad interest. Second, recent proposals advocate slowing down the pace of trading precisely on the grounds that high speed arbitrageurs raise liquidity suppliers' risk of trading at stale quotes (see, e.g., Budish et al. (2014)). However, there is yet no evidence on the sensitivity of illiquidity to this risk. Measuring this sensitivity is not straightforward because it is not the level of arbitrage activity per se that should affect illiquidity. Rather, as shown in this paper, illiquidity should depend on the *composition* of arbitrage opportunities and arbitrageurs *relative speed* of reaction to these opportunities. Specifically, illiquidity should be higher when (i) the fraction of arbitrage opportunities that are toxic or (ii) the likelihood that a toxic arbitrage opportunity terminates with an arbitrageur's trade are higher.

We first check that these two predictions are well-grounded using a model of cross-market arbitrage. In the model, arbitrage opportunities can be either toxic (due to asynchronous price adjustments to news) or non-toxic (due to liquidity shocks). As in the data, an arbitrage opportunity terminates either with an arbitrageur's trade or a market maker's quote update, depending on whoever observes the opportunity first. We solve for equilibrium bid-ask spreads in each asset and traders' optimal speed of reaction to arbitrage opportunities. Thus, in equilibrium, illiquidity and the duration of arbitrage opportunities (a measure of pricing efficiency) are jointly determined.

The model generates predictions (i) and (ii) (see above) and two additional ones about the durations of arbitrage opportunities. First, when the arbitrage mix becomes more toxic, arbitrage opportunities should be shorter, even though bid-ask spread costs borne by arbitrageurs should be higher. The reason is that market makers have more incentive to react fast to arbitrage opportunities (by updating their quotes) when they expect more of them to be toxic. This

Nasdaq dealers with stale quotes by using Nasdaq's Small Order Execution System (a system that guaranteed automatic execution of market orders up to a certain size at Nasdaq dealers quotes). See Harris and Schultz (1997) and Foucault et al. (2003).

[3] Our definition of a toxic trade follows Easley et al. (2012). They write (p.1458): *"Order flow is regarded as toxic when it adversely selects market makers who may be unaware that they are providing liquidity at a loss."*

effect also induces arbitrageurs to be faster so that overall arbitrage opportunities are shorter. Second, by a similar logic, a technological change allowing traders to be faster should reduce the duration of arbitrage opportunities, even though it may increase illiquidity by raising the odds that arbitrageurs pick off dealers' stale quotes.

We then test these predictions using data on triangular arbitrage opportunities for three currency pairs (dollar-euro, dollar-pound, and pound-euro).[4] Although our predictions and methodology apply to any type of high frequency arbitrage opportunities, we focus on triangular arbitrage opportunities for a couple of reasons.

The first one is practical. For our tests, we must accurately measure when an arbitrage begins, when it terminates, how it terminates (with a trade or a quote update), and we must track prices after the arbitrage terminates (to identify toxic arbitrage opportunities; see below). This requires data on pairs of related assets (not just one asset) with a level of precision that is not easily available to researchers. Our data has the required granularity: for three currency rates, we observe all orders and trades from January 2003 to December 2004 in Reuters D-3000 (one of the two interdealer trading platforms used by foreign exchange dealing banks at the time of our sample) with a time stamp accuracy of 10 milliseconds.[5] Moreover, empirical studies of cross-market arbitrage are often plagued with asynchronicities in price reporting for assets traded in different venues. We avoid this problem because all data in our study are reported by the same trading platform.

Second, strategies exploiting triangular arbitrage opportunities are not hindered by taxes, short-selling constraints, or funding constraints and the risk of these strategies is very limited. Hence, standard limits to arbitrage cannot explain why triangular arbitrage opportunities are not eliminated immediately (see Pasquariello (2014)).[6] The most likely explanation is that, as in our model, technological constraints limit the speed at which traders react to arbitrage opportunities. Thus, triangular arbitrage opportunities are very similar to other high speed opportunities: they are (i) frequent (we observe more than 37,000 in our sample), (ii) very short-lived (they last less than one second on average), (iii) more efficiently exploited by machines

[4] One can buy dollars with euros in two ways: (i) directly by trading in the dollar-euro market or (ii) indirectly by first buying pounds with euros and then dollars with pounds. If the price (in euros) of these two strategies differs then a triangular arbitrage opportunity exists.

[5] Kozhan and Tham (2012) use the same data to measure the profitability of triangular arbitrage opportunities.

[6] Pasquariello (2014) finds that none of the usual proxies for limits to arbitrage explain the size of triangular arbitrage opportunities in his sample (see Table 2 in Pasquariello (2014)). For other relatively short lived opportunities, these constraints can be more important. For instance, Gagnon and Karolyi (2010) find that holding costs (e.g., idiosyncratic risk) explain the size of arbitrage opportunities between home and U.S. stock prices for stocks cross-listed in the U.S.

than by humans, and (iv) they deliver razor blade profits per opportunity (1 to 2 basis points in our sample).[7]

As any other arbitrage opportunities, triangular arbitrage opportunities arise for two reasons: (i) asynchronous price adjustments of different currency pairs to new information or (ii) price pressures effects. By definition, price pressure effects are followed by reversals, while asynchronous price adjustments are eventually followed by permanent shifts in exchange rates.[8] Thus, we use price patterns following the occurrence of arbitrage opportunities to sort them into two groups: toxic (followed by permanent changes in exchange rates) and non-toxic (followed by price reversals). With this approach, we obtain 15,908 toxic arbitrage opportunities (about 32 per day), i.e., about 41% of all arbitrage opportunities in the sample. Moreover, we find that about 75% of toxic arbitrage opportunities terminate with an arbitrageur's trade.

As predicted we find a positive and significant relationship between the fraction of arbitrage opportunities that are toxic and illiquidity. Specifically, on days in which this fraction is higher, illiquidity is higher (after controlling for standard determinants of illiquidity). For instance, a one standard deviation increase in the fraction of arbitrage opportunities that are toxic is associated with a 3% increase in average effective spreads for the currencies in our sample. In other words, the arbitrage mix matters: illiquidity is higher when arbitrage opportunities are more frequently due to asynchronous price adjustments than price pressures.

Our second prediction is that an increase in the likelihood that a toxic arbitrage opportunity terminates with an arbitrageur's trade should raise illiquidity. Our theory shows that this likelihood is jointly determined with illiquidity. To account for this endogeneity, we use an instrumental variable approach. Until July 2003, traders had to manually submit their orders to Reuters D-3000. In July 2003, Reuters introduced the "AutoQuote API" functionality (API means "Application Programming Interface"). Traders using this functionality can directly feed

[7] Arbitrage opportunities in the foreign exchange market (either violations of covered interest parity or triangular arbitrage) are well documented. See, for instance, Akram et al. (2008), Fong et al. (2008), Fenn et al. (2009), Mancini-Griffoli and Ranaldo (2011), Marshall et al. (2008), Kozhan and Tham (2012), Ito et al. (2013), Chaboud et al. (2014), and Pasquariello (2014).

[8] Suppose that euro/dollar dealers receive information that calls for an appreciation of the euro and raise their bid and ask quotes (expressed in dollars per euro). If this appreciation is large enough and dealers in, say, the dollar/pound market are slow to adjust their quotes, there is a triangular arbitrage opportunity: one can indirectly buy dollars with euros at a price less than the current bid price in the dollar/euro market. The arbitrage vanishes when dealers in the dollar/pound market adjust their quotes or arbitrageurs hit stale quotes in this market. In either case, there is no reversal of the initial euro/dollar rate. Alternatively, if euro/dollar dealers temporarily accumulate a large short position in euro, they will mark up the value of the euro against the dollar to attract sellers of euros and reduce their risk exposure. If this price pressure effect is large enough, a triangular arbitrage opportunity arises. However, in this case, as dealers' short position decreases, their quotes will revert (see Ho and Stoll (1981), Grossman and Miller (1988), or Hendershott and Menkveld (2014)).

their algorithms to Reuters D-3000 data and let these algorithms submit orders automatically. AutoQuote API marked therefore the onset of algorithmic trading on Reuters D-3000, allowing traders to react faster to triangular arbitrage opportunities.[9] Thus, we instrument arbitrageurs' relative speed with AutoQuote API.

The first stage of the IV regression shows that the likelihood that a toxic arbitrage opportunity terminates with a trade (rather than a quote update) increases by about 4% following the introduction of AutoQuote API. Thus, Autoquote API raised arbitrageurs' relative speed of reaction to arbitrage opportunities, consistent with anecdotal evidence that algorithmic trading in foreign exchange markets was initially used for exploiting triangular arbitrage opportunities (see Chaboud et al. (2014)).

More importantly for our purpose, the second stage shows that the likelihood that, as predicted, a toxic arbitrage opportunity terminates with a trade has a positive effect on illiquidity. Specifically, a 1% increase in this likelihood is associated with a 0.08 basis points increase in quoted bid-ask spreads (3 to 6% of the average bid-ask spread depending on the currency pair). The economic size of this effect is significant given the trading volume for the currency pairs in our sample (we estimate that a 0.08 basis points increase in quoted spread raises the total cost of trading for the currency pairs in our sample by about $157,000 per day). We find similar effects when we measure illiquidity with effective spreads or the slope of limit order books.

Thus, consistent with our predictions, (i) the fraction of arbitrage opportunities that are toxic and (ii) the frequency with which these opportunities are effectively exploited by arbitrageurs are positively associated with illiquidity. Additional predictions of our model are supported by the data as well: (a) the duration of arbitrage opportunities is negatively related to the daily fraction of arbitrage opportunities that are toxic and (b) the introduction of AutoQuote API (an increase in arbitrageurs' relative speed) coincides with a 6.5% (about 60 milliseconds) decrease in the average duration of arbitrage opportunities in our sample. In a similar vein, Chaboud et al. (2014) find that algorithmic trading is associated with less frequent triangular arbitrage opportunities.

It is well known that liquidity facilitates arbitrage (see Holden et al. (2014), Section 5.2). The reverse relation, i.e., the effect of arbitrageurs on liquidity, which is our focus here, has

[9] Hendershott et al. (2011) use the implementation of the NYSE "autoquote" software in 2003 as an instrument for the level of algorithmic trading in NYSE stocks. The NYSE autoquote functionality is different from Reuters AutoQuote API since the former automates the dissemination of updates in best quotes for NYSE stocks while the latter automates order entry. Automation of order entry clearly accelerates the speed at which traders can react to market events.

received much less attention.[10] Kumar and Seppi (1994) is an exception. They model cross-market arbitrageurs as informed traders and show that, in their model, the effect of the number of arbitrageurs on liquidity is non-monotonic. They do not study how the composition of arbitrage opportunities affect illiquidity (all arbitrage opportunities are due to stale quotes in their model).

Several papers argue that fast traders raise adverse selection costs for slower traders.[11] Our empirical finding that illiquidity is higher on days in which toxic arbitrage opportunities terminate more frequently with an arbitrageur's trade is consistent with this view. Our contribution, however, is **not** to test whether adverse selection in general is a source of illiquidity: this is well known. What is new is to show that high speed *arbitrage* is a source of adverse selection and that, for this reason, the relative likelihood of toxic and non-toxic arbitrage opportunities (the "arbitrage mix") in an asset pair is a determinant of its liquidity. To our knowledge, these empirical findings about arbitrage are novel. They contribute to the burgeoning literature on short lived arbitrage opportunities in various assets, such as currencies, ETFs, cross-listed stocks, etc. (see, for instance, Akram et al. (2008), Ben-David et al. (2012), Gagnon and Karolyi (2010), or Shive and Schultz (2010)).

The next section derives our testable hypotheses. Section 3 describes the data, explain how we classify arbitrage opportunities in toxic and non-toxic arbitrage opportunities and present the main empirical findings of the paper. Additional results are presented in Section 4. We conclude in Section 5.

2 Hypotheses Development

In this section, we present the model that guides our empirical analysis. The model builds upon Foucault et al. (2003) but it allows for cross market arbitrage (Foucault et al. (2003) consider reactions to public information arrival in a single asset).

[10] Roll et al. (2007) show that there exist two-way relations between index futures basis and stock market liquidity. In particular, a greater index futures basis Granger-causes greater stock market illiquidity. Roll et al. (2007) argue that this effect could be due to arbitrageurs's trades but do not specifically show that these trades explain the relation. Rosch (2014) uses the size of arbitrage opportunities in Depositary Receipts as an inverse proxy for arbitrage activity and finds a positive association between arbitrage activity and liquidity.

[11] See, for instance, Biais et al. (2014), Budish et al. (2014), Foucault et al. (2003), Foucault et al. (2014), Garvey and Wu (2010), Hendershott and Moulton (2011), Hoffmann (2014), Jovanovic and Menkveld (2012), Menkveld and Zoican (2014).

2.1 A model of cross-market arbitrage

The model has two risky assets, labeled X and Y, three dates, $t \in \{0, 1, 2\}$, two market makers (called X and Y), and one arbitrageur. Market maker $j \in \{X, Y\}$ is specialized in asset j. At $t = 2$, the payoffs of the assets are realized. The payoff of asset Y is $\tilde{\theta}_Y = \tilde{v}_Y + \varepsilon$ where $\varepsilon = 1/2$ or $\varepsilon = -1/2$ with equal probabilities and ε and $\tilde{v}_Y$ are independent. The payoff of asset X is $\tilde{\theta}_X = \sigma \times \tilde{\theta}_Y$. Thus, the payoffs of the two assets are perfectly correlated and any news about the payoff of one asset is relevant for the other.

Market Makers. At $t = 0$, with probability $\alpha \times \varphi$, market maker Y privately observes ε (news about the payoff of asset Y) and his valuation for asset Y becomes $m_Y = E(\tilde{\theta}_Y \mid \varepsilon) = E(\tilde{v}_Y) + \varepsilon$. With probability $\alpha \times (1 - \varphi)$, market maker Y receives no news but experiences a liquidity shock δ that affects his utility from holding the asset. Specifically, market maker Y derives a utility of $(\tilde{\theta}_Y + \delta)$ per share of asset Y owned at date 2 where δ equals 1/2 or $-1/2$ with equal probabilities. Market-maker Y' valuation for the asset is then $m_Y = E(\tilde{\theta}_Y) + \delta = E(\tilde{v_Y}) + \delta$. As in Duffie et al. (2005), this liquidity shock can be interpreted as a need for cash (in which case $\delta < 0$) or an hedging need due to risk management concerns. Finally with probability $(1 - \alpha)$, market maker Y receives no news and experiences no liquidity shock. In this case, his valuation for the asset is $m_Y = E(\tilde{\theta}_Y)$. Market maker X receives no news and no liquidity shock. Thus, his valuation for asset X is $m_X = E(\tilde{\theta}_X)$ at $t = 0$.

At $t = 1$, market makers X and Y simultaneously post an ask price, a_j, and a bid price, b_j, for $j \in \{X, Y\}$ such that:

$$a_j = m_j + \frac{S_j}{2}, \quad and \quad b_j = m_j - \frac{S_j}{2}, \tag{1}$$

where m_j is market maker j's valuation for the asset. Thus, S_j is the bid-ask spread for asset j. In asset X, quotes are for $Q_X = 1$ share and in asset Y for $Q_Y = \sigma$ shares.

Arbitrage Opportunities. Suppose that $S_X + \sigma S_Y < \sigma$ (this condition will be satisfied in equilibrium; see below). If there is a shock to market maker Y's valuation (either a liquidity shock or the arrival of news) at $t = 0$ then there is an arbitrage opportunity. For instance, suppose that this shock is positive. The arbitrageur can then buy one share of asset X at price $a_X = E(\tilde{\theta}_X) + \frac{S_X}{2}$ and sell σ shares of asset Y at price $b_Y = E(\tilde{v_Y}) + 1/2 - S_Y/2$. As the payoff of this portfolio is zero with certainty, the arbitrageur locks in a sure gain of:

$$ArbProfit = \sigma \times b_Y - a_X = (\sigma - S_X - \sigma S_Y)/2, \tag{2}$$

where the second equality follows from the fact that $E(\tilde{\theta}_X) = \sigma E(\tilde{\theta}_Y) = \sigma E(\tilde{v}_Y)$. By symmetry, the arbitrageur's profit is identical if market maker Y's valuation shock is negative. If there is no shock to market maker Y's valuation for the asset then the arbitrageur cannot trade profitably trade (the arbitrage portfolio costs $(S_X + \sigma S_Y)/2$ and has a zero payoff).

Market maker X's expected profit on a trade with the arbitrageur depends on the nature of market maker Y's valuation shock. Consider again the previous example. If the shock to market maker Y's valuation is due to news then market maker X will lose money if he trades with the arbitrageur. Indeed, conditional on positive news for asset Y, asset X's expected payoff is $E(\tilde{\theta}_X \mid \varepsilon = 1/2) = \sigma E(\tilde{\theta}_Y \mid \varepsilon = 1/2) = \sigma(E(\tilde{v}_Y) + 1/2) = E(\tilde{\theta}_X) + \sigma/2$ and the market maker sells the asset at $a_X = E(\tilde{\theta}_X) + \frac{S_X}{2}$. Thus, the market maker's expected profit is $(S_X - \sigma)/2$, which must be negative if the arbitrageur finds it profitable to trade (i.e., if $S_X + \sigma S_Y < \sigma$). If instead, the shock to market maker Y's valuation is due to a liquidity shock then market maker X earns a profit equal to half his bid-ask spread when he trades with the arbitrageur.

In sum, with probability α, a jump in market maker Y's valuation causes an arbitrage opportunity. With probability φ, this jump is due to the arrival of news about asset Y. The arbitrage opportunity is then *toxic* for market maker X since he is exposed to the risk of trading at stale quotes with the arbitrageur. With probability $(1 - \varphi)$, the arbitrage opportunity is due to a liquidity shock private to market maker Y. This type of arbitrage opportunity is *non-toxic* since it is not due to a change in the fair values of assets Y and X. It captures arbitrage opportunities due to transient price pressures (see the introduction). Thus, parameter φ characterizes the composition of arbitrage opportunities (the "arbitrage mix").

With probability $(1 - \alpha)$, there is no shock to market maker Y's valuation and no arbitrage opportunity. In this case, we assume that a liquidity trader arrives in the market to buy or sell Q_X or Q_Y shares of asset X or Y, with equal probabilities. Thus, parameter α should be interpreted as the number of arbitrage opportunities per trade.

Speed. When an arbitrage opportunity occurs, we assume that it takes D^a and D^m units of time for the arbitrageur and market maker X, respectively, to observe the arbitrage opportunity (superscripts m and a refer to market maker X and the arbitrageur, respectively). These reaction times are exponentially distributed with parameter γ (for the arbitrageur) and λ (for market maker X). Hence, the market maker's and the arbitrageur's average reaction times are $1/\lambda$ and $1/\gamma$, respectively. The higher are λ and γ, the faster the traders react to events. We therefore refer to λ and γ as traders' speeds.

If $D^a < D^m$, the arbitrageur detects first the arbitrage opportunity and she exploits it. Otherwise market maker X is first to observe the opportunity. If this opportunity is toxic, market maker X cancels his quote to avoid trading at a loss with the arbitrageur (in reality, market makers will then resubmit new quotes around their updated valuation). If instead the arbitrage opportunity is non toxic, market maker X does nothing since his quotes are not stale relative to the fair value of the asset. Hence, in this case, the arbitrageur trades with probability one. The trading round terminates either with a trade by the arbitrageur or a cancellation of his quotes by market maker X.

Conditional on the arrival of news in asset Y, the likelihood, denoted π, that an arbitrage opportunity terminates with a trade by the arbitrageur is:

$$\pi \equiv \Pr\left(D^a < D^m\right) = \frac{\gamma}{\lambda + \gamma}, \tag{3}$$

because D^m and D^a are exponentially distributed. Thus, π measures of the arbitrageur's *relative* speed since it increases when $\frac{\gamma}{\lambda}$ goes up.

In practice, traders' response time to market events (e.g., an arbitrage opportunity) is stochastic (e.g., the time required by platforms to process orders is affected by a myriad of random factors). However, with technological investments (e.g., in hardware and software, fast data feed, dedicated communication lines, co-location etc.) and attention, traders can reduce their response times to market events (e.g., they can reduce latencies in communicating with trading platforms). Investments and attention are costly.[12] Hence, we assume that if the market maker operates at speed λ then he bears a cost $C^m(\lambda) = \frac{c^m\lambda}{2}$. Similarly, if the arbitrageur operates at speed γ then she bears a cost $C^a(\gamma) = \frac{c^a\gamma}{2}$.[13]

Allowing for the possibility that $c^a \neq c^m$ is useful to analyze the effects of exogenous shocks

[12] Attention can be interpreted literally as the effort that human traders must exert to follow prices in different markets. It can also represent the computing capacity that traders allocate to a particular task, e.g., detecting an arbitrage opportunity in a specific pair of assets. Allocating greater capacity to this specific task reduces the capacity available for other tasks, which generates an opportunity cost.

[13] We assume that costs of speed are linear to obtain closed form solutions for the equilibrium. Implications however do not crucially depend on this assumption. For instance, we have investigated the case in which the costs of speeds are quadratic in speeds and obtained identical implications.

to traders' relative speed (see below). However, our predictions do **not** depend on whether the market maker or the arbitrageur has the smallest cost of being fast. In practice, differences in costs of speed between the arbitrageur and the market maker might stem from differences in opportunity costs of attention for each type of activity. Alternatively, these costs could be identical but, for the same investment in speed, market design could enable one type of trader to react more quickly to market events (jumps in the value of asset Y).[14] For instance, suppose that $c^a = c^m = c$ but an investment of $\hat{\gamma}$ in speed for the arbitrageur produces an actual speed of only $\gamma = \kappa\hat{\gamma}$ with $\kappa < 1$. The optimal speed for the arbitrageur is then identical to the case in which his marginal cost of speed is c/κ.[15] Thus, an increase in κ is observationally equivalent to a reduction in $\frac{c^a}{c^m}$.

Traders' expected profits. Given our assumptions, the expected profit of market makers X, Y, and the arbitrageur are:

$$\Pi^X(S_X;\lambda,\gamma) = -\frac{\varphi\alpha\pi}{2}(\sigma - S_X) + \left(\frac{1-\alpha(2\varphi-1)}{4}\right)S_X - \frac{c^m\lambda}{2}, \tag{4}$$

$$\Pi^Y(S_Y) = \left(2\alpha(1-(1-\pi)\varphi) + (1-\alpha)\right)\frac{\sigma S_Y}{4}, \tag{5}$$

$$\Pi^a(S_X, S_Y;\lambda,\gamma) = \alpha\varphi\pi\frac{(\sigma - S_X - \sigma S_Y)}{2} + \alpha(1-\varphi)\frac{(\sigma - S_X - \sigma S_Y)}{2} - \frac{c^a\gamma}{2}. \tag{6}$$

The first term in eq. (4) is market maker X's expected losses when he trades with the arbitrageur in a toxic arbitrage opportunity. The second term is his expected profit when he trades with liquidity traders or with the arbitrageur in non-toxic arbitrage opportunities. Finally, the last term is the cost of speed for the market maker. Market maker Y trades $Q_Y = \sigma$ shares of asset Y with probability $(\alpha(1-(1-\pi)\varphi) + (1-\alpha)/2)$ and earns $S_Y/2$ on each share traded. He does not invest in speed since he is never exposed to the risk of trading at stale quotes. The expression for market Y's expected profit (eq.(5)) follows. The first two terms in the arbitrageur's expected profit (eq.(6)) are the arbitrageur's expected gain on toxic and non-toxic arbitrage opportunities, respectively. The last term is her cost of speed.

Equilibrium. For given bid-ask spreads S_X and S_Y, equilibrium speeds, $\gamma^*(S_X; c^a, c^m)$ and $\lambda^*(S_X; c^a, c^m)$, are such that (i) λ^* maximizes market-maker X 's expected profit (eq. (4)) when $\gamma = \gamma^*$ and (ii) γ^* maximizes the arbitrageur's expected profit (eq. (6)) when $\lambda = \lambda^*$.

[14] For instance, Hendershott and Moulton (2011) find that changes in the trading technology used by the NYSE in 2006 (the introduction of the so-called "Hybrid Market") increased the execution speed of market orders submitted by off-floor traders by a factor of 2. See Figure 2 in Hendershott and Moulton (2011).

[15] This follows directly from the first order conditions that characterize traders' optimal choices of speed.

That is, equilibrium speeds form a Nash equilibrium. Competitive equilibrium bid-ask spreads (S_X^*, S_Y^*) are such that market makers just break even given their choice of equilibrium speeds, i.e., $\Pi^X(S_X^*; \lambda^*, \gamma^*) = 0$ and $\Pi^Y(S_Y^*; \lambda^*, \gamma^*) = 0$. Eq.(5) immediately implies that $S_Y^* = 0$. The equilibrium bid-ask in asset X is derived in the next section.

2.2 Testable predictions

Traders' speed $\lambda^* > 0$ and $\gamma^* > 0$ form a Nash equilibrium if they solve the following system of first order conditions $\frac{\partial \Pi^a(S_X, 0; \lambda^*, \gamma^*)}{\partial \gamma} = 0$ and $\frac{\partial \Pi^X(S_X; \lambda^*, \gamma^*)}{\partial \lambda} = 0$.[16] We deduce that the unique equilibrium speeds for the arbitrageur and market maker X are:

$$\lambda^*(S_X; c^a, c^m) = \frac{\varphi\alpha(\sigma - S_X)\rho}{c^m(1+\rho)^2}, \tag{7}$$

$$\gamma^*(S_X; c^a, c^m) = \frac{\varphi\alpha(\sigma - S_X)\rho}{c^a(1+\rho)^2}, \tag{8}$$

where $\rho \equiv \left(\frac{c^m}{c^a}\right)$. Equilibrium speeds are strictly positive if $S_X \leq \sigma$, i.e., if arbitrage is profitable. This will be the case when S_X is competitive (see below).

Not surprisingly, traders' optimal speeds decreases in their cost of speed (e.g., γ^* decreases in c^a). They also decrease in market maker X's bid-ask spread. Indeed, an increase in the bid-ask spread reduces the transfer from market maker X to the arbitrageur when a toxic arbitrage opportunity arises and thereby both traders' incentive to react fast to an arbitrage opportunity. This is consistent with Deville and Riva (2007) who find that deviations from put-call parity last longer for less liquid options and Chordia et al. (2008) who find that short-horizons (five minutes) returns predictability (using past trades and returns to forecast future returns) is higher when bid-ask spreads are higher.

Substituting λ^* and γ^* by their expressions given by eq.(7) and eq.(8) in eq.(3), we deduce that:

$$\pi^*(\rho) = \frac{\gamma^*}{\lambda^* + \gamma^*} = \frac{\rho}{1+\rho}. \tag{9}$$

Hence, in equilibrium, the likelihood that a toxic arbitrage terminates with an arbitrageur's trade, π^*, increases in ρ, the ratio of the market maker's cost of speed to the arbitrageur's cost of speed. If $c^m = c^a$ then $\rho = 1$ and therefore $\pi^* = 1/2$.

Substituting the expression for λ^* (eq.(7)) in market maker X's expected profit (eq.(4)), we

[16] Clearly, traders' expected profits are concave in their speed. Hence, solving the first order conditions is sufficient to obtain equilibrium speeds.

solve for the zero profit bid-ask spread in asset X and we obtain:

$$S_X^* = \left(\frac{2\varphi\alpha\pi^*(\rho)(2-\pi^*(\rho))}{2\varphi\alpha\pi^*(\rho)(2-\pi^*(\rho)) + (1-\alpha(2\varphi-1))} \right) \times \sigma. \tag{10}$$

The competitive bid-ask spread increases with σ, the size of arbitrage opportunities and is always less than σ, as conjectured previously.[17] Market maker's X's bid-ask spread also increases when arbitrage opportunities are more likely to be toxic (φ is higher) or, holding this likelihood constant, when the arbitrageurs is faster ($\pi^*(\rho)$ is higher). Our two main testable implications follow.

Implication 1 *Consider a pair of assets X and Y linked by a no-arbitrage relation. An increase in the likelihood that an arbitrage opportunity is toxic (φ) causes an increase in the bid-ask spread of asset X.*

Implication 2a. *Consider a pair of assets X and Y linked by a no-arbitrage relation. The bid-ask spread of asset X increases in the likelihood (π^*) that a toxic arbitrage opportunity terminates with an arbitrageur's trade.*

The likelihood, π^*, that a toxic arbitrage opportunity terminates with an arbitrageur's trade (rather than a quote update) is endogenous and jointly determined with the bid-ask spread. In testing Implication 2a, one must therefore account for the endogeneity of π^*. The model suggests to use shocks to the arbitrageur's relative cost of speed ($\rho = c^a/c^m$) as a source of exogenous variations for π^*. Indeed, a decrease in the relative cost of speed for the arbitrageur (i.e., an increase in ρ) triggers an increase in π^* and, through this channel only, an increase in the bid-ask spread. This yields the following implication.

Implication 2b. *Consider a pair of assets X and Y linked by a no-arbitrage relation. A reduction in arbitrageurs' relative cost of speed (i.e., an increase in ρ) triggers an increase in π^* – the probability of an arbitrageur's trade, conditional on the occurrence of a toxic arbitrage – and, through this channel, it increases the bid-ask spread of asset X.*

As explained previously, changes in market structures that affect traders' speeds are similar to variations in $\rho \equiv c^m/c^a$. Implication 2b shows that these changes provide good instruments to measure the effect of π^* on the bid-ask spread, i.e., to test Implication 2a, because they

[17] This is the case even if $\alpha = 1$. In this case, trades happen only between market makers and the arbitrageur. The market does not break down because market maker X makes money when the arbitrageur exploits arbitrage opportunities due to a liquidity shock to market maker Y's valuation. In these opportunities, the net gain from trade (i.e., the sum of all traders' expected profit) is equal to $\sigma/2$, i.e., the size of market maker Y's liquidity shock times the number of shares traded (σ) by this market maker.

should affect liquidity only through their effects on traders' speeds and therefore π^*. We use this approach in Section 3.5.

The model has also predictions for the durations of arbitrage opportunities. This is interesting because we can measure these durations empirically and they are obviously related to the speed at which traders react to events. Let D be the duration of an arbitrage opportunity (the "time-to-efficiency.") The expected time-to-efficiency is:

$$E(D) = \varphi E\left(\mathrm{Min}\{D^a, D^m\}\right) + (1-\varphi)E(D^a) = \frac{(1+\rho)-\varphi}{\gamma^*(S_X^*)(1+\rho)}, \tag{11}$$

As speed is costly, equilibrium speeds are never infinite. Thus, in equilibrium, arbitrage opportunities do not immediately vanish, i.e., $E(D) > 0$. However, the duration of arbitrage opportunities goes to zero when c^a or c^m go to zero because the arbitrageur or the market maker become increasingly fast in reacting to the opportunity. Thus, arbitrage opportunities can be very short-lived in equilibrium, as observed in our data.

Now, consider a technological change that reduces c^a and c^m but for which the effect on c^a is bigger. Hence ρ increases. As c^a and c^m are lower, the arbitrageur and the market maker react faster to arbitrage opportunities, holding the bid-ask spread constant (see eq. (7) and (8)). This effect reduces the average duration of arbitrage opportunities. However, the bid-ask spread becomes larger in equilibrium because ρ is larger (Implication 2b). This indirect effect reduces the arbitrageur's incentive to be fast and therefore the market maker's incentive to be fast as well. Nevertheless, using eq.(11), we show in the appendix that the direct effect always dominates the indirect effect. This yields our next testable implication.

Implication 3: *Consider a decrease in c^a and c^m that triggers an increase in $\rho = \frac{c^m}{c^a}$ (i.e., the decrease in c^a is larger). This change should trigger a reduction in the average duration of arbitrage opportunities (see the appendix for a formal proof).*

Chaboud et al. (2014) find that algorithmic trading leads to fewer triangular arbitrage opportunities per second in the FX market. They also find that this reduction is mainly due to algorithmic arbitrageurs hitting quotes of slower (human) traders. These findings are consistent with the model. Indeed, suppose that algorithmic trading reduces the cost of being fast for arbitrageurs. Then, ρ, and therefore π^*, increases. As a result, arbitrage opportunities terminate more frequently with arbitrageurs' trades when algorithmic trading increases, as found by Chaboud et al. (2014). Moreover, the duration of arbitrage opportunities should decrease

(Implication 3). Hence, if one checks for the presence of arbitrage opportunities at fixed time intervals (say, every second), one should observe a negative relationship between algorithmic trading and the frequency of arbitrage opportunities (again as found by Chaboud et al. (2014)).

Finally, consider the effect of φ (the likelihood that an arbitrage opportunity is toxic conditional on the occurrence of an opportunity) on the duration of arbitrage opportunities. The direct effect of an increase in φ is to induce the arbitrageur and the market maker to react faster to arbitrage opportunities (holding S_X constant, γ and λ increase in φ; see eq. (7) and (8)). The indirect effect however is that the equilibrium bid-ask spread increases (Implication 1). In general, the first effect dominates, except when α and φ are both large.

Implication 4: *The average duration of arbitrage opportunities decreases with the likelihood that an arbitrage opportunity is toxic, φ if $\varphi \leq 1/2$ or, if $\varphi > 1/2$, when $\alpha \leq (4\varphi - 1)^{-1}$ (see the appendix for a formal proof).*

In our data, the number of arbitrage opportunities relative to the total number of trades is small and well below 1/3, which is a sufficient condition for $\alpha < (4\varphi - 1)^{-1}$ when $\varphi > 1/2$. Hence, according to Implication 4, we expect the duration of arbitrage opportunities to be negatively related to φ.

[Insert Figure I about here]

Figure I illustrates the four testable implications of the model. We set $\alpha = 0.1$ and $\sigma = 3.5$ basis points. We measure time in seconds. We fix the cost of speed for the market maker at $c^m = 0.056$ and let the arbitrageur's cost, c^a, vary from 0 to 0.1. When $c^a = 0.02$, this implies that $\rho = 2.8$ and $\pi = 74\%$, which is the average value of π in the data (see Table II below). We assume that $\varphi = 41\%$, which also is its average value in our data (see Table II). When $c^a = 0.02$ and $\varphi = 41\%$, the average duration of an arbitrage opportunity is about 700 milliseconds in equilibrium, which is close to the average duration in our data. For these values of the parameters, the equilibrium bid-ask spread is $S_X^* = 0.24$ basis points, which is about 8% to 15% of the average bid-ask spread in our data (depending on the currency). By design, the model only captures the component of bid-ask spreads that is a compensation for adverse selection costs due to toxic arbitrage opportunities. These costs should only be a fraction of total costs for market makers and cannot therefore be a too large fraction of the bid-ask spread.

2.3 Extensions

In this section, we discuss some extensions of the model. For brevity, we omit full derivations of the equilibrium in each case discussed below. They are available upon request. We have checked that our testable implications still hold in each case.

Competition among arbitrageurs. The baseline model features a single arbitrageur. The case with $M > 1$ arbitrageurs is straightforward to analyze and delivers identical implications.[18] Not surprisingly, as M increases, the equilibrium bid-ask spread becomes larger because the likelihood that one arbitrageur reacts first to the arbitrage opportunity becomes higher. Moreover, if $\rho \leq \frac{M}{(M-1)}$ then, in equilibrium, the market maker always chooses a zero speed (this never happens when $M = 1$). Intuitively, when the number of arbitrageurs increases, each increment in the market maker's speed has a smaller effect on the likelihood that he can update his quotes before being hit by an arbitrageur. As a result the marginal benefit of speed is lower for the market maker. If the number of arbitrageurs is large enough relative to ρ then the market maker is better off not investing in speed at all.

Market-makers as arbitrageurs. We assumed that arbitrageurs and market makers are distinct agents. This is not required. For instance, suppose that there is free entry in market making and arbitrage activities ($M \to \infty$). Just before $t = 0$, each firm decides whether to be a market maker (post quotes) or an arbitrageur. As there is free entry in both market making and arbitrage, arbitrageurs' expected profit is zero in equilibrium.[19] Market-makers' expected profit is also zero since the bid-ask spread is competitive. Thus, all firms are indifferent between each role and one can be randomly selected to be a market maker.

Shocks to market maker X's valuation. In the baseline model, arbitrage opportunities are only due to shifts in the price of asset Y (i.e., always originates in asset Y). Suppose instead that a shock to a market maker's valuation at $t = 0$ can affect either market maker X or Y with equal probabilities and that $\sigma = 1$. Under these assumptions, the two markets for the two assets are completely symmetric. Thus, in equilibrium bid-ask spreads are strictly positive and identical in each asset ($S_Y^* = S_X^* > 0$). In this case, there is no closed form solution for the competitive bid-ask spread in closed form.[20] The equilibrium can be solved numerically,

[18] In this case, an arbitrageur who chooses a speed γ_i exploits the arbitrage opportunity with probability $\frac{\gamma_i}{\lambda+\sum_{j=1}^{j=M}\gamma_j}$.

[19] When M goes to infinity, each arbitrageur's speed (γ^*) goes to zero but arbitrageurs' aggregate speed ($M\gamma^*$) remains strictly positive in equilibrium.

[20] In this more general case, π^* is a non linear decreasing function of the bid-ask spread (in the baseline model, π^* only depends on ρ for any value of the bid-ask spread). As π^* is non linear in the bid-ask spread,

however. Numerical simulations show that our predictions still hold in this more general case. In particular, Implications 1, 2a, and 2b hold for asset Y as well.

3 Empirical Analysis

3.1 Data

Our tests use tick-by-tick data from Reuters D-3000 for three currency pairs: US dollar/euro (dollars per euro), US dollar/pound sterling (dollars per pound), and pound sterling/euro (pounds per euro); hereafter USD/EUR, USD/GBP and EUR/GBP, respectively. Reuters D-3000 is one of the two electronic trading platforms for interdealer spot trading in the foreign exchange (FX) market over our sample period (the other one being Electronic Broking Services (EBS)).[21] The sample period contains 498 days from January 2, 2003 to December 30, 2004.[22] The Bank for International Settlement (BIS, 2004) estimates that the currency pairs in our sample account for 60 percent of all FX spot transactions at the time of our sample.

Reuters D-3000 is an electronic limit order book market similar to that used in major equity markets. On this system, foreign exchange dealing banks ("FX dealers") can post quotes (by submitting limit orders) or hit quotes posted by other dealers (by submitting market orders).[23] Trade sizes are only allowed in multiple of millions of the "base" currency. Our dataset contains all orders (limit and market) submitted to Reuters D-3000 over the sample period.

Our dataset has several attractive features for our tests. First, it is very detailed. For each order submitted to Reuters D-3000, the dataset reports (i) the currency pair in which the order is submitted; (ii) the order type (limit or market); (iii) the time at which the order is entered;

market makers' expected profits are non linear in their bid-ask spread and competitive bid-ask spreads solve a cubic polynomial.

[21] See Pierron (2007), Osler (2008), King and Rime (2010), and King et al. (2012) for excellent descriptions of participants, market structure, and recent developments in foreign exchange markets. The FX market is a two-tier market. In the first tier, FX dealers trade exclusively with end-users (e.g., hedge funds, mutual funds, pension funds, corporations, etc.). The second-tier is an interdealer market. In this market, dealers can trade (i) bilaterally (by calling each other), (ii) through voice brokers, or (iii) electronic broker systems (EBS and Reuters D-3000). In the last decade, the market share of EBS and Reuters D-3000 has considerably increased and was already large at the time of our sample (see Pierron (2007) and King et al. (2012)). At the time of our sample, interdealer trades for currency pairs in our sample account for about 53% of all interdealer trading in foreign exchange markets. King et al. (2012) note that the frontiers between the two-tiers of the FX market have been breaking down in recent years.

[22] The foreign exchange (FX) market operates around the clock, all year long. However, trading activity in this market considerably slows down during weekends and certain holidays. Hence, as is standard (see, for instance, Andersen et al. (2003)), we exclude the following days from our sample: weekends, the U.S. Independence Day (July 4 for 2003 and July 5 for 2004), Christmas (December 24 - 26), New Years (December 31 - January 2), Good Friday, Easter Monday, Memorial Day, Thanksgiving and the day after and Labor Day.

[23] Dealers use both types of orders. Using data from Reuters' trading platform, Bjønnes and Rime (2005) (Table 11) show that some market makers frequently use market orders to build up speculative positions and limit orders to reduce their position.

(iv) the size of the order; and (v) the price attached to limit orders. We also know for each transaction whether the market (or marketable) order initiating the transaction is a buy order or a sell order. As each order has a unique identifier, we can track it over its life. Thus, we can reconstruct the entire limit order book of each currency at any point in time. In this way, we can use the slope of the book as a measure of market illiquidity, in addition to standard measures such as bid-ask spreads. Furthermore, and more importantly for our purpose, we can identify whether an arbitrage opportunity terminates with a trade (the submission of a market order) or quote updates in limit order books for the three currencies in our sample. We can therefore accurately measure the frequency with which a toxic arbitrage opportunity terminates with a trade, that is, π^*, in the model.

Second, the time stamp of the data has an accuracy of one-hundredth of a second. Hence, we can accurately measure when an arbitrage opportunity begins, when it finishes, and its duration. Furthermore, we can track the evolution of prices after an arbitrage opportunity terminates. As explained below, we use this feature to classify arbitrage opportunities in toxic and non-toxic opportunities, which is another requirement for our tests.

At the time of our sample, Reuters D-3000 had a dominant market share in USD/GBP and EUR/GBP but its competitor, EBS, had the Lion's share of trades in the USD/EUR pair. This is not a problem for our tests because we exclusively focus on triangular arbitrage opportunities *within* Reuters D-3000. When an arbitrageur exploits a toxic arbitrage opportunity in Reuters D-3000, it inflicts a loss on market makers with stale quotes on this platform. Hence, quotes in Reuters D-3000 should reflect this risk, as predicted by our model. We will however underestimate the frequency of triangular arbitrage opportunities in the currency pairs in our sample, as some might arise between Reuters D-3000 and EBS and within EBS. However, estimating this frequency is not our goal.

For a given currency pair, measures of market liquidity on Reuters D-3000 and EBS are correlated because they are affected by common factors. We will therefore use liquidity measures from EBS to control for systematic time-series variations in liquidity in these pairs. Our EBS data, acquired from ICAP, are similar to those for Reuters D-3000 with one important difference. The time-stamps of quotes, trades etc. are accurate only up to the second. In particular, all trades occurring within the same second receive the same time stamp. Thus, data from EBS cannot be used to accurately measure when a triangular arbitrage opportunity (across trading systems or within EBS) starts, and when and how it terminates.[24] For this reason, we just use EBS data as controls in our regressions (see below).

3.2 Toxic and non-toxic arbitrage opportunities

In this section, we explain how we identify triangular arbitrage opportunities in our data and how we classify them into two groups: toxic and non-toxic opportunities.

3.2.1 Triangular arbitrage opportunities

Let $A_t^{i/j}$ be the number of units of currency i required, at time t, to buy one unit of currency j and $B_t^{i/j}$ be the number of units of currency i received for the sale of one unit of currency j. These are the best ask and bid quotes posted by market makers in currency i vs. j at time t. A trader can buy one unit of currency j with currency i *directly*, at cost $A_t^{i/j}$ or *indirectly* by first buying $A_t^{k/j}$ units of currency k with currency i and then buying one unit of currency j at $A_t^{k/j}$ in the market for currency k vs. j. The cost of this alternative strategy is $\hat{A}_t^{i/j} \equiv A_t^{i/k} \times A_t^{k/j}$. Similarly, a trader with one unit of currency j can sell it directly in exchange of $B_t^{i/j}$ units of currency i by trading in the market for i vs. j. Alternatively, he can obtain $\hat{B}_t^{i/j} = B_t^{i/k} \times B_t^{k/j}$ units of currency i by first selling currency j for $B^{k/j}$ units of currency k and then by selling these units of currency k for $B_t^{i/k} \times B_t^{k/j}$ units of currency i. We refer to $\hat{A}_t^{i/j}$ and $\hat{B}_t^{i/j}$ as the synthetic quotes for currency j in the i vs. j market.

A triangular arbitrage opportunity exists when

$$\hat{A}_t^{i/j} < B_t^{i/j} \quad \text{or,} \tag{12}$$

$$\hat{B}_t^{i/j} > A_t^{i/j}. \tag{13}$$

In the first case, one can secure a risk free profit equal to $(B_t^{i/j} - \hat{A}_t^{i/j})$ units of currency i by selling one unit of currency j in the market of currency j vs. i while simultaneously buying it at price $\hat{A}^{i/j}$ with two transactions in other currency pairs. In the second case, one can secure a risk free profit equal to $(\hat{B}_t^{i/j} - A_t^{i/j})$ units of currency i by buying one unit of currency j in the market of currency j vs. i while simultaneously selling it at price $\hat{B}^{i/j}$. These two arbitrage opportunities cannot occur simultaneously because (12) and (13) cannot both be true at the

[24] For example, suppose two market orders and two limit orders are submitted in a second in which an arbitrage opportunity occurs and that the arbitrage starts and terminates within this second. EBS data do not allow us to identify whether the arbitrage terminates due to a market order (a trade) or a limit order (a quote update). Hence, we cannot compute our proxy for π^* using EBS data.

same time; see Kozhan and Tham (2012).

Reuters D-3000 charge brokerage and membership fees. To account for these costs and other possible unobserved frictions, we say that a triangular arbitrage opportunity exists at time t iff one of the two following inequalities is satisfied:

$$\frac{B_t^{i/j} - \hat{A}_t^{i/j}}{\hat{A}^{i/j}} - z \quad > \quad 0, \tag{14}$$

$$\frac{\hat{B}_t^{i/j} - A_t^{i/j}}{\hat{B}_t^{i/j}} - z \quad > \quad 0, \tag{15}$$

where z is the fee per trade for exploiting an arbitrage opportunity (expressed as a fraction of synthetic quotes). Chaboud et al. (2014) note that this cost is well below one basis points. Hence, we, conservatively, set $z = 1$ basis point.

Table I provides an example. It gives best ask and bid prices for the three currency pairs in our sample at a given point in time. These quotes are such that there is no triangular arbitrage opportunity, even for $z = 0$.

Table I: Triangular Arbitrage Opportunities: An Example

Quotes	$B^{i/j}$	$A^{i/j}$	Mid-Quote $((A^{i/j} + B^{i/j})/2)$
\$/€	1.0770	1.0780	1.0775
\$/£	1.6255	1.6265	1.6260
£/€	0.6622	0.6632	0.6627

Now suppose that the best quotes in the USD/EUR market become $A^{\$/€} = 1.075$ and $B^{\$/€} = 1.074$ (a depreciation of the euro against the dollar). If other rates are unchanged, we have $\hat{B}^{\$/€} = 1.0764$. As $\hat{B}^{\$/€} > A^{\$/€} = 1.075$, there is a triangular arbitrage opportunity. An astute arbitrageur can buy at least 1 million euros at 1.075 dollar per euro and resell them instantaneously at 1.0764 dollar per euro (with two transactions in other currencies). If the arbitrageur successfully executes the three transactions required for this arbitrage before quotes are updated, she makes a profit of, at least, $\$1,400$.

There are two ways in which the arbitrage opportunity can be eliminated. The first possibility is that market makers update their quotes before an arbitrageur actually profits from the opportunity. For instance, market makers in the USD/GBP market may update their quotes and post new ones at, say, $A^{\$/£} = 1.6215$ and $B^{\$/£} = 1.6213$. The second possibility is that an arbitrageur exploits the arbitrage opportunity by, as we just explained, submitting (i) buy market orders in the \$/€ market and (ii) sell market orders in the \$/£, and £/€ markets.

Thus, in our empirical tests, we identify how arbitrage opportunities start and terminate as follows.

1. Starting from a state in which there is no-arbitrage opportunity (i.e., a state in which (14) and (15) do not hold), we record the latest quoted best bid and best ask prices for the three currency pairs each time a new limit order is submitted and we check whether a triangular arbitrage opportunity exists using (14) and (15).[25]

2. If an arbitrage opportunity exists, we deduce that the limit order arrival created the arbitrage opportunity. We therefore record the order arrival time, t_0, as the time at which the arbitrage opportunity begins. We call the currency pair in which the limit order was submitted the "initiating currency" since the arbitrage opportunity is triggered by a price revision in this currency.

3. We then record the first time t_1 at which the triangular arbitrage opportunity disappears and the duration of the arbitrage opportunity $((t_1 - t_0))$. We also record whether the arbitrage opportunity terminates with a trade or quote updates.

3.2.2 Classifying arbitrage opportunities

For our tests, we must measure the fraction (φ) of arbitrage opportunities that are toxic and the likelihood that a toxic arbitrage opportunity terminates with a trade from an arbitrageur (π). For this, we must first classify triangular arbitrage opportunities into two subgroups: toxic and non-toxic.

We proceed as follows. As in Shive and Schultz (2010), we consider that an arbitrage opportunity is due to a price pressure effect (i.e., is non-toxic) if the price change at the origin of this opportunity reverts after the opportunity terminates.[26] If instead this price change

[25] The arrival of a market order cannot create an arbitrage opportunity. For instance, suppose that $A_t^{i/j} > \hat{B}_t^{i/j}$ and $\hat{A}_t^{i/j} > B_t^{i/j}$, so that, at time t, there is no-arbitrage opportunity. If a buy (resp. sell) market order arrives in the market for currency i vs. j, it executes at $A_t^{i/j}$ (resp. $B_t^{i/j}$) and possibly higher (resp. lower) prices (if the size of the order is larger than the quantity available) at $A_t^{i/j}$. Thus, just after the arrival of a buy (sell) market order (at time t^+), the new best offer (bid) price is $A_{t^+}^{i/j} \geq A_t^{i/j} > \hat{B}_t^{i/j}$ ($B_{t^+}^{i/j} \leq B_t^{i/j} < \hat{A}_t^{i/j}$), so that no arbitrage opportunities exist if there is none at date t. In contrast, if a market maker posts a new ask price $A_{t^+}^{i/j}$ at time t^+ such that $A_{t^+}^{i/j} < \hat{B}_t^{i/j}$ then he creates an arbitrage opportunity. The same is true for orders arriving in other currency pairs and affecting the synthetic quotes.

[26] Shive and Schultz (2010) show that profitable arbitrage opportunities exist in dual-class stocks because the bid price of the voting share sometimes exceeds the ask price of the non-voting share. They also find that these arbitrage opportunities arise either from price pressures effects or asynchronous price adjustments, the former case being more frequent than the latter (as in our sample; see below).

persists after the arbitrage opportunity terminates, we consider that the arbitrage opportunity is due to asynchronous permanent price adjustments in the rates of the three currencies. We classify this opportunity as being toxic.

More specifically, for each triangular arbitrage opportunity in our sample, we compare the exchange rate for the three currency pairs when the arbitrage opportunity begins (time t) and when it terminates (time $t+\tau$). If these rates are identical at dates t and $t+\tau$ or if they do not move in a direction consistent with a toxic triangular arbitrage opportunity, we classify them as being non-toxic (that is, due to a price pressure effect). Remaining arbitrage opportunities are classified as toxic.

Insert Figure II about here.

Figure II illustrates this methodology by considering four arbitrage opportunities that actually occurred in our sample. In Panels A and B, the solid and the dashed lines show the evolution of bid and ask quotes ($A^{i/j}$ and $B^{i/j}$) and synthetic quotes ($\hat{A}^{i/j}$ and $\hat{B}^{i/j}$) during these arbitrage opportunities, respectively. In Panel A, actual and synthetic quotes of the currency pairs initiating the arbitrage opportunity shift permanently to a new level when the arbitrage opportunity terminates. The pattern is consistent with the arrival of information regarding fundamentals (e.g., macroeconomic news announcements or headlines news on Reuters). Thus, we classify these opportunities as toxic. In contrast, in Panel B, only the quotes of the initiating pair change during the arbitrage opportunity. Moreover, these quotes revert to their initial level when the arbitrage opportunity terminates. This pattern (reversal and the absence of changes in the synthetic quotes) is consistent with price movements arising from price pressure effects. Accordingly, we classify these arbitrage opportunities as non-toxic.

Insert Figure III about here

Using this methodology, we identify $37,689$ triangular arbitrage opportunities in our sample, of which $15,908$ are classified as toxic. Panel A of Figure III shows the time-series of the daily number of (a) all triangular arbitrage opportunities (light grey line) and (b) toxic arbitrage opportunities (black line) in our sample. There is substantial daily variation in the number of arbitrage opportunities with some days having a high number of arbitrage opportunities (e.g., in May or June 2003) and other days having much fewer opportunities. There are on average 32 (s.d=20.83) toxic triangular arbitrage opportunities and 45 non-toxic arbitrage opportunities per day.

Panel B of Figure III shows average intra-day patterns in the number of arbitrage opportunities. The bulk of the activity for currency pairs in our sample occurs when European markets are open, that is, from 7:00 GMT (European markets open) until 17:00 GMT (European markets close). Not surprisingly, most arbitrage opportunities occur during this period, with peaks when trading activity in the U.S. and in Europe overlap (13:00 to 17:00). Hence, for the variables used in our tests (see below), we only retain observations from 7:00 to 17:00 GMT.

3.3 Variables of interest

In this section, we describe the variables used in our empirical analysis. On each day t, we define $\hat{\varphi}_t$ as the ratio of the number toxic arbitrage opportunities on day t to the total number of arbitrage opportunities on this day:

$$\hat{\varphi}_t = \frac{\text{No. of toxic arbitrage opportunities on day } t}{\text{No. of all arbitrage opportunities on day } t}. \tag{16}$$

This ratio is a proxy for φ, i.e., the likelihood that an arbitrage opportunity is toxic. It characterizes the composition of arbitrage opportunities ("arbitrage mix") on day t.

Another important variable in our model is π, the likelihood that a toxic arbitrage opportunity terminates by an arbitrageur's trade. We proxy π with the daily fraction of toxic arbitrage opportunities that terminate with a trade. We denote this fraction by $\hat{\pi}_t^{tox}$ on day t:

$$\hat{\pi}_t^{tox} = \frac{\text{No. of toxic arbitrage opport. that terminate with a trade on day } t}{\text{No. of toxic arbitrage opportunities on day } t}. \tag{17}$$

Similarly, on each day, we also compute the fraction, denoted $\hat{\pi}_t^{nontox}$, of non-toxic arbitrage opportunities that terminate with an arbitrageur's trade:

$$\hat{\pi}_t^{nontox} = \frac{\text{No. of non-toxic arbitrage opport. that terminate with a trade on day } t}{\text{No. of non-toxic arbitrage opportunities on day } t}. \tag{18}$$

The likelihood of occurrence of an arbitrage opportunity (α) and the size of arbitrage opportunities, σ, also affect illiquidity according to the model (see eq. (10)). As explained previous, α is best interpreted as the ratio of the number of arbitrage opportunities relative to the total number of trades in the model. Thus, we proxy α on day t by by $\hat{\alpha}_t$, the number of all arbitrage opportunities on day t divided by the total number of trades on this day.

We measure σ empirically as follows. Suppose that a toxic arbitrage opportunity occurs at time τ on date t. For each currency pair (say i/j), let $f_{\tau,t}^{i/j} = \frac{A_{\tau,t}^{i/j}+B_{\tau,t}^{i/j}}{2}$ and $\hat{f}_{\tau,t}^{i/j} = \frac{\hat{A}_{\tau,t}^{i/j}+\hat{B}_{\tau,t}^{i/j}}{2}$

be the mid-quotes based on actual quotes and synthetic quotes, respectively, at the time of the arbitrage opportunity. We then compute, $\hat{\sigma}_t^{tox}$, the daily average *percentage* absolute difference between $f_{\tau,t}^{i/j}$ and $\hat{f}_{\tau,t}^{i/j}$ for all currency pairs and use it as a proxy for the size of toxic arbitrage opportunities on day t.

In our tests, we also use currency-specific controls known to be correlated with measures of market illiquidity: the average daily trade size in each currency (denoted $trsize_{it}$ in currency i on day t); the daily realized volatility, i.e., the sum of squared five minutes mid-quote returns in each currency (denoted vol_{it}); and the daily number of orders (entry of new limit and market orders as well as limit order updates) denoted $nrorders_{it}$. This variable measures the level of activity on Reuters D-3000 on each day.

We use three different measures of illiquidity in each currency pair i: (i) the average daily *percentage* quoted bid-ask spread ($spread_{i,t}$ for currency pair i on day t), that is, the absolute quoted spread divided by the mid-quote; (ii) the average daily effective spread ($espread_{i,t}$), i.e., twice the average absolute difference between each transaction price and the mid-quote at the time of the transaction; and (iii) the average daily slope of the limit order book ($slope_{it}$). For currency i, $slope_{it}$ is the average of: (i) the ratio of the difference between the second best ask price and the first best ask price at date t to the number of shares offered at the best ask price and (ii) the same ratio using quotes on the buy side of the limit order book. Hence, $slope_{it}$ is higher when the number of shares offered at the best quotes is lower and the second best prices in the book are further away from the best quotes. A higher $slope_{it}$ is associated with a less liquid market.

Finally, we compute the duration of each arbitrage opportunity and we denote the average duration (Time-To-Efficiency) of toxic (resp., non-toxic) arbitrage opportunities on day t by TTE_t^{tox} (resp., TTE_t^{nontox}).

3.4 Summary statistics

Table II presents descriptive statistics for all variables used in our analysis. Panels A and B present the characteristics of toxic and non-toxic arbitrage opportunities. Both types of arbitrage opportunities vanish very quickly: they last on average for about 0.894 seconds (standard deviation: 0.301) and 0.518 seconds (s.d.: 0.199), respectively.

Insert Table II here

On average, $\hat{\varphi} = 41.5\%$ (s.d.= 10%). That is, about 41% of all arbitrage opportunities on a given day are toxic on average. The average size of a toxic arbitrage opportunity, $\hat{\sigma}^{tox}$, is 3.535 basis points (s.d. = 0.757). The average daily arbitrage profit (expressed in percentage term after accounting for trading costs as in (14) and (15)) on a toxic arbitrage opportunity is 1.427 basis points (s.d.= 0.277). These statistics are similar for non-toxic arbitrage opportunities.

Quotes are valid for at least one million of basis currency on Reuters. Thus, the minimum average profit opportunity on a toxic (resp. non-toxic) triangular arbitrage opportunity is $143 (resp., $161) or $4,576 ($8,583.42) per day. As a point of comparison, Brogaard et al. (2014) report that, after accounting for trading fees, high frequency traders in their sample earn $4,209.15 per stock-day (see Table 4 in Brogaard et al. (2014)) on their market orders (i.e., liquidity taking orders) in large-cap stocks and much less in small-cap stocks. This is of the same order of magnitude as daily revenues on triangular arbitrage opportunities in our sample.[27]

The likelihood that a toxic arbitrage opportunity terminates with a trade ($\hat{\pi}^{tox}$) is 74.1% on average (s.d.= 0.110). The likelihood that a non-toxic arbitrage opportunity ($\hat{\pi}^{nontox}$) terminates with a trade is higher on average (80.7%). This is consistent with the model: liquidity suppliers in non-initiating currencies have no incentive to cancel their quotes when a non-toxic arbitrage opportunity occur. Thus, $\hat{\pi}^{nontox}$ should be larger than $\hat{\pi}^{tox}$, as we find. There are several possible reasons why yet $\hat{\pi}^{nontox}$ is less than 1, in particular the price pressure at the origin of the arbitrage opportunity might disappear because the market maker in the initiating currency trades on EBS. Alternatively, this might reflect errors in our classification of arbitrage opportunities.

In any case, arbitrage opportunities terminate in more than 2/3 of the cases with a trader hitting quotes posted in limit order books rather than with traders updating their quotes. This is consistent with Chaboud et al. (2014) who find (i) a significant negative relationship between liquidity taking (i.e., market orders) algorithmic orders and the frequency of triangular arbitrage opportunities over one minute intervals and (ii) no such relationship between liquidity making (i.e., limit orders) algorithmic orders (see Table II in Chaboud et al. (2014)).

[27] Another benchmark for daily profits on triangular arbitrage opportunities are FX dealers' daily profits. Bjønnes and Rime (2005) find an average profit of about $12,000 per day for four currency dealers (we infer this number from the weekly profits they report on page 597 of their paper). Hence, profits on triangular arbitrage opportunities would not appear negligible for the trading desks studied by Bjønnes and Rime (2005).

Panel C of Table II reports summary statistics for our various measures of illiquidity, separately for the Reuters and the EBS trading platforms. Average quoted and effective bid-ask spreads are very tight (between 1 basis point and 5 basis points). The most illiquid currency pair is GBP/USD. For instance, on Reuters, the average effective spread for this pair is 2.073 basis point while the effective spread for EUR/GBP (the most liquid pair on Reuters) is 0.966 basis points. Measures of illiquidity are much higher on EBS for the GBP/USD and EUR/GBP pairs. For instance, in EUR/GBP, quoted spreads on Reuters are equal to 1.352 basis points on average vs. 2.52 basis points for EBS. In contrast, EBS is more liquid for the EUR/USD pair. These differences in liquidity between the two platforms are consistent with their respective market shares in each currency pair. Finally, Panel D presents descriptive statistics for the distribution (mean, standard deviation, min and max values, etc.) of other control variables used in our regressions.

Insert Table III about here

Table III reports the correlation of the main variables used in our tests. Consistent with Implication 1, measures of illiquidity for the three currency pairs in our sample are positively and significantly correlated with $\hat{\varphi}$, the fraction of arbitrage opportunities that are toxic. Moreover, as expected, all measures of illiquidity are also positively and significantly correlated with the size of toxic arbitrage opportunities, ($\hat{\sigma}^{tox}$).

The correlation between $\hat{\pi}^{tox}$ (our proxy for π^*) and measures of illiquidity is positive, consistent with Implication 2a) in the model. This correlation is not significantly different from zero, however. This is not surprising given that the endogeneity of π^* should work to weaken the correlation between illiquidity and π^*. Indeed, when bid-ask spreads are small, arbitrageurs have more incentive to react fast to arbitrage opportunities since they are more profitable. This effect works to make π^* higher when bid-ask spreads are smaller (it arises in the theory when jumps in prices can occur either in asset X or Y; see Footnote 20). This highlights the importance of accounting for endogeneity issues in analyzing the effect of arbitrageurs' relative speed on illiquidity.

The correlation between $\hat{\pi}^{nontox}$ and measures of illiquidity is in general significantly negative for all currency pairs in the sample. This is expected because, as discussed earlier, market makers benefit from non-toxic arbitrage trades. First, market makers who initiate the arbitrage opportunity can share risks with arbitrageurs. Furthermore, market makers in non-initiating currencies earn the bid-ask spread on trades with arbitrageurs (as, from their standpoint, non-toxic arbitrage trades are uninformed).

The correlation between $\hat{\pi}^{nontox}$ and $\hat{\pi}^{tox}$ is slightly positive and statistically significant. This low correlation indicates that daily variations in $\hat{\pi}^{nontox}$ and $\hat{\pi}^{tox}$ are not driven by the same factors. This is consistent with the idea that market makers in non-initiating pairs should behave differently in toxic and non-toxic arbitrage opportunities: they should update their quotes as quickly as possible in toxic arbitrage opportunities while they have no reason to do so when arbitrage opportunities are non-toxic.

The duration of toxic arbitrage opportunities (TTE^{tox}) and the various measures of market illiquidity are positively correlated. That is, toxic arbitrage opportunities last longer on average when the market for the three currency pairs is more illiquid. This observation is consistent with the model. Other things equal, a higher bid-ask spread induces arbitrageurs (and therefore market makers) to react more slowly to arbitrage opportunities, which eventually results in more persistent arbitrage opportunities (see the discussion following eq.(11)).

3.5 Tests

3.5.1 Is illiquidity sensitive to the arbitrage mix (φ)?

In this section, we test whether an increase in the likelihood that an arbitrage opportunity is toxic, φ_t, is a source of illiquidity, as predicted by Implication 1. To test this implication, we estimate the following equation:

$$illiq_{it} = \omega_i + \xi_{t,m} + b_1\hat{\varphi}_t + b_2\hat{\alpha}_t + b_3\hat{\sigma}_t^{tox} + b_4 vol_{it} + b_5 trsize_{it} + b_6 nrorders_{it} + b_7 illiq_{it}^{EBS} + \varepsilon_{it}, \quad (19)$$

where $illiq_{it}$ is one of our three proxies for illiquidity for currency i on day t, ω_i is a currency pair fixed effect, and $\xi_{t,m}$ is a monthly fixed effect (a dummy equal to one if day t is in month m). Coefficient b_1 measures the sensitivity of illiquidity to our proxy for φ and should be positive according to Implication 1. The model also implies that the bid-ask spread of each currency should be larger on days in which the number of arbitrage opportunities per trade ($\hat{\alpha}_t$) or the size of toxic arbitrage opportunities ($\hat{\sigma}_t^{tox}$) are higher. Thus, we expect b_2 and b_3 to be positive. Other control variables are described in Section 3.3. For each illiquidity measure, we also include its EBS counterpart ($illiq_{it}^{EBS}$) as an explanatory variable to control for market-wide unobserved variables that create daily variations in the illiquidity of the currency pairs in our sample (and which therefore should affect illiquidity similarly on Reuters D-3000 and EBS). We estimate equation (19) with OLS using standard errors robust to heteroscedasticity and time series autocorrelation.

Insert Table IV

Results are reported in Columns (1), (2), and (3) of Panel B in Table IV. For all illiquidity measures, we find that illiquidity is higher when the fraction of arbitrage opportunities that are toxic is higher, as predicted by Implication 1. This effect is both statistically and economically significant. For instance, for the effective bid-ask spread, we find that $b_1 = 0.505$ (t-stat=7.62). Thus, an increase in $\hat{\varphi}$ by one standard deviation (i.e., 0.1) is associated with an increase of 0.0505 basis points for the bid-ask spread (i.e., about a 3% increase in the spread). We also find that the effect of the size of toxic arbitrage opportunities is positive and statistically significant. The effect of the number of arbitrage opportunities per trade is positive but statistically significant only when illiquidity is measured by the effective bid-ask spread. Other control variables have the usual sign. For instance, daily changes in illiquidity are positively associated with realized volatility and negatively associated with trading activity measured by the number of orders.

In the model, market participants are assumed to know the likelihood, φ, that a given arbitrage opportunity is toxic. In Panel A of Table IV, we show that the fraction of arbitrage opportunities that are toxic on day t, $\hat{\varphi}_t$, can be forecast using information from past trading days. Specifically, we estimate a model in which $\hat{\varphi}_t$ depends on its 20 past realizations (to capture persistence in the level of $\hat{\varphi}_t$) and various market characteristics on day $t-1$ for the three currencies in our sample (namely, their average quoted spreads on day $t-1$, their realized volatility on day $t-1$, the number of orders submitted in each currency on day $t-1$, and the average trade size in each currency on day $t-1$). The forecasting model predicts $\hat{\varphi}_t$ well with an adjusted R^2 of 44%. Using this forecasting model, we decompose $\hat{\varphi}_t$ into an anticipated component and an unanticipated component, $v_{\hat{\varphi}_t}$ (the residual of the forecasting equation) and we reestimate equation (19) with both components as explanatory variables. We report estimates in Column (4), (5), and (6) of Table IV (Panel B). Both the anticipated and unanticipated components of $\hat{\varphi}_t$ are positively and significantly associated with all measures of illiquidity.

3.5.2 Is illiquidity higher when arbitrageurs are faster?

We now test whether an increase in the likelihood, π^*, that a toxic arbitrage terminates with an arbitrageur's trade has a positive effect on illiquidity (Implication 2a). To this end, we include π^{tox} (our proxy for π) as an additional explanatory variable in our baseline regression (19). Specifically, we estimate:

$$illiq_{it} = \omega_i + \xi_{t,m} + b_1\hat{\pi}_t^{tox} + b_2\hat{\varphi}_t + b_3\hat{\alpha}_t + b_4\hat{\sigma}_t^{tox} + b_5 vol_{it} + b_6 trsize_{it} + b_7 nrorders_{it} + b_8 illiq_{it}^{EBS} + \varepsilon_{it}. \tag{20}$$

Implication 2a states that b_1 should be significantly positive: illiquidity should be larger when a toxic arbitrage opportunity is more likely to terminate with an arbitrageur's trade. In the model, π is endogenous and simultaneously determined with the bid-ask spread. Hence, we use an instrumental variable (IV) approach to estimate b_1. Implication 2b shows that a shock to arbitrageurs' relative cost of speed (parameter ρ in the model) can serve as an instrument to identify the effect of $\hat{\pi}^{tox}$ on illiquidity. Indeed, in the theory, this shock affects π^* without directly affecting illiquidity (the equilibrium bid-ask spread is related to ρ only through the effect of ρ on π^*; see eq. (10)).

Hence, as an instrument for π, we use a technological shock that affects the speed at which traders can react to arbitrage opportunities on Reuters D-3000. In July 2003, Reuters D-3000 introduced a new functionality, "Reuters AutoQuote API" (Application Programming Interface). This functionality allows traders to automate order submission on their Reuters terminal instead of manually typing trading instructions using a Reuters keyboard, as was done until July 2003. Hence, "Reuters AutoQuote API" marked the beginning of algorithmic trading on Reuters since it allowed traders to input Reuters datafeed in their algorithms and let these trade accordingly.[28] Pierron (2007) emphasizes the interest of APIs for arbitrageurs: "*This allows a full benefit from algorithmic trading, since it enables the black box to route the order to the market with the best prices and potential arbitrage across markets despite the fragmentation of the various pools of liquidity in the FX market.*" Similarly, Chaboud et al. (2014) note that (on p.2058): "*From conversations with market participants, there is widespread anecdotal evidence that in the very first years of algorithmic trading in this [FX] market, a fairly limited number of strategies were implemented with triangular arbitrage among the most prominent.*"

[28] The electronic broker EBS launched a similar service, "EBS Spot Ai", in 2004; see King and Rime (2010) and Chaboud et al. (2014).

Initially, Reuters allowed only a limited number of clients to use the AutoQuote API functionality because of capacity constraints (APIs' users consume more bandwidth than manual users). This does not in itself invalidate our tests because our implications hold even if there is only one fast arbitrageur (as in the baseline model). However, a low usage of Autoquote makes it more difficult for us to detect its effect.

Insert Figure IV about here.

Automated trading activity often generates a higher order-to-trade ratio, i.e., the number of orders to the number of trades in a market (see Hendershott et al. (2011) for instance). Figure IV presents a time series of the daily value of this ratio on Reuters D-3000 for the three currencies in our sample. Dashed lines in this figure indicate the average levels of the ratio before and after July, 1^{st} 2003. There is a clear upward shift in the order-to-trade ratio on Reuters D-3000 in July 2003, consistent with increased automation of trading due to the introduction of Autoquote.

AutoQuote API (henceforth AutoQuote for brevity) enables all traders (arbitrageurs and market makers) to react faster to changes in limit order books and therefore arbitrage opportunities. It should therefore either increase or decrease arbitrageurs' *relative* speed of reaction (ρ in the model). Thus, the sign of the effect of AutoQuote on $\hat{\pi}_t^{tox}$ can be positive or negative. However, as long as this effect exists (i.e., AutoQuote affects $\hat{\pi}_t^{tox}$), the introduction of AutoQuote on Reuters D-3000 can be used as an instrument to test whether the effect (b_1) of $\hat{\pi}^{tox}$ on illiquidity is positive. A second condition for the introduction of AutoQuote to be a valid instrument is that it satisfies the exclusion restriction, i.e., the introduction of AutoQuote should not be correlated with the error term in eq.(20). In other words, the introduction of AutoQuote should affect illiquidity only through its effect on $\hat{\pi}^{tox}$ (after controlling for other variables appearing in eq.(20)). This assumption is plausible because, to our knowledge, there is no known mechanism by which a technological change such as AutoQuote API should affect the usual determinants of bid-ask spreads, namely, order processing costs, inventory holding costs, or adverse selection costs due to private information.

The first stage of the IV is:

$$\hat{\pi}_t^{tox} = \omega_i + \xi_{t,m} + a_1 AD_t + a_2\hat{\varphi} + a_3\hat{\alpha}_t + a_4\hat{\sigma}_t^{tox} + a_5 vol_{it} + a_6 trsize_{it} + a_7 nrorders_{it} + a_8 illiq_{it}^{EBS} + u_t, \tag{21}$$

where AD_t is our instrument (a dummy equal to 1 after July 2003 and zero before). Estimates for this regression are reported in Columns (1), (3), and (5) of Table V Panel A.[29] We find that the introduction of AutoQuote on Reuters D-3000 has a significant positive effect on the likelihood that an arbitrage opportunity terminates with a trade rather than a quote update. The coefficient (a_1) on the dummy variable, AD_t, is equal to 0.04 and statistically significant. Thus, the likelihood that a toxic arbitrage opportunity terminates by an arbitrageur's trade ($\hat{\pi}^{tox}$) increases by about 4% after July 2003. The instrument is not weak since the F statistics is around 18 in all specifications. Overall, the first stage regression indicates that AutoQuote increased arbitrageurs' relative speed of reaction to toxic arbitrage opportunities in the currency pairs in our sample. This is consistent with the view that algorithmic trading in the FX market is, at least at the time of our sample, predominantly used to exploit triangular arbitrage opportunities (see Chaboud et al. (2014))

Estimates for the second stage of the IV are reported in Columns (2), (4), and (6) of Table V. As in Table IV, for all illiquidity measures, we find a positive and statistically significant relation between the likelihood that an arbitrage opportunity is toxic ($\hat{\varphi}$) and illiquidity. The size of the effect of $\hat{\varphi}$ on illiquidity is not different from those reported in Table IV.

As predicted, the effect of $\hat{\pi}^{tox}$ on illiquidity, b_1, is positive and statistically significant at the 1% level for all measures of illiquidity. For instance, a 1% increase in $\hat{\pi}_t^{tox}$ raises the quoted bid-ask spread by 0.07746 basis points (t-stat = 4.04).[30] The economic size of this effect is significant as well since 0.07746 basis points represents about 4% of the average bid-ask spread (about 2 basis points) for currencies in our sample. Hence the effect of AutoQuote on $\hat{\pi}_t^{tox}$ (4% on average) raised bid-ask spreads by about 16% for currency pairs in our sample.

Another way to evaluate the economic significance of this finding is to consider the effect of an increase of 1% in $\hat{\pi}^{tox}$ on daily trading costs in the three currencies in our sample. The average trade sizes for the currencies considered in our sample are: \$2.390 million in GBP/USD (dollar value of £1.386 million), \$1.655 million in EUR/USD (dollar value of €1.401), and \$1.831 million in EUR/GBP (dollar value of €1.548) (see Table II). Moreover, the average number

[29] The first stage regression is slightly different for each illiquidity measure because one of the control (the illiquidity of EBS) varies with the illiquidity measure. Estimates of coefficients for the first stage are very similar across all illiquidity measures, however.

[30] We have also estimated eq. (20) with OLS. In this case, we find that $\hat{\pi}^{tox}$ is positively related with illiquidity but the relationship is not statistically significant. This suggests that the endogeneity bias for b_1 is negative. This direction is consistent with intuition: arbitrageurs are likely to pay more attention to arbitrage opportunities (and therefore be faster) on days in which bid-ask spreads are smaller because they can earn larger profits on these days. This effect should partially offset the true effect of an exogenous increase in arbitrageurs' relative speed on spreads.

of transactions per day is 4,692 in GBP/USD, 2,365 in EUR/USD, and 2,841 in EUR/GBP. Hence, according to our estimates, a 1% increase in $\hat{\pi}^{tox}$ raises total daily trading costs by $0.07746bps \times (\$2.390 \times 4,692 + \$1.655 \times 2,365 + \$1.831 \times 2,841) = \$157,475$ for the three markets in total or about $40 million per year. Thus, even a small increase in arbitrageurs' speed can be rather costly for other market participants.

As a robustness check, we have also estimated eq. (20) using hourly estimates of each variable in our regressions rather than daily estimates. The findings in this case are qualitatively very similar and, in economic terms, they are stronger. For instance, an increase of 1% in $\hat{\pi}^{tox}$ triggers an increase of 0.08 basis points for the effective spread when we estimate (20) at the hourly frequency vs. 0.034 basis points at the daily frequency. For brevity, we do not report the estimates obtained with tests at the hourly frequency. They are available upon request.

3.5.3 Time-to-efficiency

We now test our auxiliary predictions regarding the duration of arbitrage opportunities (Implications 3 and 4). Implication 3 states that a decrease in the cost of being fast for arbitrageurs and market makers should reduce the duration of arbitrage opportunities even if the decrease is relatively larger for arbitrageurs (so that $\hat{\pi}^{tox}$ and therefore illiquidity increase). Hence, we should observe a negative effect of AutoQuote on the duration of arbitrage opportunities even though it raises trading costs (see previous section). Implication 4 further predicts that the duration of arbitrage opportunities should be shorter on days in which the fraction of arbitrage opportunities that are toxic is higher.

We test these two predictions by estimating the following regression:

$$\log(TTE_t) = \omega_i + \xi_{t,m} + a_1 AD_t + a_2 \hat{\varphi}_t + a_3 \hat{\alpha}_t + a_4 \hat{\sigma}_t^{tox} + a_5 vol_{it} + a_6 trsize_{it} + a_7 nrorders_{it} + u_t, \quad (22)$$

where TTE_t is the average duration of arbitrage opportunities on day t. Estimates are reported in Table VI. In Column 1, the dependent variable is the duration of toxic arbitrage opportunities while in Column 2 the dependent variable is the duration of all arbitrage opportunities.[31]

Consistent with Implication 3, AutoQuote API is associated with a decrease in the duration of toxic arbitrage opportunities by about 6.5%.[32] Similar estimates are obtained when we use

[31] The expected duration of toxic arbitrage opportunities is $E(\text{Min}\{D^a, D^m\}) = \frac{\rho}{\gamma^*(S_X^*)(1+r)}$ in the model. It is easily checked that Implications 3 and 4 also hold for $E(\text{Min}\{D^a, D^m\})$.

[32] As the dependent variable is $\log(TTE_t)$, a_1 measures the percentage change in time-to-efficiency after the

the average duration of all arbitrage opportunities. We also find that, on average, the duration of toxic arbitrage opportunities is shorter when the fraction of arbitrage opportunities that are toxic, $\hat{\varphi}$, is higher, which is consistent with Implication 4. In contrast, the effect of $\hat{\varphi}$ on the duration of all arbitrage opportunities is statistically insignificant. This might reflect the fact that the duration of *non toxic* arbitrage opportunities are driven by factors that are not captured by our model, which mainly focuses on toxic arbitrage opportunities.

In sum, faster arbitrage can enhance pricing efficiency at the expense of market liquidity. The introduction of AutoQuote (algorithmic trading) or an increase in the fraction of arbitrage opportunities that are toxic induce arbitrageurs (and market makers) to correct arbitrage opportunities more quickly. This effect reduces the duration of arbitrage opportunities and therefore improves pricing efficiency. Yet, faster arbitrageurs raise market makers' exposure to the risk of being picked off, which impairs market liquidity.

4 Additional Tests

4.1 Heterogeneous effects across currencies

The likelihood that one currency pair initiates a toxic arbitrage opportunity is not equal across all currency pairs: 51% of all toxic arbitrage opportunities can be traced back to a jump in the EUR/USD rate while the GBP/USD and EUR/GBP pairs initiate only 28.68% and 20.32% of all toxic arbitrage opportunities, respectively. These findings suggest that the EUR/USD pair leads other currency pairs in our sample in terms of price discovery (i.e., shocks to fundamentals are reflected first the EUR/USD market). Thus, the EUR/USD pair frequently plays the role of asset Y in our model while the two other pairs (GBP/USD and EUR/GBP) play the role of asset X. Hence, we expect market makers in the EUR/USD pair to be less exposed to toxic arbitrage trades than in other currencies. Accordingly, illiquidity in this pair should be less sensitive to (i) the likelihood of occurrence of toxic arbitrage opportunities ($\hat{\varphi}$) than the other pairs and (ii) the likelihood that market makers are picked off when a toxic arbitrage opportunity occurs ($\hat{\pi}^{tox}$).

To test these additional predictions, we re-run our IV regression for each individual currency pair separately. That is, we allow the effect of AutoQuote to be currency specific. Panels A,

introduction of AutoQuote. The average time-to-efficiency in toxic arbitrage opportunities is 0.89 seconds in our sample. Hence, AutoQuote coincides with a reduction of about 62 milliseconds in the duration of toxic arbitrage opportunities.

B, and C of Table VII present the results for each currency pair in our sample. Realizations of control variables in the first stage and the second stage are specific to each currency pair (e.g., the average trade size in a given day is specific to each currency). Yet, consistent with our earlier results, we find that AutoQuote is associated with a significant increase in $\hat{\pi}^{tox}$, which is roughly the same across all currencies.

As expected, the fraction of arbitrage opportunities that are toxic ($\hat{\varphi}$) still has a positive and significant effect (at the 5% level) on the illiquidity of GBP/USD and EUR/GBP but it has only a mildly positive significant effect on the quoted spread of the EUR/USD pair (and no significant effect on other illiquidity measures in this pair). We also find that the effect of $\hat{\pi}^{tox}$ on the quoted bid-ask spread is much weaker for the EUR/USD pair than in other currencies. In contrast, and surprisingly, the effect of $\hat{\pi}^{tox}$ on other measures of illiquidity for the EUR/USD pair seems sometimes stronger than in other currency pairs.

Insert Tables VII about here

4.2 Exposure to toxic arbitrage trades or other forms of adverse selection?

By picking off stale quotes, arbitrageurs expose market-makers to adverse selection. This form of adverse selection is similar to that highlighted in other models of market making with informed investors (e.g., Copeland and Galai (1983)). An important difference, however, is that arbitrageurs' advantage does not stem from private information or a superior ability to process existing information. Rather, their profit only stems from speed: a quicker reaction than other market participants to publicly available and easy to process information (in the case of our paper, the existence of an arbitrage opportunity).

A natural question is whether our measures of market makers' exposure to arbitrageurs' picking off risk ($\hat{\varphi}$ and $\hat{\pi}^{tox}$) are distinct from other measures of adverse selection. We consider two alternative measures. First, the immediate period following a macro-economic announcement is often associated with an increase in informational asymmetries because some market participants are better at processing information. For instance, Green (2004) find that the informational content of trades in treasury bond markets increases in the few seconds following scheduled macro-economic announcements. Accordingly, market makers face larger adverse selection costs just after macro-economic announcements. This effect is naturally stronger when macro-economic announcements are more surprising (that is, differ more from traders' forecasts).

If $\hat{\varphi}$ and $\hat{\pi}^{tox}$ proxies for informational asymmetries associated with macro-economic announcements, we would expect their effects on illiquidity to be weaker when we control for surprises in macro-economic announcements. To test whether this is the case, we use data from Money Market Survey (MMS), provided by InformaGM, to construct macroeconomic announcement surprises in the different geographical areas (EMU, U.K., and U.S.) relevant for our currency pairs.

The MMS data provide median forecasts of all macro-economic announcements by market participants (collected on the Thursday prior to the announcement week) and their actual realization on the day of the announcement. Announcement surprises are measured as the realized announced value minus the median forecast. Following Andersen et al. (2003), we standardize announcements surprises by their standard deviation. Specifically, the surprise $N_{k\tau}$ of announcement type k (e.g., nonfarm payroll, CPI, unemployment, etc.) on day τ is,

$$N_{k\tau} = \frac{A_{k\tau} - F_{k\tau}}{\sigma_k},$$

where $A_{k\tau}$ and $F_{k\tau}$ are the actual announcement value and median forecast of this value, respectively (σ_k is the standard deviation of $A_{k\tau} - F_{k\tau}$). For each area, we build on each day a macro-economic announcement variable (namely $macro_t^{EMU}$, $macro_t^{UK}$, and $macro_t^{US}$) equal to the sum of all macro-economic announcements surprises in this area. Macro-economic announcements in at least one geographical area are frequent in our sample so that there are only 102 days *without* any macro-announcements.

Easley et al. (2011) and Easley et al. (2012) advocate the use of VPIN ("volume-synchronised probability of informed trading") as a measure of high-frequency order flow toxicity (adverse selection). Following Easley et al. (2012), in each trading day t, we group successive trades into 50 equal volume buckets of size V_t^i, where V_t^i is equal to the trading volume in day t for currency i divided by 50. For robustness, we have also grouped trades into different volume buckets of 30, 75 and 100 in the construction of VPIN. Results reported below are not sensitive to the size of grouping. The $VPIN_t^i$ metric for currency pair i and day t is then

$$VPIN_t^i = \frac{\sum_{\tau=1}^{50} \left| V_\tau^{i,S} - V_\tau^{i,B} \right|}{50 \times V_t^i},$$

where $V_\tau^{i,B}$ and $V_\tau^{i,S}$ are the amount of base currency purchased and sold, respectively, within the τ^{th} bucket for currency pair i.[33] We observe a significant a positive correlation among the VPIN measures for the three currency pairs in our sample. In contrast, the correlation between $\hat{\pi}^{tox}$ and the VPIN of each currency are much lower and significantly different from zero only for the EUR/USD (correlation equal to 0.07) and the EUR/GBP (correlation equal to 0.09) pairs. This already suggests that $\hat{\pi}^{tox}$ and $VPIN$ do not capture the same information about market makers' exposure to adverse selection.

Table VIII reports the result of the IV regression when we control for surprises in macroeconomic announcements ($macro^{EMU}$, $macro^{UK}$, $macro^{US}$), and $VPIN$ for each currency. Consistent with Green (2004), we find that macroeconomic surprises are positively associated with illiquidity. Effects of surprises on our measures of illiquidity, however, are only marginally significant (at the 10% level).[34] We also find a positive and marginally significant relation between VPIN and quoted or effective bid-ask spreads. However, and more importantly, the effects of $\hat{\varphi}$ and $\hat{\pi}^{tox}$ on illiquidity remain positive and significant for all measures of illiquidity. Furthermore, estimates of the effect of these variables are very similar to those reported in Table V. Hence, $\hat{\varphi}$ and $\hat{\pi}^{tox}$ contain information about market makers' exposure to adverse selection, that is not captured by VPIN and not associated with informational asymmetries around macro-economic announcements.

Insert Table VIII about here

5 Conclusions

Short lived arbitrage opportunities arise when the prices of asset pairs do not adjust to information at the same speed. These opportunities are toxic because they expose liquidity suppliers to the risk of trading with arbitrageurs at stale quotes (the risk of "being picked off"). Hence, more frequent toxic arbitrage opportunities and a faster arbitrageurs' response to these opportunities can impair liquidity and therefore investors' welfare. We provide evidence for these effects using a sample of triangular arbitrage opportunities. Specifically we find that bid-ask spreads for the currency pairs in our sample are larger on days in which (a) the fraction of

[33]Computation of VPIN requires classifying market orders into two groups: buys and sales so that $V_\tau^{i,B}$ and $V_\tau^{i,S}$ can be computed. This is straightforward with our data because we observe whether a market order is a buy or a sale. Hence, we do not need to infer the direction of market orders from price changes as in Easley et al. (2012).

[34]Informational asymmetries created by surprises in macroeconomic announcements quickly fade away. Hence, it is difficult to detect their effect using daily measures of market illiquidity.

arbitrage opportunities that are toxic is higher and (b) the frequency with which arbitrageurs successfully exploit these opportunities is higher.

These empirical findings provide support for proposals to curb high frequency traders' ability to exploit short lived arbitrage opportunities. To this end, Budish et al. (2014) propose to use frequent batch auctions (say, one per second) instead of continuous limit order books. In effect, this change in market design should considerably weaken arbitrageurs' ability to pick off liquidity providers' quotes (this is similar to a sharp decrease in π in our model). Another possibility is to delay the execution of marketable orders by a small random delay (of the order of a few milliseconds).[35] This approach also reduces the likelihood of trading at stale quotes with arbitrageurs for liquidity suppliers.

Our findings imply that the effect on liquidity of limiting high frequency arbitrageurs' speed depends on the composition of arbitrage opportunities in a given asset pair (φ in our model). When most arbitrage opportunities are due to price pressures effects, "speed bumps" for arbitrageurs should not have much impact on liquidity. In contrast, when a large fraction of arbitrage opportunities are toxic, they should significantly improve liquidity. This implication is important for at least two reasons. First, it suggests that analyzing systematically how the composition of arbitrage opportunities varies *across* asset pairs is necessary to identify markets where limits on arbitrageurs' speed are most likely to be effective. Second, empiricists should expect to find strong positive effects of these limits on liquidity only when they apply to asset pairs for which toxic arbitrage opportunities are relatively frequent.

[35] This approach has been implemented in some foreign exchange trading platforms, in particular EBS, Reuters, and ParFX, precisely on the ground that it would limit liquidity suppliers exposure to the risk of being picked off by arbitrageurs. See '*'Life in the slow lane*," Automated Trader, 30, 2013 or "*EBS takes new steps to rein in high frequency traders*," Reuters, August 23, 2013. It is also advocated by some participants in U.S. equity markets (see "Interactive Brokers Group Proposal to Address High Frequency Trading" available at: `https://www.interactivebrokers.com/download/SEC_proposal_high_frequency_trading.pdf`.

Appendix

Derivation of Implications 3 and 4

First, substituting the expression for the equilibrium spread, S^*_X (given in (10)) in (8), we obtain the arbitrageur's speed, $\gamma^*(S^*_X)$, in equilibrium:

$$\gamma^*(S^*_X) = \frac{(\varphi\alpha\rho)(1-\alpha(2\varphi-1))}{(c^a(1+\rho)^2)(2\varphi\alpha\pi^*(\rho)(2-\pi^*(\rho))+(1-\alpha(2\varphi-1)))}\sigma \tag{23}$$

Substituting (23) in (11), we obtain:

$$E(D) = \frac{(2\varphi\alpha(\frac{2+\rho}{1+\rho})+(1-\alpha(2\varphi-1))(\frac{2+\rho}{\rho}))(c^a(1-\varphi)+c^m)}{\varphi\alpha(1-\alpha(2\varphi-1))\sigma} \tag{24}$$

We deduce that $\frac{\partial E(D)}{\partial\rho} > 0$. Using the fact that $\rho = \frac{c^m}{c^a}$, we deduce from (24) that a decrease in c^a and c^m that eventually result in an increase in ρ lowers the expected duration of an arbitrage opportunity (Implication 3).

Using (23), we also deduce $\gamma^*(S^*_X)$ increases with φ if $\alpha(4\varphi-1) < 1$. This condition is automatically satisfied when $\varphi \leq 1/2$ or when $\varphi > 1/2$ and $\alpha < (4\varphi-1)^{-1}$. We deduce from (11) that the average duration of an arbitrage opportunity decreases with φ.

Figures

Figure I: Testable Implications

This figure shows the equilibrium bid-ask spread (in bps) as (a) a function of the likelihood of a toxic arbitrage opportunity, φ (Panel A), (b) the likelihood that an arbitrageur trades when a toxic arbitrage opportunity occurs, π^* (Panel B), (c) the relative cost of speed for the market maker ρ (Panel C). It also shows the duration of arbitrage opportunities (in seconds) as a function of (i) the cost of speed for the arbitrageur, c^a (Panel D) and (ii) the likelihood of a toxic arbitrage opportunity, φ (Panel D). In all cases, we set $\sigma = 3.5bps$ and $\alpha = 0.1$.

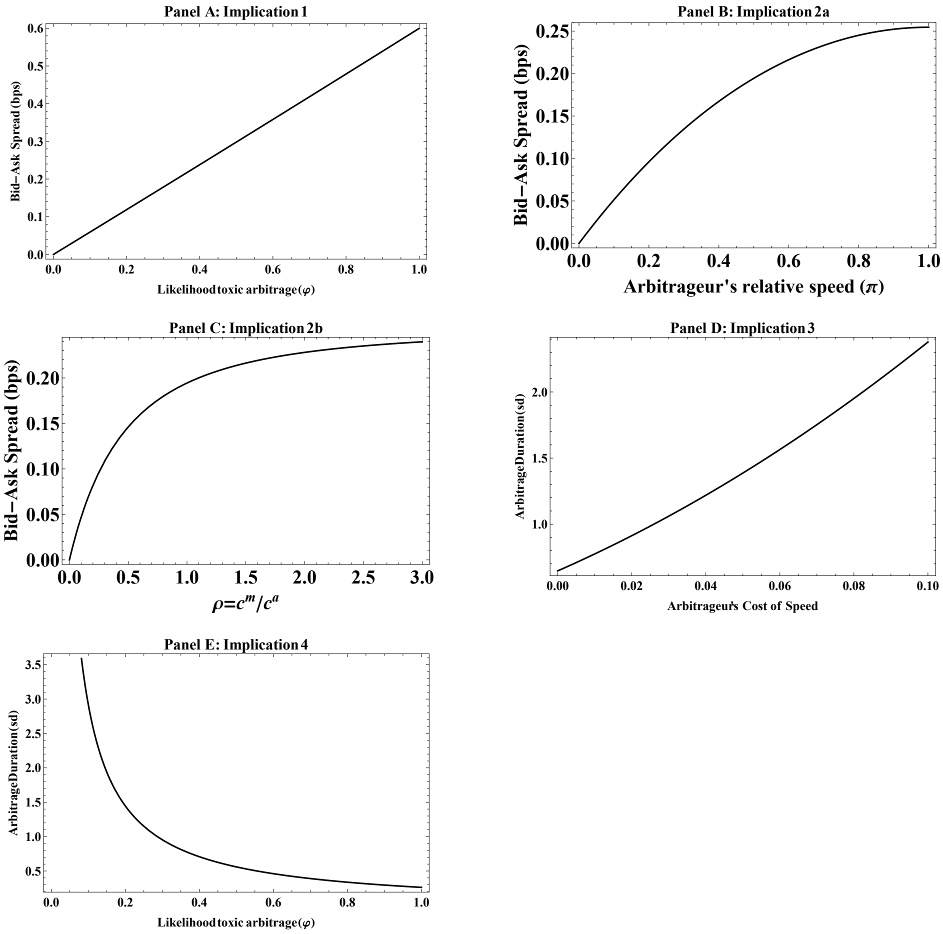

Figure II: Toxic vs. Non-Toxic Arbitrage Opportunities

This figure shows how we classify triangular arbitrage opportunities into toxic and non-toxic opportunities using four triangular arbitrage opportunities that occurred in our sample. In each panel, the arbitrage opportunity starts at time t and ends at time $t + \tau$. The solid line shows the evolution of best ask and bid prices in the currency pair that initiates the arbitrage opportunity. The dashed lines show the evolution of best bid and ask synthetic quotes. In Panel A, we provide two examples of opportunities that we classify as toxic because they are associated with permanent shifts in exchange rates. In Panel B, we provide two examples of opportunities that we classify as non-toxic because the exchange rate in the currency pair initiating the arbitrage opportunity eventually reverts to its level at the beginning of the opportunity.

Panel A: Toxic Arbitrage Opportunities

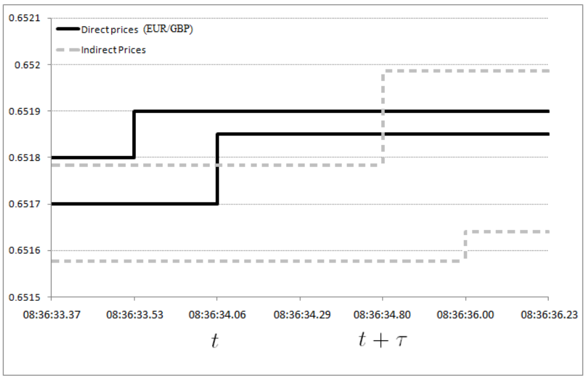

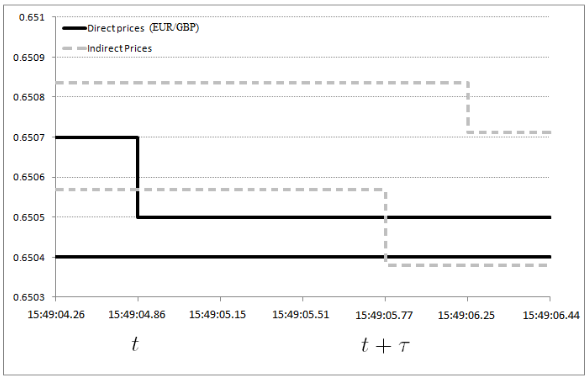

Panel B: Non-Toxic Arbitrage Opportunities

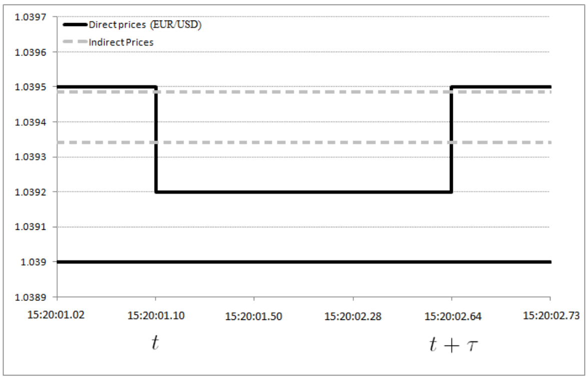

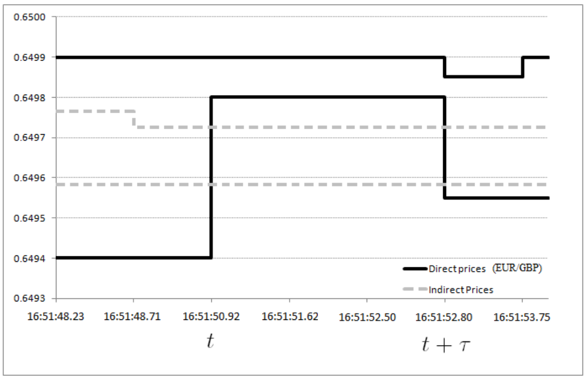

Figure III: Number of Arbitrage Opportunities

Panel A shows the time series of the daily number of all triangular arbitrage opportunities (grey line) and toxic arbitrage opportunities (black line) in our sample. Panel B shows the intra-day pattern of toxic and non-toxic arbitrage opportunities in our sample. Time is GMT.

Panel A: Daily Numbers of Arbitrage Opportunities

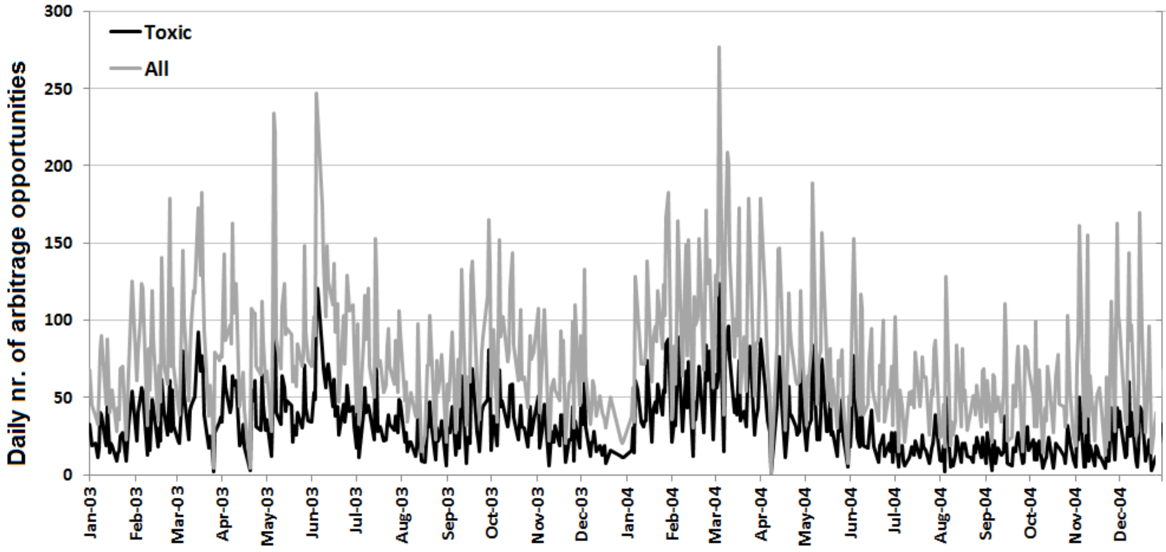

Panel B: Intraday Pattern in the Number of Arbitrage Opportunities

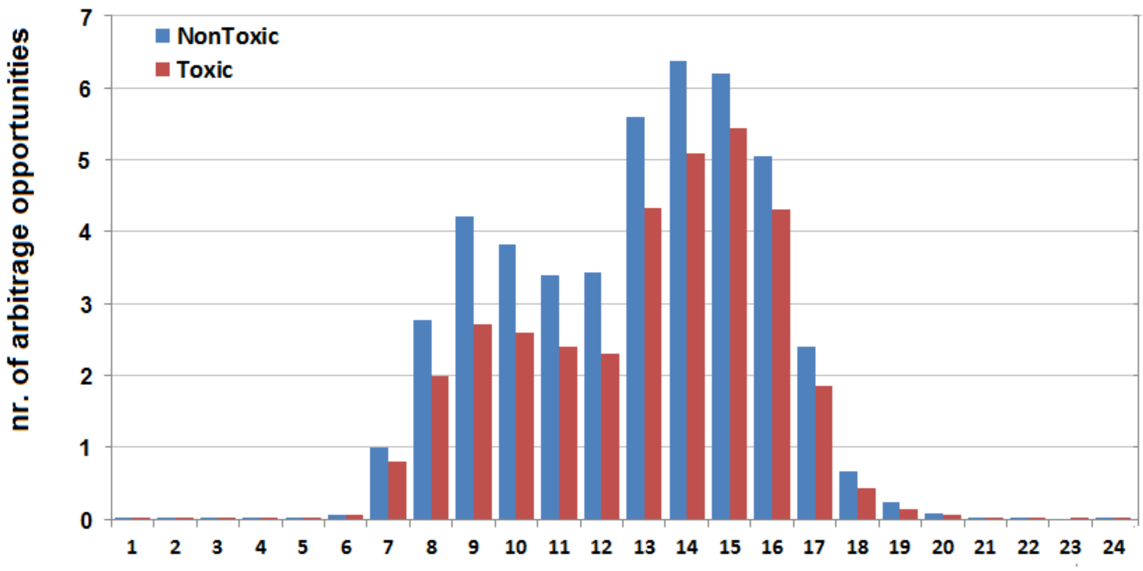

Figure IV: AutoQuote and Order to Trade Ratio.

This figure shows the evolution of the order-to-trade ratio (defined as the daily number of orders to the daily number of trades for the three currency pairs in our sample) from January 2003 to December 2004. The dashed lines indicate the average levels of the order to trade ratio before and after July, 1^{st} 2003.

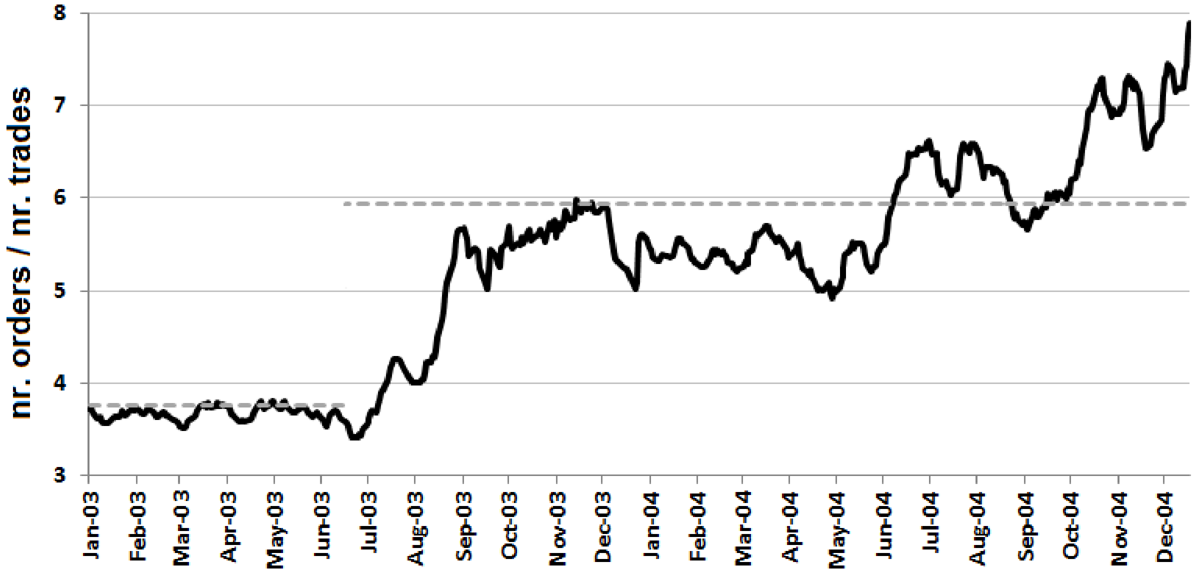

Tables

Table II: Descriptive Statistics

This table presents the descriptive statistics for the variables used in our tests for each currency pair $i \in \{GU, EU, EG\}$, where indexes GU, EU, and EG refer to the GBP/USD, EUR/USD, and EUR/GBP currency pairs, respectively. In Panel A (Panel B), we present descriptive statistics for variables that are specific to our set of toxic (non-toxic) arbitrage opportunities. TTE_t^{tox} (TTE_t^{nontox}) denotes the duration in seconds of toxic (non-toxic) arbitrage opportunities on day t; $nrarb_t^{tox}$ ($nrarb_t^{nontox}$) is the number of toxic (non-toxic) arbitrage opportunities on day t; $\hat{\pi}_t^{tox}$ ($\hat{\pi}_t^{nontox}$) is the number of toxic (non-toxic) arbitrage opportunities on day t that terminate with a trade divided by the total number of toxic (non-toxic) arbitrage opportunities on this day; $\hat{\varphi}_t$ is the number of toxic (non-toxic) arbitrage opportunities on day t divided by the number of arbitrage opportunities on this day; $\hat{\sigma}_t^{tox}$ ($\hat{\sigma}_t^{nontox}$) is the average size of toxic (non-toxic) arbitrage opportunities on day t (in basis points); $profit_t^{tox}$ ($profit_t^{nontox}$) is the average profit in basis points on toxic (non-toxic) triangular arbitrage opportunities on day t (calculated as explained in Section 3.2.1). The *t-stat* column reports t-statistics for significance of the mean differences of the variables computed using toxic and non-toxic arbitrage opportunities. Panel C presents descriptive statistics for the illiquidity measures (all expressed in basis points) used in our tests: $spread_{it}$ is the average quoted bid-ask spread in currency pair i on day t; $espread_{it}$ is the average effective spreads in currency pair i on day t; $slope_{it}$ is the slope of the limit order book in currency pair i on day t; superscript EBS is used when these variables are measured using EBS data. Panel D presents summary statistics for control variables used in our tests: $\hat{\alpha}_t$ is the number of all arbitrage opportunities on day t divided by the total number of trades on this day; vol_{it} is the realized volatility (in percentage) of 5-minutes returns for currency pair i on day t; $nrorders_{it}$ (in thousands) is the total number of orders (market, limit, or cancelations) in currency pair i on day t; $trsize_{it}$ is the average daily trade size (in million) for currency pair i on day t; $nrtr_{it}$ is the daily number of trades (in thousands) in currency pair i on day t. The sample period is from January 2, 2003 to December 30, 2004.

Variable	Mean	Std.Dev.	Min	Q1	Median	Q3	Max	t-stat
Panel A: Toxic Arbitrage								
TTE_t^{tox}	0.894	0.301	0.262	0.725	0.847	1.006	4.060	
$nrarb_t^{tox}$	32.01	20.83	1.000	17.00	28.00	43.00	124.0	
$\hat{\pi}_t^{tox}$	0.741	0.110	0.000	0.685	0.743	0.804	1.000	
$\hat{\varphi}_t$	0.415	0.100	0.080	0.361	0.431	0.483	0.691	
$\hat{\sigma}_t^{tox}$	3.535	0.757	2.224	3.112	3.439	3.843	13.61	
$profit_t^{tox}$	1.427	0.277	1.115	1.336	1.401	1.470	6.668	
Panel B: Non-toxic Arbitrage								
TTE_t^{nontox}	0.518	0.199	0.025	0.389	0.485	0.611	1.899	23.2
$nrarb_t^{nontox}$	45.22	38.40	2.000	27.00	40.00	55.00	740.0	-6.75
$\hat{\pi}_t^{nontox}$	0.807	0.082	0.412	0.755	0.800	0.867	1.000	-10.7
$\hat{\varphi}_t^{nontox}$	0.585	0.094	0.309	0.517	0.569	0.639	0.920	-29.0
$\hat{\sigma}_t^{nontox}$	3.531	0.841	2.404	3.058	3.350	3.747	9.662	0.08
$profit_t^{nontox}$	1.618	0.571	1.218	1.417	1.512	1.610	7.280	-6.72

Table II continued

Variable	Mean	Std.Dev.	Min	Q1	Median	Q3	Max
Panel C: Illiquidity Measures							
$spread_{GUt}$	2.741	0.309	2.089	2.523	2.725	2.937	5.258
$spread_{EUt}$	2.530	0.509	1.572	2.160	2.458	2.800	5.281
$spread_{EGt}$	1.352	0.259	0.922	1.184	1.331	1.473	4.421
$espread_{GUt}$	2.073	0.255	1.578	1.904	2.045	2.205	3.784
$espread_{EUt}$	1.886	0.459	1.152	1.593	1.812	2.080	5.815
$espread_{EGt}$	0.966	0.180	0.671	0.841	0.945	1.052	2.838
$slope_{GUt}$	1.120	0.162	0.774	1.011	1.109	1.217	2.635
$slope_{EUt}$	1.111	0.275	0.494	0.928	1.088	1.266	2.493
$slope_{EGt}$	0.541	0.132	0.312	0.455	0.524	0.604	1.605
$spread^{EBS}_{GUt}$	5.253	1.157	2.421	4.687	5.093	5.580	12.82
$spread^{EBS}_{EUt}$	1.139	0.046	1.050	1.103	1.136	1.164	1.336
$spread^{EBS}_{EGt}$	2.520	0.807	1.431	2.064	2.376	2.803	11.11
$espread^{EBS}_{GUt}$	5.112	3.467	2.375	3.237	4.165	5.796	27.90
$espread^{EBS}_{EUt}$	0.998	0.065	0.899	0.958	0.985	1.020	1.420
$espread^{EBS}_{EGt}$	2.082	1.847	0.901	1.245	1.589	2.189	24.23
$slope^{EBS}_{GUt}$	3.860	3.246	1.125	2.377	3.122	4.332	47.97
$slope^{EBS}_{EUt}$	0.296	0.041	0.205	0.266	0.294	0.323	0.441
$slope^{EBS}_{EGt}$	1.833	2.448	0.444	1.011	1.428	1.980	39.47
Panel D: Control Variables							
$\hat{\alpha}_t$	0.022	0.016	0.005	0.015	0.020	0.027	0.323
vol_{GUt}	0.683	0.268	0.117	0.532	0.622	0.753	2.456
vol_{EUt}	0.827	0.386	0.258	0.616	0.744	0.920	4.363
vol_{EGt}	0.387	0.094	0.203	0.325	0.381	0.440	1.256
$nrorders_{GUt}$	17.65	6.091	0.576	12.51	17.62	22.45	32.22
$nrorders_{EUt}$	19.05	6.831	0.188	14.78	18.15	22.88	44.07
$nrorders_{EGt}$	14.77	5.810	0.307	9.326	16.07	19.41	28.93
$trsize_{GUt}$	1.386	0.043	1.247	1.357	1.382	1.415	1.509
$trsize_{EUt}$	1.401	0.056	1.000	1.365	1.396	1.434	1.605
$trsize_{EGt}$	1.548	0.076	1.294	1.497	1.541	1.591	1.853
$nrtr_{GUt}$	4.692	1.505	0.175	3.634	4.523	5.639	9.611
$nrtr_{EUt}$	2.365	0.707	0.027	1.859	2.377	2.870	4.103
$nrtr_{EGt}$	2.841	0.811	0.068	2.301	2.761	3.318	6.329
Obs.				498			

Table III: Correlations

This table presents correlations between the variables used in our tests. Indexes GU, EU, and EG refer to the GBP/USD, EUR/USD, and EUR/GBP currency pairs, respectively. TTE_t^{tox} denotes the duration in seconds of toxic arbitrage opportunities on day t; $\hat{\pi}_t^{tox}$ (resp., $\hat{\pi}_t^{nontox}$) is the number of toxic arbitrage opportunities on day t that terminate with a trade divided by the total number of toxic (non-toxic) arbitrage opportunities on day t; $\hat{\varphi}_t$ is the number of toxic (resp., non-toxic) arbitrage opportunities in day t divided by the number of arbitrage opportunities on that day; $\hat{\sigma}_t^{tox}$ is the average size of toxic arbitrage opportunities on day t (in basis points); $spread_{it}$ is the average quoted bid-ask spread (in basis points) in currency pair i on day t; $espread_{it}$ is the average effective spreads (in basis points) in currency pair i on day t; $slope_{it}$ is the average slope of the limit order book in currency pair i on day t. Superscript EBS is used for illiquidity measures computed using EBS data. $illiq^{EBSt}$ reports the correlation between the Reuters and EBS illiquidity variables. The sample period is from January 2, 2003 to December 30, 2004. Bold values are significant at 5% level.

	TTE_t^{tox}	$\hat{\sigma}_t^{tox}$	$\hat{\varphi}_t$	$spread_{GUt}$	$spread_{EUt}$	$spread_{EGt}$	$slope_{GUt}$	$slope_{EUt}$	$slope_{EGt}$	$espread_{GUt}$	$espread_{EUt}$	$espread_{EGt}$	$\hat{\pi}_t^{nontox}$
$\hat{\pi}_t^{tox}$	0.067	-0.023	0.009	0.051	-0.020	0.076	0.070	-0.021	0.083	0.043	0.041	0.064	0.087
TTE_t^{tox}	1.000	**0.395**	**0.175**	**0.179**	**0.119**	**0.293**	**0.268**	0.079	**0.281**	**0.112**	**0.172**	**0.254**	-0.075
$\hat{\sigma}_t^{tox}$		1.000	**0.226**	**0.567**	**0.457**	**0.639**	**0.582**	**0.451**	**0.558**	**0.574**	**0.603**	**0.624**	**-0.174**
$\hat{\varphi}_t$			1.000	**0.271**	**0.263**	**0.289**	**0.257**	**0.274**	**0.288**	**0.348**	**0.338**	**0.345**	**-0.176**
$spread_{GUt}$				1.000	**0.647**	**0.741**	**0.953**	**0.700**	**0.671**	**0.919**	**0.703**	**0.723**	**-0.287**
$spread_{EUt}$					1.000	**0.512**	**0.607**	**0.925**	**0.491**	**0.668**	**0.909**	**0.527**	**-0.314**
$spread_{EGt}$						1.000	**0.779**	**0.529**	**0.955**	**0.680**	**0.685**	**0.964**	**-0.192**
$slope_{GUt}$							1.000	**0.632**	**0.731**	**0.853**	**0.675**	**0.755**	**-0.289**
$slope_{EUt}$								1.000	**0.495**	**0.695**	**0.868**	**0.524**	**-0.277**
$slope_{EGt}$									1.000	**0.611**	**0.636**	**0.918**	**-0.186**
$espread_{GUt}$										1.000	**0.754**	**0.717**	**-0.266**
$espread_{EUt}$											1.000	**0.716**	**-0.301**
$espread_{EGt}$												1.000	-0.196
$illiq^{EBSt}$				**0.273**	**0.735**	**0.414**	**0.174**	**0.814**	**0.233**	**0.314**	**0.594**	**0.521**	

Table IV: The Arbitrage Mix (φ) and Market Illiquidity

In Panel A, we report OLS estimates of the following regression:

$$\begin{aligned} \hat{\varphi}_t &= b_0 + b_1 spread_{GU,t-1} + b_2 spread_{EU,t-1} + b_3 spread_{GE,t-1} + b_4 nrorders_{GU,t-1} + b_5 nrorders_{EU,t-1} \\ &+ b_6 nrorders_{GE,t-1} + b_7 vol_{GU,t-1} + b_8 vol_{EU,t-1} + b_9 vol_{GE,t-1} + b_{10} trsize_{GU,t-1} + b_{11} trsize_{EU,t-1} \\ &+ b_{12} trsize_{GE,t-1} + \sum_{j=1}^{20} c_j \hat{\varphi}_{t-j} + v_{\hat{\varphi},t}, \end{aligned}$$

where indexes GU, EU, EG refer to the GBP/USD, EUR/USD, and EUR/GBP currency pairs, respectively and $\hat{\varphi}_t$ is the number of toxic arbitrage opportunities on day t divided by the number of arbitrage opportunities on this day. In panel B, we report OLS estimates of the following two regressions:
$illiq_{it} = \omega_i + \xi_{t,m} + b_1\hat{\varphi}_t + b_2\hat{\alpha}_t + b_3\hat{\sigma}_t^{tox} + b_4 vol_{it} + b_5 trsize_{it} + b_6 nrorders_{it} + b_6 illiq_{it}^{EBS} + \varepsilon_{it}$ and
$illiq_{it} = \omega_i + \xi_{t,m} + c_1 v_{\hat{\varphi},t} + c_2 fitted_{\hat{\varphi},t} + c_3\hat{\alpha}_t + c_4\hat{\sigma}_t^{tox} + c_5 vol_{it} + c_6 trsize_{it} + c_7 nrorders_{it} + c_8 illiq_{it}^{EBS} + \varepsilon_{it}$, where $illiq_{it}$ is one of our three proxies for illiquidity for currency i on day t: $spread_{it}$ is the average quoted bid-ask spread (in basis points) in currency pair i on day t; $espread_{it}$ (in basis points) is the average effective spreads in currency pair i on day t; $slope_{it}$ is the average slope of the limit order book in currency pair i on day t; superscript EBS is used for measures of these variables computed using EBS data. $fitted_{\hat{\varphi},t}$ and $v_{\hat{\varphi},t}$ are, respectively, the predicted value of $\hat{\varphi}_t$ and the residual from the regression estimated in Panel A; $\hat{\alpha}_t$ is the number of all arbitrage opportunities on day t divided by the total number of trades on this day; $\hat{\sigma}_t^{tox}$ is the average size of toxic arbitrage opportunities in day t (in basis points); vol_{it} is the realized volatility (in percentage) of 5-minutes returns for currency pair i on day t; $nrorders_{it}$ (in thousands) is the total number of orders (market, limit or cancelations) in currency pair i on day t; $trsize_{it}$ is the average daily trade size (in million) for currency pair i on day t; t-statistics in parenthesis are calculated using robust standard errors correcting for heteroscedasticity and serial correlation. The sample period is from January 2, 2003 to December 30, 2004.

Panel A: Forecasting $\hat{\varphi}$

$spread_{GU,t-1}$	-0.034 (-1.99)
$spread_{GE,t-1}$	0.060 (2.56)
$nrorders_{GU,t-1}$	0.002 (2.21)
$trsize_{GU,t-1}$	-0.185 (-1.98)
$trsize_{GE,t-1}$	0.109 (2.10)
$\hat{\varphi}_{t-1}$	0.158 (3.33)
$\hat{\varphi}_{t-2}$	0.104 (2.23)
$\hat{\varphi}_{t-3}$	0.115 (2.59)
$\hat{\varphi}_{t-7}$	0.099 (2.20)
$\hat{\varphi}_{t-12}$	0.106 (2.37)
$\hat{\varphi}_{t-13}$	0.219 (4.87)
$\hat{\varphi}_{t-15}$	0.116 (2.52)
$Adj.R^2$	44.39%
Obs.	463

Panel B: Illiquidity and Arbitrage Mix

	(1)	(2)	(3)	(4)	(5)	(6)
	spread	*espread*	*slope*	*spread*	*espread*	*slope*
$\hat{\varphi}$	0.619 (7.44)	0.505 (7.62)	0.396 (9.11)			
$v_{\hat{\varphi}}$				0.425 (4.49)	0.364 (4.21)	0.248 (4.87)
$fitted_{\hat{\varphi}}$				0.814 (5.57)	0.690 (5.43)	0.623 (8.34)
$\hat{\alpha}$	0.437 (1.18)	1.160 (2.59)	0.274 (1.43)	0.078 (0.20)	0.943 (2.17)	0.031 (0.16)
$\hat{\sigma}^{tox}$	0.156 (10.2)	0.180 (4.28)	0.072 (9.45)	0.157 (10.1)	0.183 (4.21)	0.072 (9.17)
vol	0.295 (7.86)	0.261 (4.81)	0.172 (9.17)	0.282 (7.57)	0.248 (4.54)	0.166 (9.10)
$trsize$	-0.344 (-3.37)	-0.354 (-2.43)	-0.400 (-7.46)	-0.352 (-3.34)	-0.371 (-2.38)	-0.401 (-7.22)
$nrorders$	-0.011 (-8.49)	-0.009 (-6.54)	-0.007 (-10.0)	-0.013 (-8.69)	-0.009 (-6.62)	-0.007 (-9.94)
$illiq^{EBS}$	0.025 (3.39)	0.014 (2.66)	0.012 (3.56)	0.021 (2.82)	0.011 (2.14)	0.009 (2.79)
$Adj.R^2$	86.54%	87.36%	82.99%	86.95%	87.56%	83.84%
Obs.		1,449			1,389	
Currency pair FE	YES	YES	YES	YES	YES	YES
Month dummies	YES	YES	YES	YES	YES	YES

Table V: Arbitrageurs' Relative Speed ($\hat{\pi}^{tox}$) and Market Illiquidity

This table reports estimates of the following IV regression:

$illiq_{it} = \omega_i + \xi_{t,m} + b_1\hat{\pi}_t^{tox} + b_2\hat{\varphi}_t + b_3\hat{\alpha}_t + b_4\hat{\sigma}_t^{tox} + b_5 vol_{it} + b_6 trsize_{it} + b_7 nrorders_{it} + b_8 illiq_{it}^{EBS} + \varepsilon_{it}$, for $i \in \{GU, EU, EG\}$ where indexes GU, EU, and EG refer to the GBP/USD, EUR/USD, and EUR/GBP currency pairs, respectively. We instrument $\hat{\pi}_t^{tox}$ with the introduction of AutoQuote on Reuters D-3000 (see the text). The first stage regression is: $\hat{\pi}_t^{tox} = \omega_i + \xi_{t,m} + a_1 AD_t + a_2\hat{\varphi}_t + a_3\hat{\alpha}_t + a_4\hat{\sigma}_t^{tox} + a_5 vol_{it} + a_6 trsize_{it} + a_7 nrorders_{it} + a_8 illiq_{it}^{EBS} + u_{it}$, where $illiq_{it}$ is one of our three proxies for illiquidity for currency i on day t: $spread_{it}$ is the average quoted bid-ask spread (in basis points) in currency pair i on day t; $espread_{it}$ (in basis points) is the average effective spreads in currency pair i on day t; $slope_{it}$ is the average slope of the limit order book in currency pair i on day t; superscript EBS is used for measures of these variables computed using EBS data. AD_t is a dummy variable equal to one after July 2003 and zero before; $\hat{\pi}_t^{tox}$ is the number of toxic arbitrage opportunities on day t that terminate with a trade divided by the total number of toxic arbitrage opportunities on day t; $\hat{\varphi}_t$ is the number of toxic arbitrage opportunities on day t divided by the number of arbitrage opportunities on this day; $\hat{\alpha}_t$ is the number of all arbitrage opportunities on day t divided by the total number of trades on this day; $\hat{\sigma}_t^{tox}$ is the average size of toxic arbitrage opportunities on day t (in basis points); vol_{it} is the realized volatility (in percentage) of 5-minutes returns for currency pair i on day t; $nrorders_{it}$ (in thousands) is the total number of orders (market, limit or cancelations) in currency pair i on day t; $trsize_{it}$ is the average daily trade size (in million) for currency pair i on day t.

	(1)	(2)	(3)	(4)	(5)	(6)
	spread		*espread*		*slope*	
	1^{st} stage	2^{nd} stage	1^{st} stage	2^{nd} stage	1^{st} stage	2^{nd} stage
AD	0.041 (4.23)		0.043 (4.22)		0.042 (4.25)	
$\hat{\pi}^{tox}$		7.746 (4.04)		3.334 (3.76)		4.421 (3.96)
$\hat{\varphi}$	-0.014 (-0.39)	0.702 (2.38)	-0.014 (-0.38)	0.553 (4.08)	-0.013 (-0.36)	0.452 (2.61)
$\hat{\alpha}$	0.362 (1.91)	-1.889 (-1.16)	0.357 (1.89)	0.263 (0.35)	0.364 (1.93)	-1.003 (2.61)
$\hat{\sigma}^{tox}$	-0.012 (-2.35)	0.242 (5.04)	-0.013 (-2.37)	0.219 (9.92)	-0.012 (-2.33)	0.122 (4.39)
vol	-0.012 (-1.04)	0.391 (3.89)	-0.013 (-1.05)	0.307 (6.65)	-0.013 (-1.06)	0.229 (3.87)
$trsize$	0.013 (0.25)	-0.121 (-0.29)	0.011 (0.21)	-0.228 (-1.17)	0.014 (0.28)	-0.260 (-1.09)
$nrorders$	-0.003 (-3.92)	-0.004 (-0.87)	-0.003 (-3.90)	-0.006 (-2.63)	-0.003 (-3.94)	-0.003 (-1.01)
$illiq^{EBS}$	0.002 (0.51)	0.024 (0.91)	0.001 (0.66)	-0.002 (-0.43)	0.001 (0.68)	0.001 (0.08)
$Adj.R^2$	2.52%	35.41%	2.53%	62.48%	2.54%	26.43%
$F-stats$	17.9		17.8		18.1	
Obs.	1,449					
Currency pair FE	YES		YES		YES	
Month dummies	YES		YES		YES	

Table VI: Toxic Arbitrage and Time-to-Efficiency

In this table, we present estimates of the following regression using OLS:
$\log(TTE_t) = c_i + \xi_{t,m} + a_1 AD_t + a_2 vol_{it} + a_3\hat{\varphi}_t + a_4\hat{\alpha}_t + a_5\hat{\sigma}_t + a_6 trsize_{it} + a_7 nrorders_{it} + u_{it}$, where TTE_t is the time-to-efficiency on day t of toxic arbitrage opportunities (*Toxic* column) or any (both toxic and non-toxic) arbitrage opportunity (*All* column), AD (AutoQuote Dummy) is a dummy variable equal to one after July, 2003 and 0 before; $\hat{\varphi}_t$ is the number of toxic arbitrage opportunities on day t divided by the number of arbitrage opportunities in this day; $\hat{\alpha}_t$ is the number of all arbitrage opportunities on day t divided by the total number of trades on this day; $\hat{\sigma}_t^{tox}$ is the average size of arbitrage opportunities in day t (in basis points); vol_{it} is the realized volatility (in percentage) of 5-minutes returns for currency pair i on day t; $nrorders_{it}$ (in thousands) is the total number of orders (market, limit or cancelations) in currency pair i on day t; $trsize_{it}$ is the average daily trade size (in million) for currency pair i on day t; t-statistics in parenthesis are calculated based on robust standard errors correcting for heteroscedasticity and serial correlation. The sample period is from January 2, 2003 to December 30, 2004.

Dep.Var: $log(TTE)$	Toxic	All
AD	-0.065 (-2.92)	-0.053 (-2.72)
$\hat{\varphi}$	-0.267 (-2.71)	0.022 (0.28)
$\hat{\alpha}$	0.886 (0.99)	1.274 (1.44)
$\hat{\sigma}^{tox}$	0.069 (3.31)	0.084 (4.47)
vol	-0.089 (-3.13)	-0.112 (-4.20)
$trsize$	0.019 (0.13)	0.011 (0.09)
$nrorders$	-0.012 (-6.16)	-0.010 (-6.44)
$Adj.R^2$	21.22%	33.36%
Obs	1,449	

Table VII: Currency-Level Tests

In this table, we estimate the following regression for each currency pair separately:
$illiq_{it} = \omega_i + \xi_{t,m} + b_1\hat{\pi}_t^{tox} + b_2\hat{\varphi}_t + b_3\hat{\alpha}_t + b_4\hat{\sigma}_t + b_5 vol_{it} + b_6 trsize_{it} + b_7 nrorders_{it} + b_8 illiq_{it}^{EBS} + \varepsilon_{it}, \quad for \quad i \in \{GU, EU, EG\}$, where indexes GU, EU, and EG refer to the GBP/USD, EUR/USD, and EUR/GBP currency pairs, respectively. For each currency pair, we estimate this regression using an IV approach in which we instrument $\hat{\pi}_t^{tox}$ with the introduction of AutoQuote on Reuters D-3000 (see the text). The first stage regression is:
$\hat{\pi}_t^{tox} = \omega_i + \xi_t + a_1 AD_t + a_2\hat{\varphi}_t + a_3\hat{\alpha}_t + a_4\hat{\sigma}_t + a_5 vol_{it} + a_6 trsize_{it} + a_7 nrorders_{it} + a_8 illiq_{it}^{EBS} + u_{it}$, where $illiq_{it}$ is one of our three proxies for illiquidity for currency i on day t: $spread_{it}$ is the average quoted bid-ask spread (in basis points) in currency pair i on day t; $espread_{it}$ (in basis points) is the average effective spreads in currency pair i on day t; $slope_{it}$ is the average slope of the limit order book in currency pair i on day t; superscript EBS is used when these variables are computed using EBS data. AD_t is a dummy variable equal to one after July 2003 and zero before. $\hat{\pi}_t^{tox}$ is the number of toxic arbitrage opportunities that terminate with a trade on day t divided by the total number of toxic arbitrage opportunities on day t; $\hat{\varphi}_t$ is the number of toxic arbitrage opportunities on day t divided by the number of arbitrage opportunities in this day; $\hat{\alpha}_t$ is the number of all arbitrage opportunities on day t divided by the total number of trades on this day; $\hat{\sigma}_t^{tox}$ is the average size of toxic arbitrage opportunities in day t (in basis points); vol_{it} is the realized volatility (in percentage) of 5-minutes returns for currency pair i on day t; $nrorders_{it}$ (in thousands) is the total number of orders (market, limit or cancelations) in currency pair i on day t; $trsize_{it}$ is the average daily trade size (in million) for currency pair i on day t; t-statistics in parenthesis are calculated using robust standard errors correcting for heteroscedasticity and serial correlation. The sample period is from January 2, 2003 to December 30, 2004.

Panel A: GBP/USD

	(1)	(2)	(3)	(4)	(5)	(6)
	spread		*espread*		*slope*	
	1^{st} stage	2^{nd} stage	1^{st} stage	2^{nd} stage	1^{st} stage	2^{nd} stage
AD	0.060 (7.33)		0.062 (6.56)		0.060 (7.55)	
$\hat{\pi}^{tox}$		5.136 (5.18)		2.991 (5.50)		1.985 (5.53)
$\hat{\varphi}$	-0.038 (-0.79)	0.957 (3.81)	-0.038 (-0.80)	0.596 (4.13)	-0.038 (-0.79)	0.420 (4.25)
$\hat{\alpha}$	0.414 (4.51)	-1.243 (-1.89)	0.405 (4.57)	0.613 (1.24)	0.416 (4.56)	-0.354 (-0.92)
$\hat{\sigma}^{tox}$	-0.024 (-2.55)	0.275 (5.11)	-0.025 (-2.61)	0.194 (7.49)	-0.024 (-2.46)	0.128 (7.30)
vol	0.005 (0.24)	0.248 (2.12)	0.004 (0.17)	0.243 (3.13)	0.004 (0.21)	0.096 (2.12)
$trsize$	0.145 (2.44)	-0.926 (-2.02)	0.146 (2.38)	-0.154 (-0.58)	0.146 (2.40)	-0.724 (-3.41)
$nrorders$	-0.004 (-6.40)	0.001 (0.36)	-0.004 (-5.92)	0.003 (0.86)	-0.004 (-6.72)	-0.004 (-3.05)
$illiq^{EBS}$	0.001 (0.26)	0.031 (1.40)	0.001 (0.64)	0.007 (1.22)	0.001 (0.50)	0.006 (3.04)
$Adj.R^2$	1.49%	8.97%	1.50%	17.41%	1.50%	15.63%
$F-stats$	53.7		43.0		57.0	
Obs.			483			
Month dummies	*YES*		*YES*		*YES*	

Table VII continued

Panel B: EUR/GBP

	(1)	(2)	(3)	(4)	(5)	(6)
	spread		*espread*		*slope*	
	1^{st} stage	2^{nd} stage	1^{st} stage	2^{nd} stage	1^{st} stage	2^{nd} stage
AD	0.064 (7.96)		0.069 (8.07)		0.065 (7.41)	
$\hat{\pi}^{tox}$		1.430 (4.59)		0.673 (3.55)		0.389 (1.85)
$\hat{\varphi}$	-0.058 (-1.38)	0.499 (10.3)	-0.055 (-1.27)	0.319 (13.4)	-0.057 (-1.35)	0.224 (13.8)
$\hat{\alpha}$	0.361 (4.11)	-0.265 (-0.93)	0.359 (3.99)	0.338 (1.17)	0.352 (3.67)	0.023 (0.14)
$\hat{\sigma}^{tox}$	-0.025 (-3.87)	0.097 (6.10)	-0.026 (-3.90)	0.066 (6.36)	-0.026 (-3.83)	0.036 (5.62)
vol	0.065 (2.50)	0.976 (13.2)	0.049 (1.77)	0.748 (11.6)	0.076 (2.57)	0.551 (12.2)
$trsize$	0.009 (0.20)	-0.322 (-5.67)	0.009 (0.22)	-0.096 (-2.49)	0.008 (0.18)	-0.321 (-21.0)
$nrorders$	-0.005 (-7.57)	-0.015 (-14.8)	-0.005 (-7.58)	-0.011 (-13.8)	-0.005 (-7.84)	-0.011 (-20.5)
$illiq^{EBS}$	0.004 (1.91)	0.026 (2.21)	0.004 (4.17)	0.006 (2.75)	0.002 (4.43)	0.001 (0.34)
$Adj.R^2$	1.27%	48.25%	1.44%	62.29%	1.37%	67.91%
$F-stats$	63.4		65.1		54.9	
Obs.	483					
Month dummies	*YES*		*YES*		*YES*	

Panel C: EUR/USD

	(1)	(2)	(3)	(4)	(5)	(6)
	spread		*espread*		*slope*	
	1^{st} stage	2^{nd} stage	1^{st} stage	2^{nd} stage	1^{st} stage	2^{nd} stage
AD	0.044 (3.62)		0.037 (6.42)		0.035 (3.18)	
$\hat{\pi}^{tox}$		0.344 (0.42)		2.844 (3.73)		3.833 (2.75)
$\hat{\varphi}$	0.012 (0.26)	0.200 (2.13)	0.014 (0.31)	0.387 (1.88)	0.022 (0.47)	0.043 (0.20)
$\hat{\alpha}$	0.522 (5.29)	-0.366 (-0.64)	0.560 (4.81)	-0.290 (-0.44)	0.539 (6.35)	-2.524 (-3.35)
$\hat{\sigma}^{tox}$	-0.013 (-1.30)	0.091 (5.21)	-0.009 (-1.08)	0.150 (5.52)	-0.016 (-1.64)	0.112 (2.25)
vol	-0.004 (-0.55)	0.018 (0.68)	0.013 (1.33)	-0.112 (-3.16)	-0.006 (-0.92)	0.074 (3.09)
$trsize$	0.108 (1.28)	-0.241 (-1.01)	0.127 (1.56)	-0.695 (-1.28)	0.101 (1.26)	-0.692 (-2.07)
$nrorders$	-0.002 (-2.69)	-0.009 (-3.92)	-0.001 (-1.98)	-0.018 (5.87)	-0.002 (-2.62)	0.001 (0.36)
$illiq^{EBS}$	-0.256 (-2.23)	6.680 (17.1)	-0.268 (-4.48)	3.843 (10.4)	-0.179 (-1.56)	4.563 (12.5)
$Adj.R^2$	0.71%	75.02%	0.99%	40.14%	0.40%	19.77%
$F-stats$	13.1		41.2		10.1	
Obs.	483					
Month dummies	*YES*		*YES*		*YES*	

Table VIII: Toxic Arbitrage or Other Forms of Adverse Selection?

This table reports IV estimates of the following equation for $i \in \{GU, EU, EG\}$:
$illiq_{it} = \omega_i + \xi_{t,m} + b_1\hat{\pi}^{tox}_t + b_2\hat{\varphi}_t + b_3\hat{\alpha}_t + b_4\hat{\sigma}^{tox}_t + b_5 vol_{it} + b_6 trsize_{it} + b_7 nrorders_{it} + b_8 illiq^{EBS}_{it} + b_9 VPIN_{it} + b_{10} macro^{US}_t + b_{11} macro^{UK}_t + b_{12} macro^{EMU}_t + \varepsilon_{it}, \quad for \quad i \in \{GU, EU, EG\}$. We instrument $\hat{\pi}^{tox}_t$ with the introduction of Reuters D-3000 AutoQuote. AD_t is a dummy variable equal to one after July 2003 and zero before. The first stage regression is:
$\hat{\pi}^{tox}_t = \omega_i + \xi_{t,m} + a_1 AD_t + a_2\hat{\varphi}_t + a_3\hat{\alpha}_t + a_4\hat{\sigma}^{tox}_t + a_5 vol_{it} + a_6 trsize_{it} + a_7 nrorders_{it} + a_8 illiq^{EBS}_{it} + a_9 VPIN_{it} + a_{10} macro^{US}_t + a_{11} macro^{UK}_t + a_{12} macro^{EMU}_t + u_{it}, \quad for \quad i \in \{GU, EU, EG\}$. Indexes GU, EU, and EG refer to the GBP/USD, EUR/USD, and EUR/GBP currency pairs, respectively. $illiq_{it}$ is one of our three proxies for illiquidity for currency i on day t: $spread_{it}$ is the average quoted bid-ask spread (in basis points) in currency pair i on day t; $espread_{it}$ (in basis points) is the average effective spreads in currency pair i on day t; $slope_{it}$ is the average slope of the limit order book in currency pair i on day t; superscript EBS is used when these variables are computed using EBS data. $\hat{\pi}^{tox}_t$ is the number of toxic arbitrage opportunities that terminate with a trade on day t divided by the total number of toxic arbitrages on this day; $\hat{\varphi}_t$ is the number of toxic arbitrage opportunities in day t divided by the number of arbitrage opportunities in that day; $\hat{\alpha}_t$ is the number of all arbitrage opportunities on day t divided by the total number of trades on this day; $\hat{\sigma}^{tox}_t$ is the average size of arbitrage opportunities in day t (in basis points); vol_{it} is the realized volatility (in percentage) of 5-minutes returns for currency pair i in day t; $trsize_{it}$ is the average daily trade size (in million) for currency pair i on day t; $nrorders_{it}$ (in thousands) is the total number of orders (market, limit or cancelations) in currency pair i on day t; $VPIN_{it}$ is a measure of adverse selection in currency pair i on day t (see Easley et al. (2012)); $macro^{US}$, $macro^{UK}$, and $macro^{EMU}$ are measures of surprises in macro economic announcements on day t in the U.S., the U.K., and the EMU, respectively. t-statistics in parenthesis are calculated based on robust standard errors correcting for heteroscedasticity and serial correlation. The sample period is from January 2, 2003 to December 30, 2004.

	(1)	(2)	(3)	(4)	(5)	(6)
	spread		*espread*		*slope*	
	1^{st} stage	2^{nd} stage	1^{st} stage	2^{nd} stage	1^{st} stage	2^{nd} stage
AD	0.039 (4.01)		0.040 (4.00)		0.039 (4.02)	
$\hat{\pi}^{tox}$		8.114 (3.84)		3.349 (3.56)		4.671 (3.89)
$\hat{\varphi}$	-0.016 (-0.42)	0.733 (2.38)	-0.016 (-0.45)	0.546 (4.01)	-0.016 (-0.43)	0.469 (2.68)
$\hat{\alpha}$	0.437 (-0.42)	-2.310 (-1.27)	0.431 (2.23)	0.303 (0.37)	0.439 (2.29)	-1.379 (-1.33)
$\hat{\sigma}^{tox}$	-0.022 (-3.53)	0.319 (4.73)	-0.023 (-3.81)	0.201 (6.73)	-0.022 (-3.77)	0.168 (4.37)
vol	-0.003 (-0.02)	0.333 (3.21)	-0.003 (-0.27)	0.328 (7.12)	-0.003 (-0.27)	0.193 (3.24)
$trsize$	0.038 (0.82)	-0.333 (-0.75)	0.036 (0.67)	-0.153 (-0.78)	0.038 (0.73)	-0.368 (-1.47)
$nrorders$	-0.003 (-3.71)	-0.002 (-0.34)	-0.003 (-3.38)	-0.007 (-2.74)	-0.003 (-3.39)	-0.003 (-0.83)
$illiq^{EBS}$	0.001 (0.55)	0.027 (1.00)	0.001 (0.59)	-0.001 (-0.27)	0.001 (0.52)	0.001 (0.26)
$VPIN$	0.057 (-0.45)	0.285 (0.32)	0.053 (0.52)	-0.250 (-0.64)	0.057 (0.56)	-0.159 (-0.32)
$macro^{US}$	-0.002 (-1.75)	0.016 (1.84)	-0.002 (-1.69)	0.008 (1.98)	-0.002 (-1.69)	0.009 (1.83)
$macro^{UK}$	-0.012 (-2.16)	0.091 (1.68)	-0.012 (-2.04)	0.038 (1.59)	-0.011 (-2.02)	0.050 (1.62)
$macro^{EMU}$	-0.011 (-1.80)	0.085 (1.60)	-0.011 (-1.89)	0.035 (1.47)	-0.011 (-1.91)	0.046 (1.50)
$Adj.R^2$	3.41%	33.27%	3.42%	61.54%	3.42%	24.12%
$F-stats$	16.1		16.0		16.2	
Obs.	1,449					
Currency pair FE	*YES*		*YES*		*YES*	
Month dummies	*YES*		*YES*		*YES*	

References

Akram, F., Rime, D., Sarno, L., 2008, "Arbitrage in the foreign exchange market: Turning on the Microscope," *Journal of International Economics*, 76, 237–253.

Andersen, T. G., Bollerslev, T., Diebold, F. X., Vega, C., 2003, "Micro effects of macro announcements: Real-time price discovery in foreign exchange," *American Economic Review*, 93, 38–62.

Ben-David, I., Franzoni, F., Moussawi, R., 2012, "ETFs, arbitrage, and shock propagation," *Working paper, Fisher College of Business.*

Biais, B., Foucault, T., Moinas, S., 2014, "Equilibrium fast trading," *Journal of Financial economics, forthcoming.*

Bjønnes, G., Rime, D., 2005, "Dealer behavior and trading systems in foreign exchange markets," *Journal of Financial Economics*, 75, 571–605.

Brogaard, J., Hendershott, T., Riordan, R., 2014, "High frequency trading and price discovery," *Review of Financial Studies.*

Budish, E., Cramton, E., Shim, J., 2014, "The High frequency trading arms race: Frequent batch auctions as a market design response," *Working Paper, University of Chicago.*

Chaboud, A., Chiquoine, B., Hjalmarsson, E., Vega, C., 2014, "Rise of the machine: Algorithmic trading in the foreign exchange market," *Journal of Finance*, 65, 2045–2084.

Chordia, T., Roll, R., Subrahmanyam, A., 2008, "Liquidity and market efficiency," *Journal of Financial Economics*, 87, 249–268.

Copeland, D., Galai, T., 1983, "Information effects on the bid-ask spread," *Journal of Finance*, 38, 1457–1469.

Deville, L., Riva, F., 2007, "Liquidity and arbitrage in options markets: A survival analysis approach," *Review of Finance*, 11, 497–525.

Duffie, D., 2010, "Asset price dynamics with slow moving capital," *Journal of Finance*, 65, 1237–1267.

Duffie, D., Garleanu, N., Pedersen, L., 2005, "Over-the-counter markets," *Econometrica*, 73, 1815–1857.

Easley, D., Lopez de Prado, M. M., O'Hara, M., 2011, "The microstructure of the flash crash: Flow toxicity, liquidity crashes and the probability of informed trading," *Journal of Portfolio Management*, 37, 118–128.

Easley, D., Lopez de Prado, M. M., O'Hara, M., 2012, "Flow toxicity and liquidity in a high frequency world," *Review of Financial Studies*, 25, 1457–1493.

Fenn, D., Howison, S., Mcdonald, M., Williams, S., Johnson, N., 2009, "The mirage of triangular arbitrage in the spot foreign exchange market," *International Journal of Theoretical and Applied Finance*, 12, 1105–1123.

Fong, W., Valente, G., Fung, J., 2008, "Market liquidity and FX arbitrage in emerging markets: Evidence from Hong Kong tick Data," *Working Papers 082008, Hong Kong Institute for Monetary Research.*

Foucault, T., Hombert, J., Rosu, J., 2014, "News trading and speed," *Forthcoming, Journal of Finance.*

Foucault, T., Röell, A., Sandås, P., 2003, "Market making with costly monitoring: An analysis of the SOES controversy," *Review of Financial Studies*, 16, 345–384.

Gagnon, L., Karolyi, G., 2010, "Multi-market trading and arbitrage," *Journal of Financial Economics*, 53, 53–80.

Garvey, R., Wu, F., 2010, "Speed, distance, and electronic trading: New evidence on why location matters," *Journal of Financial Markets*, 13, 367–396.

Green, T. C., 2004, "Economic news and the impact of trading on bond prices," *Journal of Finance*, 59, 1201–1234.

Gromb, D., Vayanos, D., 2002, "Equilibrium and welfare in markets with financially constrained arbitrageurs," *Journal of Financial Economics*, 66, 361–407.

Gromb, D., Vayanos, D., 2010, "Limits to arbitrage: The state of the theory," *Annual Review of Financial Economics*, 98, 251–275.

Grossman, S., Miller, M., 1988, "Liquidity and market structure," *Journal of Finance*, 43, 617–637.

Hendershott, T., Jones, C. M., Menkveld, A. J., 2011, "Does algorithmic trading improve liquidity?," *Journal of Finance*, 66, 1–34.

Hendershott, T., Menkveld, A., 2014, "Price pressures," *Journal of Financial Economics*, 114, 405–423.

Hendershott, T., Moulton, P., 2011, "Automation, speed, and stock market quality: The NYSE's hybrid," *Journal of Financial Markets*, 14, 568–604.

Ho, T., Stoll, H. R., 1981, "Optimal dealer pricing under transactions and return uncertainty," *Journal of Financial Economics*, 9, 47–73.

Hoffmann, P., 2014, "A dynamic limit order market with fast and slow traders," *Journal of Financial Economics*, 113, 159–169.

Holden, C., 1995, "Index arbitrage as cross sectional market making," *Journal of Futures Markets*, 15, 423–455.

Holden, C., Jacobsen, S., Subrahmanyam, A., 2014, "The empirical analysis of liquidity," *Foundations and Treans in Finance.*

Ito, T., Yamada, K., Takayasu, M., Takayasu, H., 2013, "Free lunch! Arbitrage opportunities in the foreign exchange markets," *Working Paper, NBER.*

Jovanovic, B., Menkveld, A., 2012, "Middlemen in limit order markets," *Working Paper, University of Amsterdam.*

King, M. R., Osler, C., Rime, D., 2012, *Foreign exchange market structure, players and evolution* . Handbook of Exchange Rates, Wiley.

King, M. R., Rime, D., 2010, "The $4 trillion question: What explains FX growth since the 2007 survey?," *BIS Quarterly Review*, pp. December, 27–40.

Kozhan, R., Tham, W. W., 2012, "Execution risk in high-frequency arbitrage," *Management Science*, 58, 2131–2149.

Kumar, P., Seppi, D., 1994, "Information and index arbitrage," *Journal of Business*, 67, 481–509.

Mancini-Griffoli, T., Ranaldo, A., 2011, "Limits to arbitrage during the crisis: funding liquidity constraints and covered interest parity," *Working paper, Swiss National Bank.*

Marshall, B., Treepongkaruna, S., Young, M., 2008, "Exploitable Arbitrage Opportunities Exist in the foreign exchange Market," In: *American Finance Association Annual Meeting, New Orleans.*

Menkveld, A., Zoican, M., 2014, "Need for speed. Exchange latency and liquidity," *Working Paper, University of Amsterdam.*

Osler, C., 2008, *Foreign exchange microstructure: A survey* . Springer Encyclopedia of Complexity and System Science, Springer.

Pasquariello, P., 2014, "Financial market dislocations," *Review of Financial Studies*, 27, 1868–1914.

Pierron, A., 2007, "Electronic Platforms in foreign exchange trading," Discussion paper, Celent, Paris.

Roll, R., Schwartz, E., Subrahmanyam, A., 2007, "Liquidity and the Law of One Price: The Case of future-cash basis," *Journal of Finance*, 62, 2201–2234.

Rosch, D., 2014, "The impact of arbitrage on market liquidity," *Working paper, Erasmus University.*

Shive, S., Schultz, P., 2010, "Mispricing of dual class shares: Profit opportunities, arbitrage, and trading," *Journal of Financial Economics*, 98, 524–549.

Sussman, A., Tabb, L., Iati, R., 2009, "U.S equity high frequency trading: Strategies, sizing and market structure," *Report tabb Group.*

U.S. Securities and Exchange Commission, 2010, "U.S. Concept Release on Equity Market Structure," *17 CFR Part 242, Release No. 34-61358*, pp. File No. S7–02–10.