Toxic Arbitrage[*]

Thierry Foucault[†]   Roman Kozhan[‡]   Wing Wah Tham[§]

November, 2014

(First draft: September 2012)

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

[†]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 cheaper 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, such as that for catastrophe risk insurance."*

Well-known frictions (e.g., short-selling costs, funding constraints etc.) explain why some arbitrage opportunities persist (see Gromb and Vayanos (2010)). However, these frictions are less likely to play out for arbitrage opportunities lasting *fractions of a second.* For such very short lived opportunities, attention costs and technological constraints on traders' speed are the main impediments to a seamlessly Law of One Price. These barriers are falling as some arbitrageurs invest massively in fast trading technologies to react ever faster to arbitrage opportunities, making them even more short lived. Their profits per opportunity are small but short lived arbitrage opportunities are very frequent because of the proliferation of derivatives securities (e.g., ETFs, options, futures) and market fragmentation. Accordingly, returns on high speed arbitrage are high: a report from the Tabb group estimates profits of high frequency arbitrageurs at $21 billion for year 2009 alone (see Sussman et al. (2009)).

This evolution raises questions about the social value of arbitrage, in particular high speed arbitrage. For instance, in its report on recent changes in market structures, the SEC asks whether arbitrage strategies "*benefit or harm the interests of long-term investors and market quality in general.[...]*" (see U.S. Securities and Exchange Commission (2010), Section B, p.51). At first glance, questioning the social value of arbitrage is surprising: economists tend to view arbitrage as beneficial because it enhances price efficiency. Moreover, arbitrageurs implicitly act as liquidity providers when they correct mispricings due to transient demand or supply shocks ("price pressures" effects). In this view, not only arbitrageurs make prices more efficient but they also reduce illiquidity costs for other investors (see, for instance, Holden (1995), Gromb and Vayanos (2002), and Gromb and Vayanos (2010)).[1]

---

[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*

For high speed arbitrage opportunities, this view is incomplete, however. Indeed, these opportunities do not arise only because of transient supply and demand shocks. Asynchronous adjustments in asset prices following new information can also cause very short lived breakdowns of the LOP. In these cases, in enforcing the LOP, arbitrageurs expose their counterparties (e.g., market makers) to the risk of trading at stale quotes ("being picked off").[2]

Consider, for instance, two "market makers" (or limit order books) A and B trading the same asset in two different trading platforms. Suppose that good news about this asset arrives. Market maker A instantaneously marks up his quotes while B is slower. If, as a result, A's bid price exceeds B's ask price momentarily, there is an arbitrage opportunity.[3] If arbitrageurs are fast enough, they buy the asset from B, before the latter updates his quotes, and they resell it to A, at a profit. As B sells the asset at a price lower than its fair value, he incurs a loss, as if he had been trading with better informed investors.

Thus, asynchronous price adjustments to information in asset pairs generate "toxic" arbitrage opportunities in the sense that these opportunities raise the risk of trading at stale quotes for market makers.[4] Dealers require a compensation for this form of adverse selection (Copeland and Galai (1983)). Thus, illiquidity should be higher when the fraction of arbitrage opportunities that are toxic is higher. Moreover, illiquidity should be higher when the likelihood that a toxic arbitrage opportunity terminates with an arbitrageur's trade is higher. Our main goal in this paper is to test these two predictions.

We check that these predictions are well-grounded using a simple model of cross-market arbitrage in pairs of similar assets with specialized market makers. 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 who observes the arbitrage 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

---

liquidity to the market."

[2]This problem is not new. In the 90s, professional day traders (so-called SOES bandits) were picking off 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]This situation is not uncommon. Shkilko et al. (2008) show that ask and bid prices for stocks listed on Nasdaq and the NYSE (and traded on multiple markets) are crossed 3.5% of the time during a day.

[4]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."*

generates our two main predictions and has two additional implications 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 in equilibrium, 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 effect also induces arbitrageurs to be faster so that overall arbitrage opportunities are shorter. A similar logic implies that 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.

Our tests use data on triangular arbitrage opportunities for three currency pairs (dollar-euro, dollar-pound, and pound-euro).[5] 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 arbitrage opportunities that are toxic; 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.

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)). 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 than by humans, and (iv) they deliver razor blade profits per opportunity (1 to 2 basis points

---

[5]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.

in our sample).[6]

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, whereas asynchronous price adjustments are eventually followed by permanent shifts in exchange rates.[7] 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 an increase of about 3% in average effective spreads for the currencies in our sample. Thus, the arbitrage mix matters: pairs of related assets are less liquid when arbitrage opportunities in these pairs 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 model shows that this likelihood is jointly determined with illiquidity in equilibrium. 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 their algorithms to Reuters D-3000 data and let these algorithms submit orders automatically. AutoQuote API marked therefore the onset of algorithmic trading on

---

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

[7]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)).

Reuters D-3000, allowing traders to react faster to triangular arbitrage opportunities.[8] 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 anedoctal evidence that algorithmic trading in foreign exchange markets was initially used for exploiting triangular arbitrage opportunities (see Chaboud et al. (2014)).

The second stage of our IV regression shows that an increase in the likelihood that a toxic arbitrage opportunity terminates with a trade has a positive effect on illiquidity, as conjectured. Specifically, a 1% increase in this likelihood is associated with a 0.08 basis points increase in quoted bid-ask spreads in our sample (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 $161,000 per day). We find similar effects when we use effective spreads and the slope of limit order books (a measure of market depth) as measures of market illiquidity.

Thus, consistent with our predictions, (a) the fraction of arbitrage opportunities that are toxic and (b) the frequency with which these opportunities are effectively exploited (a proxy for arbitrageurs' relative speed of reaction) are associated with a less liquid market. Yet, faster arbitrageurs also increase pricing efficiency. For instance, the introduction of AutoQuote API coincides with a 6.8% (about 60 milliseconds) decrease in the average duration of arbitrage opportunities in our sample. This finding is consistent with Chaboud et al. (2014), who find that algorithmic trading reduces the likelihood of observing a triangular arbitrage opportunity when one samples market data every second. Moreover, this duration of is negatively related to the daily fraction of arbitrage opportunities that are toxic, again as predicted by the model.

Several papers suggest that liquidity facilitates arbitrage and thereby enhances price efficiency (see Holden et al. (2014)). In this paper, however, we consider the reverse relation: the effect of arbitrageurs on liquidity. Kumar and Seppi (1994) emphasize the connection be-

---

[8]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.

tween cross-asset arbitrageurs and informed trading as we do in this paper. However, to our knowledge, our paper is first to test whether high speed arbitrage is a source of illiquidity.[9]

Our paper is also related to papers on the effect of speed in securities markets.[10] Several of these papers argue that fast informed traders raise adverse selection costs for slower traders. 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. However, our contribution is **not** to test whether adverse selection 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 and contribute to the question of whether high speed arbitrageurs are beneficial.

Arbitrage opportunities in the foreign exchange market (either violations of covered interest parity or triangular arbitrage) are well documented.[11] However, existing papers on these opportunities do not study the effect of arbitrageurs on liquidity. A few papers have also analyzed the extent to which market makers in FX markets are exposed to adverse selection (e.g., Lyons (1995), Bjønnes and Rime (2005)) and the source of informational asymmetries in FX markets (e.g., Bjønnes et al. (2011)). We complement them by showing that arbitrageurs' orders can be, in some circumstances, a source of adverse selection.

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.

---

[9]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.

[10]See, for instance, Biais et al. (2014), Budish et al. (2014), Foucault et al. (2003), Foucault et al. (2012), Garvey and Wu (2010), Hendershott and Moulton (2011), Hoffman (2012), Jovanovic and Menkveld (2012), Menkveld and Zoican (2014), or Pagnotta and Phillipon (2013).

[11]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).

2 Hypotheses Development

In this section, we present the model that guides our empirical analysis. The model is similar to Foucault et al. (2003) but allow for cross market arbitrage (Foucault et al. (2003) consider reactions to public information arrival in a single asset). We deliberately keep the model very simple in order to better highlight the economic forces that we seek to identify in the data. Extensions are discussed in Section 2.3.

2.1 A model of cross-market arbitrage

Consider two risky assets X and Y. The model has three dates, $t \in \{0, 1, 2\}$. At $t = 2$, each asset pays a single cash-flow, $\tilde{\theta}_X$ and $\tilde{\theta}_Y$ with $\tilde{\theta}_X = \sigma \tilde{\theta}_Y$. Thus, a portfolio with a long position of σ shares of asset Y and a short position of one share of asset X is riskless. Let v_X and v_Y be the expected payoffs of assets X and Y at $t = 0$. We have:

$$v_X = \sigma \times v_Y. \tag{1}$$

There are two market makers (called X and Y) and one arbitrageur. Each market maker is specialized in one asset (trades only this asset). At $t = 0$, market maker $j \in \{X, Y\}$ learns his valuation m_j for asset j (see below). Then, at $t = 1$, the market makers simultaneously post an ask price, a_j and a bid price b_j for $j \in \{X, Y\}$ with:

$$a_j = m_j + \frac{S_j}{2}, \tag{2}$$

and

$$b_j = m_j - \frac{S_j}{2}. \tag{3}$$

Thus, S_j is the bid-ask spread for asset j. In asset X, quotes are firm for $Q_X = 1$ share and in asset Y for $Q_Y = \sigma$ shares.

Market maker X's valuation for asset X is equal to its expected payoff, i.e., $m_X = v_X$. With probability $(1 - \alpha)$, this is also the case for market maker Y: $m_Y = v_Y$. Alternatively, with probability α, market maker Y experiences a shock to his valuation for one of two possible reasons. First, with probability φ, he receives news about the payoff of asset Y and revises his expectation about the payoff of this asset by ε, where $\varepsilon = 1/2$ (good news) or $-1/2$ (bad news) with equal probabilities. His valuation is then $m_Y = v_Y + \varepsilon$. Alternatively, with probability

$(1 - \varphi)$, market maker Y is hit by a liquidity shock δ. In this case, $m_Y = v_Y + \delta$ where $\delta = 1/2$ or $-1/2$ with equal probabilities.

This liquidity shock captures changes in liquidity providers' valuations due to risk management concerns (so called "price pressures;" see Hendershott and Menkveld (2014) for instance). For instance, a market maker with a large long position in one asset will be less willing to buy other assets with positively correlated payoffs and more eager to sell these for hedging purposes. This corresponds to a negative δ in our model.

If there is no shock to market maker Y's valuation (probability $(1 - \alpha)$) then there is no arbitrage opportunity. For instance, buying σ shares of asset Y at a_Y and selling one share of asset X at b_X yields a zero payoff portfolio but costs $(S_X + \sigma S_Y)/2$. In this case, 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.

If instead there is a shock to market maker Y's valuation then there is an arbitrage opportunity if $S_X + \sigma S_Y < \sigma$. For instance, suppose that this shock is positive. The arbitrageur can then buy one share of asset X at price $a_X = v_X + \frac{S_X}{2}$ and sell σ shares of asset Y at price $b_Y = v_Y + 1/2 - S_Y/2$. With this trade, she locks in a sure gain of:

$$ArbProfit = \sigma \times b_Y - a_X = (\sigma - S_X - \sigma S_Y)/2. \tag{4}$$

By symmetry, the arbitrageur's profit is identical if market maker Y experiences a negative valuation shock.

The expected profit of market maker X on a trade with the arbitrageur depends on the nature of market maker Y's valuation shock. Consider again the previous example. If the valuation shock is due to news then market maker X will lose money on his trade with the arbitrageur. Indeed, in this case, asset X is worth $v_{tX} + \sigma/2$ and the market maker sells it at $a_X = v_X + \frac{S_X}{2}$. Thus, the market maker's expected profit is $(S_X - \sigma)/2$, which must be negative if the arbitrageur chooses 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, arbitrage opportunities due to news arrival about asset Y expose market maker X to a form of adverse selection (the risk of trading a stale quotes) while arbitrage opportunities due to liquidity shocks in asset Y generate profits for market maker X. For this reason, we

refer to the former type of arbitrage opportunity as *toxic* and to the latter type as *non-toxic*. Hence, parameter φ characterizes the composition of arbitrage opportunities (the "arbitrage mix"): the higher is φ, the higher is the likelihood that a given arbitrage opportunity is toxic.

To complete the model, we need to specify when and how an arbitrage opportunity terminates. We assume that it takes a time D^a and D^m for the arbitrageur and market maker X, respectively, to observe the arbitrage opportunity. 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 are the traders. We therefore refer to λ and γ as traders' speed, respectively.

If the arbitrageur detects first the opportunity ($D^a < D^m$), she exploits it. If instead market maker X is first to observe the opportunity, he cancels his quote if the arbitrage opportunity is toxic (to avoid trading at a loss with the arbitrageur) and does nothing if the arbitrage opportunity is non toxic (to earn a profit by trading with the arbitrageur).[12] 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 (in reality, the market maker will then resubmit new quotes around his updated valuation; this can easily be incorporated in the model by extending it to multiple trading rounds).

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

$$\pi \equiv \ \text{Pr}\left(D^a < D^m\right) = \frac{\gamma}{\lambda + \gamma}, \tag{5}$$

because D^m and D^a are exponentially distributed. Thus, π increases when the arbitrageur becomes *relatively* faster, i.e., when $\frac{\gamma}{\lambda}$ increases. Thus, π can be seen as a measure of the arbitrageur's relative speed.

In practice, traders' response time to market events (e.g., an arbitrage opportunity) is called *latency* (see Moallemi and Saglam (2013) or Pagnotta and Phillipon (2013)). Latencies are stochastic because traders cannot control the time required by platforms to process their orders and this time will be affected by a myriad of factors (e.g., trading activity). However, with

---

[12]For simplicity, we assume that market maker X learns the nature of the shock in asset Y upon observing the arbitrage opportunity. For instance, he can check whether the opportunity coincides with news to infer the nature of market maker Y's valuation shock.

investments in technologies (e.g., in hardware and software, datafeed, dedicated communication lines, co-location etc.) and attention (e.g., computer capacity), traders can reduce the time it takes for them to communicate with a trading platform and receive messages from the platform. These investments are costly, however.[13] Hence, we assume that if the market maker operates at speed λ then he bears a cost $\Psi^m(\lambda) = \frac{c^m \lambda_t}{2}$. Similarly, if the arbitrageur operates at speed γ then she bears a cost $\Psi^a(\gamma) = \frac{c^a \gamma_t}{2}$. The cost of speed for market maker Y is irrelevant: this market maker has no reason to be fast since arbitrage opportunities always originates in asset Y.

Allowing for the possibility that $c^a \neq c^m$ will be useful to analyze the effects of differential shocks to traders' speed. 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 marginal 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 new information (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 $\frac{c^a}{c^m}$.

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

$$\Pi^X(S_X; \lambda, \gamma) = -\frac{\varphi \alpha \pi_t}{2}(\sigma - S_X) + \left(\frac{1 - \alpha(2\varphi - 1)}{4} \right) S_X - \frac{c^m \lambda}{2}, \qquad (6)$$

$$\Pi^Y(S_Y) = (2\alpha(1 - (1 - \pi)\varphi) + 1 - \alpha)Q_Y \frac{S_Y}{4}, \qquad (7)$$

$$\Pi^a(S_X, S_Y; \lambda, \gamma) = \alpha \varphi \pi \frac{(\sigma - S_X - \sigma S_Y)}{2} + \alpha(1 - \varphi)\frac{(\sigma - S_X - \sigma S_Y)}{2} - \frac{c^a \gamma}{2}. \qquad (8)$$

The first term in (6) is market maker X's expected losses when he trades with the arbitrageur

---

[13]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.

[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. See the next section.

at stale quotes 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 with probability $(\alpha(1 - (1 - \pi)\varphi) + (1 - \alpha)/2)$ and earns the half bid-ask spread for asset Y in this case. He does not invest in speed since he is never exposed to the risk of trading at stale quotes. The first two terms in the arbitrageur's expected profit are the arbitrageur's expected gain on toxic and non toxic arbitrage, respectively. The last term is her cost of speed.

2.2 Testable predictions

We focus on equilibria in which market makers set competitive bid-ask spreads. Thus, $S_Y = 0$. In asset X, the competitive spread depends on traders' speeds (since they affect π and the market maker's cost). We require that the speeds chosen by the arbitrageur and the market maker maximize their expected profit for each possible level of the bid-ask spread in asset X. That is, for a given bid-ask spread S_X and for $S_Y = 0$, equilibrium speeds, γ^* and λ^* are such that (i) λ^* maximizes (6) when $\gamma = \gamma^*$ and (ii) γ^* maximizes (8) when $\lambda = \lambda^*$.

Thus, when $\lambda^* > 0$ or $\gamma^* > 0$, traders' speeds solve the following system of first order conditions $\frac{\partial \Pi^a}{\partial \gamma} = 0$ and $\frac{\partial \Pi^m}{\partial \lambda} = 0$.[16] The unique equilibrium is such that:

$$\lambda^*(S_X; c^a, c^m) = \frac{\varphi\alpha(\sigma - S_X)r}{c^m(1+r)^2}, \tag{9}$$

$$\gamma^*(S_X; c^a, c^m) = \frac{\varphi\alpha(\sigma - S_X)r}{c^a(1+r)^2}, \tag{10}$$

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

Not surprisingly, the traders' optimal speed decreases in their cost of speed (e.g., γ^* decreases in c^a). As expected, they also decrease in the 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. Equations (9) and (10) imply that:

$$\frac{\gamma^*(S_X; c^a, c^m)}{\lambda^*(S_X; c^a, c^m)} = r = c^m/c^a. \tag{11}$$

---

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

Thus, traders' relative speeds, and therefore π, are only determined by their marginal costs in equilibrium. Equations (5) and (11) yield:

$$\pi^*(r) = \frac{\gamma^*}{\lambda^* + \gamma^*} = \frac{r}{1+r}. \tag{12}$$

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

Substituting λ^* in the expected profit of market maker X (eq. (6)), we solve for the zero profit bid-ask spread in asset X. We obtain:

$$S_X^* = \frac{2\varphi\alpha\pi^*(r)(2 - \pi^*(r))\sigma}{2\varphi\alpha\pi^*(r)(2 - \pi^*(r)) + (1 - \alpha(2\varphi - 1))}. \tag{13}$$

The equilibrium competitive spread increases with σ, the size of arbitrage opportunities and is always less than σ, as conjectured previously. Market maker's X's exposure to the risk of being picked off by arbitrageurs also depends on φ and $\pi^*(r)$. This risk increases when arbitrage opportunities are more likely to be toxic (φ is higher) or, holding this likelihood constant, when arbitrageurs are faster ($\pi^*(r)$ is higher). Our two main testable implications follow.

Implication 1a. *Consider a pair of asset 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 1b. *Consider a pair of asset X and Y linked by a no-arbitrage relation. An increase in the likelihood (π^*) that a toxic arbitrage opportunity terminates by an arbitrageur's trade causes an increase in the bid-ask spread of asset X.*

The economics for these two predictions is straightforward: when arbitrageurs exploit arbitrage opportunities due to asynchronous price reactions to news, they adversely select market makers who adjust their quotes slowly. Yet, to our knowledge they are novel and have not been tested so far. The likelihood, π^*, that a toxic arbitrage opportunity terminates with an arbitrageur's trade (rather than a quote update) is endogenous and potentially determined by the bid-ask spread (since traders' speeds are). In testing Implication 1b, one must therefore account for the endogeneity of π^*. The model suggests to use shocks to the arbitrageur's relative cost of speed (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 r) triggers an increase in π^* and, through

this channel only, an increase in the bid-ask spread. This yields the following implication.

Implication 2. *Consider a pair of asset X and Y linked by a no-arbitrage relation. A reduction in arbitrageurs' cost of speed (c^a) relative to market makers' cost of speed (c^m) 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 bid-ask spreads.*

As explained previously, changes in market structures that affect traders' speeds are equivalent to variations in $r = c^m/c^a$. Insofar as these changes only affect liquidity through their effect on traders' speeds and therefore π^*, Implication 2 shows that they provide good instruments to measure the effect of π^* on the bid-ask spread, i.e., to test Implication 1b. 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:

$$E(D) = \varphi E\left(\mathrm{Min}\{D^a, D^m\}\right) + (1 - \varphi)E(D^a) = \frac{(1 + r) - \varphi}{\gamma^*(S_X^*)(1 + r)}. \tag{14}$$

As speed is costly, equilibrium speeds are never infinite. Thus, in equilibrium, arbitrage opportunities do not immediately vanish, so that $E(D) > 0$. The duration of arbitrage opportunities however goes to zero as 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.

An exogenous increase in the bid-ask spread of asset X translates into arbitrage opportunities that last longer because they induce the arbitrageur to invest less in speed (γ^* decreases in S_X). 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. However, bid-ask spreads and durations are determined simultaneously in equilibrium. For this reason, there exist cases in which a change in the environment can simultaneously increase bid-ask spreads and yet make the duration of arbitrage opportunities shorter because arbitrageurs have more incentive to react fast, despite the larger bid-ask spread.

Consider for instance a technological change that reduces c^a and c^m but for which the effect on c^a is bigger. Hence r 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 (9) and (10)). However, the bid-ask spread becomes larger in equilibrium because r is larger (Implication 2). This indirect effect reduces the arbitrageur's incentive to be fast and therefore the market maker's incentive to be fast as well. Nevertheless, the direct effect always dominates the indirect effect, which yields our next testable implication.

Implication 3: *Consider a decrease in c^a and c^m that eventually triggers an increase in r (i.e., the decrease in c^a is larger). This decrease 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 the action of algorithmic arbitrageurs hitting quotes of slower, human, traders. These findings are consistent with the logic of the model. Indeed, suppose that algorithmic trading reduces the cost of being fast for arbitrageurs. Then, r and therefore π^* increases. As a result, one should observe that 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 tests 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, even though the true likelihood of these opportunities (α in the model) is unaffected by algo trading.

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 (9) and (10)). 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 when $\varphi \leq 1/2$ or $\alpha \leq (4\varphi - 1)^{-1}$. (see the appendix for a formal proof).*

In the model, parameter α controls the fraction of arbitrage trades relative to the number

of trades occurring for other reasons (liquidity motives in the model). In our data, the number of arbitrage opportunities relative to the total number of trades is small. Thus, α is likely to be small empirically and well below $1/3$ which is a sufficient condition for $\alpha < (4\varphi - 1)^{-1}$ when $\varphi > 1/2$. Hence, empirically, 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, which is close to our estimate of the size of triangular arbitrage opportunities in our data (see below). 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 $r = 2.8$ and $\pi = 74\%$, which is the average value of π in the data (see Table II below). Moreover, we estimate the average value of φ to be 41% (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). In reality, market makers certainly bear other costs than just adverse selection costs due to arbitrageurs' trades in toxic opportunities (those captured by our model). In fact these costs should only be a fraction of total adverse selection costs for market makers and cannot therefore be a too large fraction of the bid-ask spread.[17]

2.3 Extensions

We briefly 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

---

[17]Bjonnes and Rime (2005) find that adverse selection costs account for about 70% to 80% of foreign exchange dealers's bid-ask spreads in their sample.

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

higher. Moreover, if $r \leq \frac{M}{(M-1)}$ 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 she can update her quote before being hit by an arbitrageur.[19] As a result the marginal benefit of speed is lower for the market maker. If the number of arbitrageurs and r are large enough 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 ($M \rightarrow \infty$). 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.[20] Market-makers' expected profit is also zero since the bid-ask spread is competitive. Thus, all firms are indifferent between both roles and one can be randomly selected to be a market maker.

Shocks in assets X and Y. Suppose that at $t = 0$, the shock to market makers' valuation happens in asset Y or X with probabilities $\alpha\beta$ and $\alpha(1 - \beta)$, respectively. In the baseline model, $\beta = 1$. When $0 < \beta < 1$, market makers in each asset are exposed to toxic arbitrage trades and therefore each market features a bid-ask spread. If the arbitrageur takes a long-short position in each asset, as in the baseline model, her expected profit depends on both S_X and S_Y (see (4)). Accordingly her optimal speed now depends on both spreads and for this reason the bid-ask spreads in each market become interdependent. Solving for the equilibrium in closed form becomes significantly more complex. The equilibrium can be solved numerically, however. Numerical simulations show that our predictions still hold in this more general case.

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

---

[19]This likelihood is $\frac{\lambda}{\lambda + M\gamma}$ when each arbitrageur chooses a speed of γ. Thus, a marginal increase in λ increases the market maker's chance of being first by $\frac{M\gamma}{(\lambda + M\gamma)^2}$. This decreases with M.

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

D-3000 is one of the two electronic trading platforms for interdealer spot trading in the 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 currency pairs in our sample account for 60 percent of all foreign exchange (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 at least two attractive features for our tests. First, it is very rich. For each order submitted to Reuters D-3000, the dataset reports the currency pair in which the order is submitted, the order type (limit or market), the time at which the order is entered, the size of the order, and the price attached to the order for a limit order. 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

---

[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.

with a trade, that is, π^*, in the model. This is important for testing Implication 1b.

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, Electronic Broking Service (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.[24] 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.[25] 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. This classification is an important step for our tests.

---

[24]Representatives from EBS told us that they do not provide data at a more granular level.

[25]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.

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 bid and ask 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 $A^{k/j}$ 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

$$
\begin{aligned}
\hat{A}_t^{i/j} &< B_t^{i/j} &\text{or,} && (15) \\
\hat{B}_t^{i/j} &> A_t^{i/j}. &&& (16)
\end{aligned}
$$

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 (15) and (16) cannot both be true at the same time; see Kozhan and Tham (2012).

Reuters D-3000 charge brokerage and membership fees. These fees may vary across traders and some are fixed (e.g., the subscription fee to the Reuters D-3000 platform). 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:

$$
\begin{aligned}
\frac{B_t^{i/j} - \hat{A}_t^{i/j}}{\hat{A}^{i/j}} - z &> 0, && (17) \\
\frac{\hat{B}_t^{i/j} - A_t^{i/j}}{\hat{B}_t^{i/j}} - z &> 0, && (18)
\end{aligned}
$$

where z is the cost of exploiting an arbitrage opportunity (expressed as a fraction of synthetic quotes) over and above bid-ask spread costs. Chaboud et al. (2014) argue 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 quotes (ask and bid) for the three currency pairs (EUR/USD, GBP/USD, and EUR/GBP) 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

Exchange rate (i/j)	Bid	Ask	Mid-Quote ((Bid + Ask)/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 \$/€ become $A^{\$/€} = 1.075$ and $B^{\$/€} = 1.074$ (a depreciation of the euro against the dollar). If the quotes of the other currency pairs 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 (17) and (18) 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 (17) and (18).[26]

---

[26]The arrival of a market order cannot create an arbitrage opportunity. For instance, suppose that $A_t^{i/j} > \hat{B}_t^{i/j}$

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 we record $(t_1 - t_0)$ as the duration of the arbitrage opportunity. 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.[27] If instead this price change 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.

---

and $\hat{A}_t^{i/j} > B_t^{i/j}$, so that, at time t, there is no-arbitrage opportunity. Buy (sell) market orders in the market for currency i vs. j can only push $A_t^{i/j}$ upward ($B_t^{i/j}$ downward) because they deplete the limit order book for currency i vs. j on the sell (buy) side. Thus, buy and sell market orders in the market for currency i vs. j do not create arbitrage opportunities. 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.

[27]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).

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.[28] 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 IV 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 IV 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 IV 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 and U.S. 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 constructing the variables used in our tests (see below), we only retain observations from 7:00 to 17:00 GMT.

---

[28]Empirical studies of foreign exchange markets have found that macroeconomic news announcements or headlines news on Reuters are associated with jumps in exchange rates (see Andersen et al. (2003) or Evans and Lyons (2008)).

In robustness checks (not reported), we have checked that our empirical findings are not sensitive to extending the window of observation for price reversals after an arbitrage opportunity terminates.

3.3 Variables of interest

We now 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:

$$\hat{\varphi}_t = \frac{\text{No. of toxic arbitrage opportunities on day } t}{\text{No. of all arbitrage opportunities on day } t}. \tag{19}$$

This ratio is a proxy for φ, i.e., the likelihood that an arbitrage opportunity is toxic. It characterizes the 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. As a proxy for π, we use the daily fraction of toxic arbitrage opportunities that terminate with a trade (the submission of market orders). We denote this fraction by $\hat{\pi}_t^{tox}$ on day t:

$$\hat{\pi}_t^{tox} = \frac{\text{No. of toxic arbitrage opport. that terminate with a trade on day } t}{\text{No. of toxic arbitrage opportunities on day } t}. \tag{20}$$

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: $\hat{\pi}_t^{nontox}$. That is:

$$\hat{\pi}_t^{nontox} = \frac{\text{No. of non-toxic arbitrage opport. that terminate with a trade on day } t}{\text{No. of non-toxic arbitrage opportunities on day } t}. \tag{21}$$

Another important variable in our model is σ (see (13)), the size of toxic arbitrage opportunities. To obtain a proxy for σ, we proceed 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 will 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 of the 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.

<center>Insert Table II here</center>

On average, the daily fraction of toxic arbitrage opportunities is $\hat{\varphi} = 41.5\%$ (s.d.= 10%). 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 (17) and (18)) 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. (2013) 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. (2013)) on their market orders (i.e., liquidity taking orders) in large-cap stocks and much less in small caps. This is of the same order of magnitude as daily revenues on triangular arbitrage opportunities in our sample.[29]

The likelihood that a toxic arbitrage opportunity terminates with a trade ($\hat{\pi}^{tox}$) is 74.1% on average (s.d.= 0.110). Thus, arbitrageurs are on average relatively faster than traders with posted quotes in the limit order book. 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: traders with posted quotes in non initiating currencies have no incentive to cancel their quotes when a non toxic arbitrage opportunity occur. This implies that $\hat{\pi}^{nontox}$ should be larger than $\hat{\pi}^{tox}$, as we find. There are several possible reasons why yet $\hat{\pi}^{nontox}$ is less than 1: (i) market makers with posted quotes in non initiating currencies might sometimes wrongly believe than an arbitrage opportunity is toxic, (ii) the price pressure at the origin of the arbitrage opportunity disappears because the market maker in the initiating currency trades elsewhere (e.g., on EBS), and (iii) our classification of arbitrage opportunities is imperfect.

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)).

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 0.96

---

[29]Another benchmark for daily profits on triangular arbitrage profits are actual daily profits by dealers in FX markets. Bjønnes and Rime (2005) find an average daily profit of about $12,000 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).

basis points. Measures of illiquidity are much higher on EBS for the GBP/USD and EUR/GBP pairs, probably because they are more heavily traded on Reuters D-3000 at the time of our sample. For instance, in EUR/GBP, quoted spreads on Reuters are equal to 1.35 basis points on average vs. 2.5 basis points for EBS. In contrast, EBS is more liquid for the EUR/USD 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 unconditional correlation of the variables used in our tests. Consistent with Implication 1a, measures of illiquidity for the three currency pairs in our sample are positively and significantly correlated with $\hat{\varphi}$, the fraction of arbitrage opportunities. 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_t^*) and measures of illiquidity is positive (consistent with Implication 1b) in the model) but not significantly different from zero. This may stem from the fact that $\hat{\pi}^{tox}$ and bid-ask spreads are jointly determined. For instance, an increase in market activity can simultaneously increase arbitrageurs' speed relative to market makers (as market makers' attention constraints are more likely to be binding) and reduce bid-ask spreads (as market makers are contacted more frequently by liquidity traders). In contrast, 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 $\hat{\pi}^{tox}$ contain different information and are not driven by the same factors. This again is consistent with the idea that market makers in non-initiating pairs should behave differently in toxic and non-toxic arbitrage opportunities. According to our model, they should update their quotes as quickly as possible in toxic arbitrage opportunities whereas 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 also 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 (14)).

3.5 Tests

3.5.1 Is toxic arbitrage a source of illiquidity?

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 1a. To test this implication, we estimate the following equation:

$$illiq_{it} = \omega_i + \xi_t + b_2 vol_{it} + b_3 \hat{\varphi}_t + b_4 \hat{\sigma}_t^{tox} + b_5 trsize_{it} + b_6 nrorders_{it} + b_7 illiq_{it}^{EBS} + \varepsilon_{it}, \quad (22)$$

where $illiq_{it}$ is one of our three proxies for illiquidity for currency i on day t and ω_i and ξ_t are, respectively, currency and time fixed-effects (dummies for each month in our sample). Coefficient b_3 measures the sensitivity of illiquidity to our proxy for φ and should be positive according to Implication 1a. The model also implies that the bid-ask spread of each currency should be larger on days in which the size of toxic arbitrage opportunities, proxied by $\hat{\sigma}_t^{tox}$, is higher. Thus, we expect b_4 to be positive. In estimating these effects, we control for various variables that are known to affect bid-ask spreads: the daily realized volatility (vol_{it}), the daily average trade size ($trsize_{it}$), and the daily number of orders ($nrorders_{it}$). For each illiquidity measure, we also include its EBS counterpart ($illiq_{it}^{EBS}$) in our set of explanatory variables 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 (22) 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

28

toxic is higher, as predicted by Implication 1a. This effect is both statistically and economically significant. For instance, for the effective bid-ask spread, we find that $b_3 = 0.485$ (t-stat=7.38). Thus, an increase of $\hat{\varphi}$ by one standard deviation (i.e., 0.1) is associated with an increase of 0.0485 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, as implied by (13) in the model. 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%.

We then use the forecasting model to decompose $\hat{\varphi}_t$ into an anticipated component and an unanticipated component, $v_{\hat{\varphi}_t}$ (the residual of the forecasting equation). We then reestimate equation (22) with both components as explanatory variables and reports estimates in Column (4), (5) and (6) of Table IV (Panel B). We find that both the anticipated and unanticipated components of $\hat{\varphi}_t$ are positively and significantly associated with all measures of illiquidity.

We cannot discard the possibility that omitted variables drive both variations in our measures of illiquidity and $\hat{\varphi}_t$. However, the fact that illiquidity is also positively related to the unanticipated component of $\hat{\varphi}_t$ alleviates this concern.

3.5.2 Are faster arbitrageurs a source of illiquidity?

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 1b). To this end, we include π^{tox} (our proxy for π) as an additional explanatory variable in our baseline regression (22).

Specifically, we estimate:

$$illiq_{it} = \omega_i + \xi_t + b_1\hat{\pi}_t^{tox} + b_2vol_{it} + b_3\hat{\varphi}_t + b_4\hat{\sigma}_t + b_5trsize_{it} + b_6nrorders_{it} + b_7illiq_{it}^{EBS} + \varepsilon_{it}. \quad (23)$$

Implication 1b 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. To address this issue, we use an instrumental variable (IV) approach to estimate b_1. Implication 2 shows that a shock to arbitrageurs' relative cost of speed (parameter r 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 r only through the effect of r on π^*; see equation (22)).

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 use an algorithm to submit orders on their Reuters terminal rather than enter trading instructions manually using a Reuters keyboard, as was done until July 2003. The introduction of "Reuters AutoQuote API" therefore marked the beginning of algorithmic trading on Reuters since it allowed traders to input Reuters datafeed in their algorithms and let the algorithms make trading decisions.[30] 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 anedoctal 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."*

Initially, Reuters allowed only a limited number of clients to use the AutoQuote API functionality because of capacity constraints as 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

---

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

it more difficult for us to detect its effect.

<center>Insert Figure III about here.</center>

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 III 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, 1st 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 might therefore either increase or decrease arbitrageurs' *relative* speed of reaction (r 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 Auto-Quote 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 (23). 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 (23)). This is plausible because, to our knowledge, there is no obvious mechanism by which technological changes 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:

$$\hat{\pi}_t^{tox} = \omega_i + \xi_t + a_1 AD_t + a_2 vol_{it} + a_3 \hat{\varphi} + a_4 \hat{\sigma}_t + a_5 trsize_{it} + a_6 nrorders_{it} + a_7 illiq_{it}^{EBS} + u_{it}, \quad (24)$$

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.[31] We find that the introduction of AutoQuote on Reuters D-3000 has a significant positive effect on the likelihood

---

[31]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.

<center>31</center>

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 does not seem to be weak (the F statistics is around 16 in all specifications).[32] 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. Moreover, consistent with the model, an increase in the size of arbitrage opportunities ($\hat{\sigma}$) has also a significant and positive effect on all measures of illiquidity. The size of these effects are not different from those reported in Table IV (Panel B).

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 0.07934 basis points (t-stat = 3.7).[33] The economic size of this effect is significant as well since 0.07934 basis points represents 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 size for the currencies considered in our sample is quite large: $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, there are about 4,692 transactions per day in GBP/USD, 2,365 in EUR/USD and 2,841 in EUR/GBP. Hence, the total increase in trading cost per day due to a 1% increase in $\hat{\pi}^{tox}$ is at

---

[32]Bound et al. (1995) (p.446) mention that "F statistics close to 1 should be cause for concern".

[33]We have also estimated equation (24) 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.

least $0.07934 bps \times (\$2.390 \times 4,692 + \$1.655 \times 2,365 + \$1.831 \times 2,841) = \$161,296$ for the three markets in total or about \$40 million per year.[34] Thus, even a small increase in arbitrageurs' speed can be rather costly for other market participants.

As a robustness check, we have also estimated (24) using hourly estimates of each variable in our regressions rather than daily estimates. Qualitatively, the findings are 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 (23) 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 $\hat{\pi}^{tox}$). 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:

$$\log(TTE_t) = \omega_i + \xi_t + a_1 AD_t + a_2 vol_{it} + a_3 \hat{\varphi}_t + a_4 \hat{\sigma}_t^{tox} + a_5 trsize_{it} + a_6 nrorders_{it} + u_{it}, \quad (25)$$

where TTE_t is the average duration of arbitrage opportunities. 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.[35]

Consistent with Implication 3, AutoQuote is associated with a decrease in the duration of toxic arbitrage opportunities by about 6.8%.[36] Similar estimates are obtained when we use the

---

[34]As a point of comparison, Naranjo and Nimalendran (2000) estimate the annualized increased in trading costs due to adverse selection created by unanticipated interventions of the Bundesbank and the FED in currency markets to about \$55 million per year.

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

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

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 again is consistent with the model. In contrast, the effect of $\hat{\varphi}$ on the duration of all arbitrage opportunities is statistically non significant.

Overall, the empirical findings suggest that pricing efficiency can come 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. Specifically, 51% of all toxic arbitrage opportunities can be traced back to a jump in the EUR/USD rate. In contrast, 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, 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

---

opportunities.

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 effect of the fraction of toxic arbitrage opportunities in the three currency pairs ($\hat{\varphi}$) 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 significant effect on the quoted and the effective spreads of the EUR/USD pair (and no significant effect on the slope of the limit order book 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 seems sometimes stronger than in other currency pairs.

<center>Insert Tables VII about here</center>

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 (a textbook arbitrage opportunity).

A natural question is whether our measures of market makers' exposure to arbitrageurs' picking off risk ($\hat{\varphi}$ and $\hat{\pi}^{tox}$) is distinct from other existing 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 require a greater compensation for 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

<center>35</center>

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., non-farm 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).[37] Thus, on each day t, we compute a VPIN metric ($VPIN_t^i$) for each currency pair i in our sample. Specifically, following Easley et al. (2012)'s methodology, 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. The $VPIN_t^i$ metric for currency pair i and day t is then

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

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.[38] We find a significant a positive correlation among the

---

[37]VPIN is an alternative to the PIN measure that has been extensively used in finance; see Easley et al. (2012).

[38]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).

VPIN measures for the three currency pairs in our sample. In contrast, 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 macro-economic 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).[39] 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

The Law of One Price frequently breaks down at high frequency. The faster arbitrageurs correct deviations from the Law of One Price, the higher is pricing efficiency. This conclusion does not necessarily hold for liquidity. Indeed, a fast response of arbitrageurs to opportunities caused by asynchronous price adjustments of related assets to information ("toxic arbitrage opportunities") is a source of illiquidity because it increases liquidity suppliers' risk of trading at stale quotes. Through this channel, high speed arbitrage can impair market liquidity. We provide evidence for this channel 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 frequency of toxic arbitrage opportunities is higher and (b) the frequency with which arbitrageurs successfully exploit these opportunities is higher.

---

[39]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.

One important take away of our analysis is that the nature of short lived arbitrage opportunities matters. If, in today's markets, these opportunities are predominantly due to stale quotes rather than price pressures then there is ground for regulatory intervention. In particular, limiting arbitrageurs' speed (e.g., by randomizing the execution time of their market orders) should improve liquidity, at the cost of a small increase in the duration of arbitrage opportunities.[40] Future empirical research should therefore explore more systematically whether short lived arbitrage opportunities are mainly due to price pressure effects or stale quotes. This is a pre-requisite to better understand the social value of high speed arbitrage.

---

[40]The negative effect of arbitrage on liquidity should be a concern because, eventually, it reduces investors' ability to realized their gains from trade and therefore investors' welfare.

Appendix

Derivation of Implications 3 and 4

First, substituting the expression for the equilibrium spread, S_X^* (given in (13)) in (10), we obtain the arbitrageur's speed, $\gamma^*(S_X^*)$ in equilibrium:

$$\gamma^*(S_X^*) = \frac{(\varphi \alpha r)(1 - \alpha(2\varphi - 1))}{(c^a(1+r)^2)(2\varphi \alpha \pi^*(r)(2 - \pi^*(r)) + (1 - \alpha(2\varphi - 1)))} \sigma \qquad (26)$$

Substituting (26) in (14), we obtain:

$$E(D) = \frac{(2\varphi \alpha(\frac{2+r}{1+r}) + (1 - \alpha(2\varphi - 1))(\frac{2+r}{r}))(c^a(1 - \varphi) + c^m)}{\varphi \alpha(1 - \alpha(2\varphi - 1))\sigma} \qquad (27)$$

We deduce that $\frac{\partial E(D)}{\partial r} > 0$. Using the fact that $r = \frac{c^m}{c^a}$, we deduce from (27) that a decrease in c^a and c^m that eventually result in an increase in r lowers the expected duration of an arbitrage opportunity (Implication 3).

Using (26), we also deduce $\gamma^*(S_X^*)$ increases with φ if $\alpha(4\varphi - 1) < 1$. This condition is automatically satisfied when $\varphi \le 1/2$ or when $\varphi > 1/2$ and $\alpha < (4\varphi - 1)^{-1}$. We deduce from (14) 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 r (Panel C). It also shows the duration of arbitrage opportunities (in seconds) as (i) a function of the likelihood of a toxic arbitrage opportunity, φ (Panel D) and (b) the cost of speed for the arbitrageur, c^a (Panel E). In all cases, we set $\sigma = 3.5bps$ and $\alpha = 0.1$.
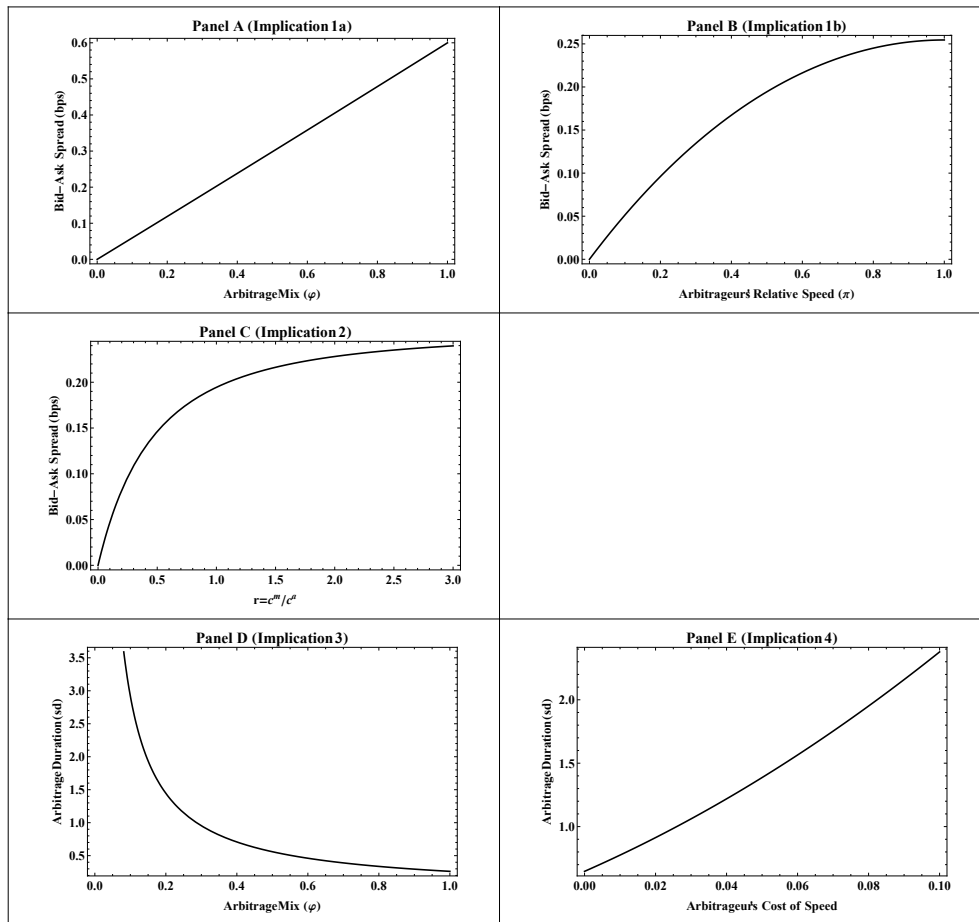
Figure II: Toxic vs. Non-Toxic Arbitrage Opportunities

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
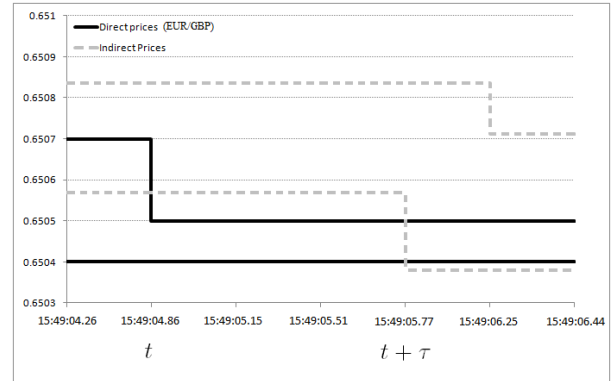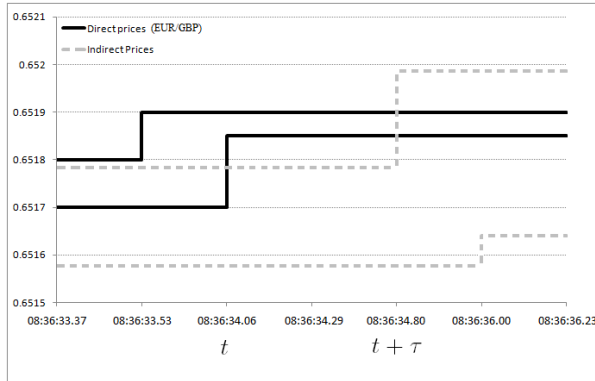
Panel B: Non-Toxic Arbitrage Opportunities
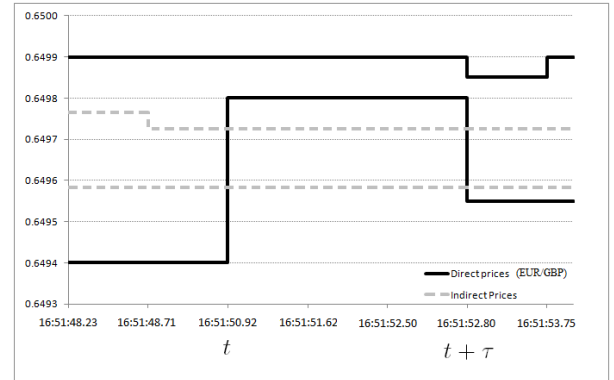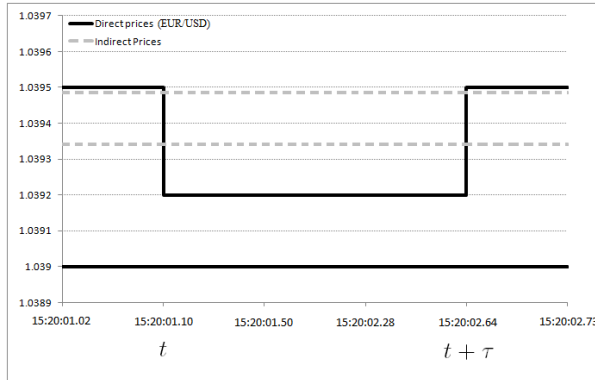
Figure III: AutoQuote and Order to Trade Ratio.

This figure shows the evolution of the order to trade ratio the ratio (defined as daily number of orders to 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.
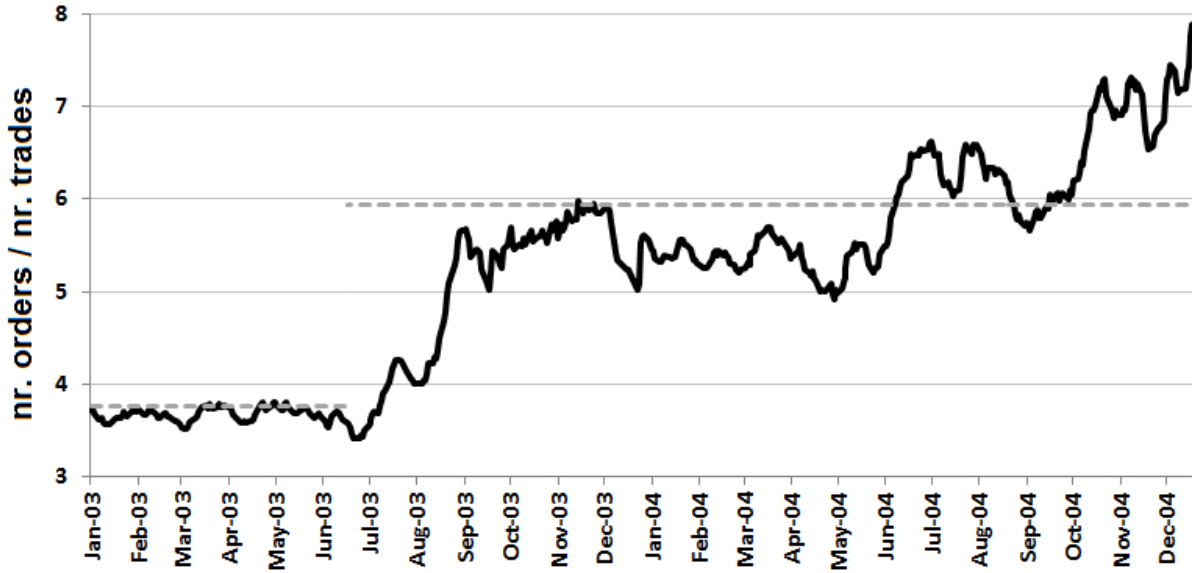
Figure IV: Number of Arbitrage Opportunities

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
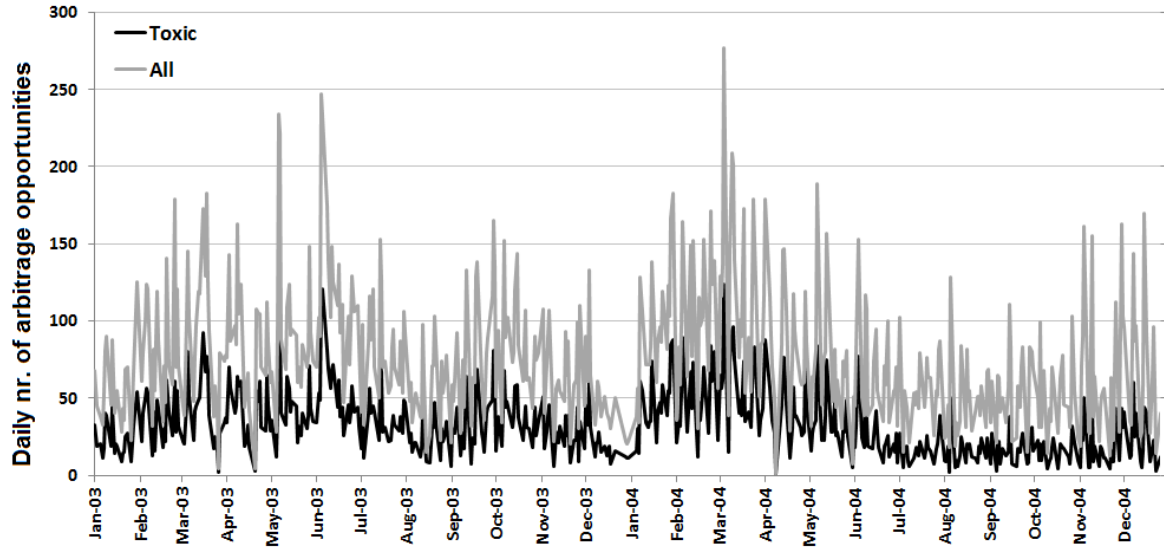
Panel B: Intraday Pattern in the Number of Arbitrage Opportunities

43
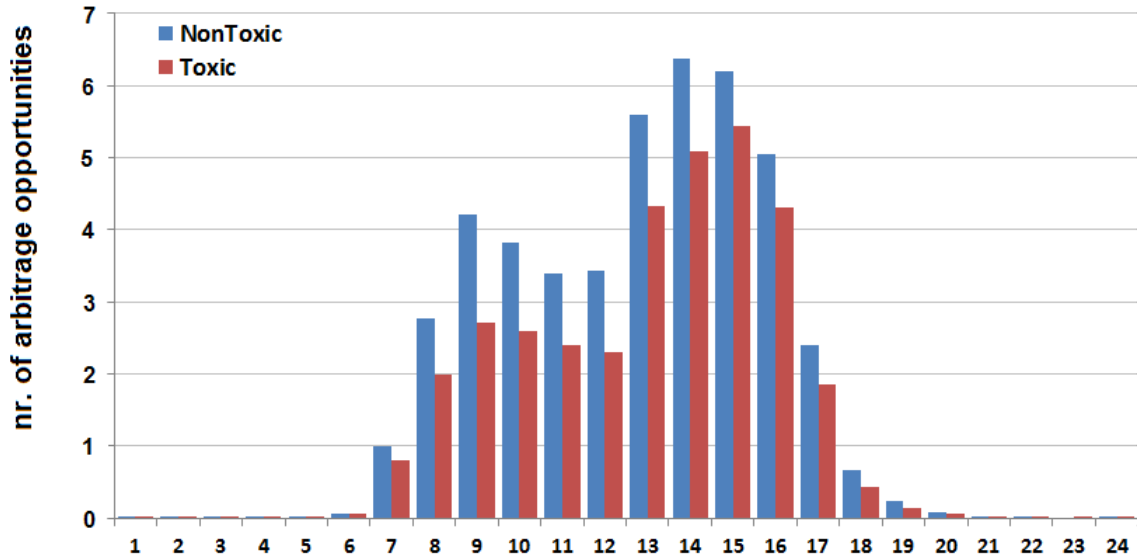
Tables

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 in day t; $\hat{\pi}_t^{tox}$ ($\hat{\pi}_t^{nontox}$) is the number of toxic (non-toxic) arbitrage opportunities in day t that terminate with a trade divided by the total number of toxic (non-toxic) arbitrage opportunities in this day; $\hat{\varphi}_t$ is the number of toxic (non-toxic) arbitrage opportunities in day t divided by the number of arbitrage opportunities in this day; $\hat{\sigma}_t^{tox}$ ($\hat{\sigma}_t^{nontox}$) is the average size of toxic (non-toxic) arbitrage opportunities in 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 in 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: vol_{it} is the realized volatility (in percentage) of 5-minutes returns for currency pair i in 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	17	28	43	124	
$\hat{\pi}_t^{tox}$	0.741	0.110	0	0.685	0.743	0.804	1.000	
$\hat{\varphi}_t$	0.415	0.100	0.080	0.354	0.429	0.482	1.000	
$\hat{\sigma}_t^{tox}$	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	27	40	55	740	-6.75
$\hat{\pi}_t^{nontox}$	0.807	0.082	0.412	0.755	0.800	0.867	1.000	-10.7
$\hat{\varphi}_t^{nontox}$	0.617	0.119	0.340	0.538	0.598	0.676	1.577	-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
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	
$spread_{GUt}^{EBS}$	5.253	1.157	2.421	4.687	5.093	5.580	12.82	
$spread_{EUt}^{EBS}$	1.139	0.046	1.050	1.103	1.136	1.164	1.336	
$spread_{EGt}^{EBS}$	2.520	0.807	1.431	2.064	2.376	2.803	11.11	
$espread_{GUt}^{EBS}$	5.112	3.467	2.375	3.237	4.165	5.796	27.90	
$espread_{EUt}^{EBS}$	0.998	0.065	0.899	0.958	0.985	1.020	1.420	
$espread_{EGt}^{EBS}$	2.082	1.847	0.901	1.245	1.589	2.189	24.23	
$slope_{GUt}^{EBS}$	3.860	3.246	1.125	2.377	3.122	4.332	47.97	
$slope_{EUt}^{EBS}$	0.296	0.041	0.205	0.266	0.294	0.323	0.441	
$slope_{EGt}^{EBS}$	1.833	2.448	0.444	1.011	1.428	1.980	39.47	

Table II continued.

Variable	Mean	Std.Dev.	Min	Q1	Median	Q3	Max
Panel D: Common Variables							
vol_{GU}	0.683	0.268	0.117	0.532	0.622	0.753	2.456
vol_{EU}	0.827	0.386	0.258	0.616	0.744	0.920	4.363
vol_{EG}	0.387	0.094	0.203	0.325	0.381	0.440	1.256
$nrorders_{GU}$	17.65	6.091	0.576	12.51	17.62	22.45	32.22
$nrorders_{EU}$	19.05	6.831	0.188	14.78	18.15	22.88	44.07
$nrorders_{EG}$	14.77	5.810	0.307	9.326	16.07	19.41	28.93
$trsize_{GU}$	1.386	0.043	1.247	1.357	1.382	1.415	1.509
$trsize_{EU}$	1.401	0.056	1.000	1.365	1.396	1.434	1.605
$trsize_{EG}$	1.548	0.076	1.294	1.497	1.541	1.591	1.853
$nrtr_{GU}$	4.692	1.505	0.175	3.634	4.523	5.639	9.611
$nrtr_{EU}$	2.365	0.707	0.027	1.859	2.377	2.870	4.103
$nrtr_{EG}$	2.841	0.811	0.068	2.301	2.761	3.318	6.329
Nr. days				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 in day t that terminate with a trade divided by the total number of toxic (non-toxic) arbitrage opportunities in day t; $\hat{\sigma}_t^{tox}$ is the number of toxic (resp., non toxic) arbitrage opportunities in day t divided by the number of arbitrage opportunities in that day; $\hat{\varphi}_t$ is the average size of toxic arbitrage opportunities in 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 slope of the limit order book in currency pair i on day t. Superscript EBS is used for illiquidity measures computed using EBS data. $illiq^{EBS}t$ 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
$illiq^{EBS}t$				**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:

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

where indexes GU, EU, EG refer to the GBP/USD, EUR/USD, and EUR/GBP currency pairs, respectively and φ_t is the number of toxic arbitrage opportunities in day t divided by the number of arbitrage opportunities in this day. In panel B, we report OLS estimates of the following two regressions:

$$illiq_{it} = \omega_i + \xi_t + b_1 \hat{\varphi}_t + b_2 vol_{it} + b_3 \hat{\sigma}_t^{tox} + b_4 trsize_{it} + b_5 nrorders_{it} + b_6 illiq_{it}^{EBS} + \varepsilon_{it} \text{ and}$$

$$illiq_{it} = \omega_i + \xi_t + c_1 v_{\hat{\varphi},t} + c_2 fitted_{\hat{\varphi},t} + c_3 vol_{it} + c_4 \hat{\sigma}_t^{tox} + c_5 trsize_{it} + c_6 nrorders_{it} + c_7 illiq_{it}^{EBS} + \varepsilon_{it},$$

where $fitted_{\varphi,t}$ and $v_{\hat{\varphi},t}$ are, respectively, the predicted value of φ_t and the residual from the regression estimated in Panel A; $\hat{\sigma}_t^{tox}$ is the average size of toxic arbitrage opportunities in day t (in basis points); $spread_{it}$ is the average quoted bid-ask spread (in basis points) in currency pair i on day t; $espread_{it}$ (in basis points) 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 for measures of these variables computed using EBS data; vol_{it} is the realized volatility (in percentage) of 5-minutes returns for currency pair i in 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%

Panel B

	(1)	(2)	(3)	(4)	(5)	(6)
	spread	espread	slope	spread	espread	slope
$\hat{\varphi}$	0.622 (7.47)	0.485 (7.38)	0.407 (9.34)			
$v_{\hat{\varphi}}$				0.425 (4.49)	0.342 (4.05)	0.251 (4.92)
$fitted_{\hat{\varphi}}$				0.816 (5.58)	0.646 (5.20)	0.635 (8.48)
vol	0.299 (8.01)	0.370 (7.45)	0.177 (9.44)	0.283 (7.63)	0.357 (7.13)	0.167 (9.23)
$\hat{\sigma}^{tox}$	0.158 (10.5)	0.187 (5.26)	0.074 (10.1)	0.157 (10.2)	0.189 (5.11)	0.073 (9.45)
$trsize$	-0.341 (-3.37)	-0.276 (-2.17)	-0.385 (-7.25)	-0.352 (-3.35)	-0.287 (-2.10)	-0.389 (-7.07)
$nrorders$	-0.011 (-8.52)	-0.009 (-7.00)	-0.007 (-10.0)	-0.013 (-8.72)	-0.010 (-7.11)	-0.007 (-10.0)
$illiq^{EBS}$	0.026 (3.47)	-0.004 (-1.80)	0.002 (2.16)	0.021 (2.83)	-0.004 (-2.10)	0.002 (1.85)
$Adj.R^2$	86.54%	87.81%	82.90%	86.96%	88.03%	83.80%
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 IV estimates of the following regression for $i \in \{GU, EU, EG\}$:

$illiq_{it} = \omega_i + \xi_t + b_1\hat{\pi}_t^{tox} + b_2vol_{it} + b_3\hat{\varphi}_t + b_4\hat{\sigma}_t + b_5trsize_{it} + b_6nrorders_{it} + b_7illiq_{it}^{EBS} + \varepsilon_{it}$, 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 of the IV is: $\hat{\pi}_t^{tox} = \omega_i + \xi_t + a_1AD_t + a_2vol_{it} + a_3\hat{\varphi}_t + a_4\hat{\sigma}_t + a_5trsize_{it} + a_6nrorders_{it} + a_7illiq_{it}^{EBS} + u_{it}$, where 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 arbitrages of day t; $\hat{\varphi}_t$ is the number of toxic arbitrage opportunities in day t divided by the number of arbitrage opportunities in this day; $\hat{\sigma}_t^{tox}$ is the average size of toxic arbitrage opportunities in day t (in basis points); $spread_{it}$ is the average quoted bid-ask spread (in basis points) in currency pair i on day t; $espread_{it}$ (in basis points) 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 for measures of these variables computed using EBS data; vol_{it} is the realized volatility (in percentage) of 5-minutes returns for currency pair i in 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.

	spread		espread		slope	
	1^{st} stage	2^{nd} stage	1^{st} stage	2^{nd} stage	1^{st} stage	2^{nd} stage
$\hat{\pi}^{tox}$		7.934 (3.91)		3.443 (3.70)		4.526 (3.96)
AD	0.040 (4.09)		0.042 (4.12)		0.040 (4.10)	
vol	-0.009 (-0.75)	0.374 (3.72)	-0.009 (-0.77)	0.401 (8.65)	-0.009 (-0.76)	0.220 (3.87)
$\hat{\varphi}$	-0.011 (-0.31)	0.691 (2.29)	-0.011 (-0.31)	0.511 (3.68)	-0.010 (-0.28)	0.445 (2.61)
$\hat{\sigma}^{tox}$	-0.011 (-2.14)	0.238 (4.93)	-0.012 (-2.17)	0.221 (9.94)	-0.011 (-2.11)	0.120 (4.39)
$trsize$	0.002 (0.66)	-0.128 (-0.30)	0.001 (0.84)	-0.196 (-0.98)	0.001 (0.76)	-0.265 (-1.09)
$nrorders$	0.014 (0.27)	-0.004 (-0.77)	0.012 (0.22)	-0.006 (-2.62)	0.016 (0.30)	-0.003 (-1.01)
$illiq^{EBS}$	-0.003 (-3.88)	0.021 (0.79)	-0.003 (-3.85)	-0.002 (-0.43)	-0.003 (-3.89)	0.001 (0.08)
$Adj.R^2$	2.34%	34.40%	2.34%	62.18%	2.35%	25.56%
$Fstat$	16.7		16.9		16.8	
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:

$$\log(TTE_t) = c_i + \xi_t + a_1 AD_t + a_2 vol_{it} + a_3 \hat{\varphi}_t + a_4 \hat{\sigma}_t + a_5 trsize_{it} + a_6 nrorders_{it} + u_{it},$$

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 in day t divided by the number of arbitrage opportunities in 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 in 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	-0.068 (-3.04)	-0.057 (-2.93)
vol	-0.084 (-3.15)	-0.105 (-4.53)
$\hat{\varphi}$	-0.248 (-2.95)	0.050 (0.68)
$\hat{\sigma}^{tox}$	0.070 (6.59)	0.085 (9.22)
$trsize$	0.022 (0.18)	0.015 (0.14)
$nrorders$	-0.012 (-7.29)	-0.010 (-7.40)
$Adj.R^2$	21.24%	33.33%

Table VII: Currency-Level Tests

In this table, we estimate the following regression for each currency pair separately:

$illiq_{it} = \omega_i + \xi_t + b_1\hat{\pi}_t^{tox} + b_2 vol_{it} + b_3\hat{\varphi}_t + b_4\hat{\sigma}_t + b_5 trsize_{it} + b_6 nrorders_{it} + b_7 illiq_{it}^{EBS} + \varepsilon_{it} \quad for \quad i \in \{GU, EU, EG\}$,
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 of the IV is: $\hat{\pi}_t^{tox} = \omega_i + \xi_t + a_1 AD_t + a_2 vol_{it} + a_3\hat{\varphi}_t + a_4\hat{\sigma}_t + a_5 trsize_{it} + a_6 nrorders_{it} + a_7 illiq_{it}^{EBS} + u_{it}$, where 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 arbitrages on day t; $\hat{\varphi}_t$ is the number of toxic arbitrage opportunities in day t divided by the number of arbitrage opportunities in this day; $\hat{\sigma}_t^{tox}$ is the average size of toxic arbitrage opportunities in day t (in basis points); $spread_{it}$ is the average quoted bid-ask spread (in basis points) in currency pair i on day t; $espread_{it}$ (in basis points) 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 computed using EBS data; vol_{it} is the realized volatility (in percentage) of 5-minutes returns for currency pair i in 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.

		spread		espread		slope	
		1^{st} stage	2^{nd} stage	1^{st} stage	2^{nd} stage	1^{st} stage	2^{nd} stage
GBP/USD	$\hat{\pi}^{tox}$		5.216 (5.12)		3.024 (5.71)		2.008 (5.52)
	AD	0.058 (6.79)		0.062 (6.39)		0.059 (7.01)	
	vol	0.010 (0.48)	0.231 (2.00)	0.008 (0.36)	0.345 (5.06)	0.010 (0.46)	0.091 (2.12)
	$\hat{\varphi}$	-0.033 (-0.66)	0.946 (3.63)	-0.034 (-0.68)	0.559 (3.95)	-0.033 (-0.66)	0.417 (4.09)
	$\hat{\sigma}^{tox}$	-0.023 (-2.36)	0.273 (4.91)	-0.023 (-2.44)	0.210 (7.75)	-0.023 (-2.27)	0.127 (7.05)
	$trsize$	0.137 (2.38)	-0.915 (-2.00)	0.139 (2.32)	-0.116 (-0.40)	0.139 (2.34)	-0.721 (-3.43)
	$nrorders$	-0.004 (-6.10)	0.001 (0.54)	-0.004 (-5.65)	0.002 (0.77)	-0.004 (-6.44)	-0.004 (-3.05)
	$illiq^{EBS}$	0.001 (0.48)	0.030 (1.34)	0.001 (1.10)	0.006 (1.18)	0.001 (0.53)	0.006 (2.93)
	$Adj.R^2$	1.37%	8.80%	1.40%	20.78%	1.37%	15.42%
	$Fstat$	46.1		32.6		49.1	
EUR/GBP	$\hat{\pi}^{tox}$		1.427 (4.61)		0.678 (3.64)		0.389 (1.86)
	AD	0.064 (7.74)		0.070 (7.92)		0.065 (7.23)	
	vol	0.081 (3.11)	0.964 (12.7)	0.066 (2.26)	0.782 (12.4)	0.093 (3.03)	0.552 (11.9)
	$\hat{\varphi}$	-0.059 (-1.36)	0.500 (10.2)	-0.056 (-1.24)	0.311 (13.3)	-0.058 (-1.32)	0.224 (13.7)
	$\hat{\sigma}^{tox}$	-0.024 (-3.73)	0.096 (6.06)	-0.025 (-3.77)	0.068 (6.60)	-0.024 (-3.72)	0.036 (5.79)
	$trsize$	0.014 (0.33)	-0.326 (-5.75)	0.015 (0.35)	-0.093 (-2.29)	0.013 (0.31)	-0.321 (-21.5)
	$nrorders$	-0.005 (-7.17)	-0.010 (-14.9)	-0.005 (-7.19)	-0.010 (-13.8)	-0.005 (-7.46)	-0.010 (-20.6)
	$illiq^{EBS}$	0.005 (2.02)	0.026 (2.17)	0.004 (4.24)	0.006 (2.79)	0.002 (5.61)	0.001 (0.35)
	$Adj.R^2$	1.24%	48.41%	1.41%	62.37%	1.35%	67.96%
	$Fstat$	59.9		62.7		52.3	
EUR/USD	$\hat{\pi}^{tox}$		0.402 (0.45)		2.818 (3.32)		4.679 (2.36)
	AD	0.041 (3.12)		0.034 (5.35)		0.029 (2.54)	
	vol	0.000 (-0.03)	0.015 (0.54)	0.014 (1.44)	-0.053 (-2.06)	-0.003 (-0.39)	0.062 (2.09)
	$\hat{\varphi}$	0.012 (0.27)	0.199 (2.07)	0.015 (0.31)	0.312 (1.58)	0.017 (0.36)	0.051 (0.20)
	$\hat{\sigma}^{tox}$	-0.011 (-1.13)	0.091 (5.05)	-0.008 (-0.93)	0.152 (5.81)	-0.015 (-1.50)	0.120 (1.83)
	$trsize$	0.081 (1.01)	-0.227 (-0.88)	0.096 (1.24)	-0.683 (-1.29)	0.077 (0.99)	-0.641 (-1.61)
	$nrorders$	-0.002 (-2.43)	-0.009 (-3.89)	-0.001 (-1.82)	-0.020 (-6.45)	-0.002 (-2.33)	0.001 (0.69)
	$illiq^{EBS}$	-0.234 (-1.98)	6.679 (17.1)	-0.240 (-4.23)	4.059 (11.3)	-0.121 (-1.01)	4.396 (9.48)
	$Adj.R^2$	0.41%	74.82%	0.62%	44.32%	0.08%	13.79%
	$Fstat$	9.73		28.6		6.45	
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 $i \in \{GU, EU, EG\}$:

$$illiq_{it} = \omega_i + \xi_t + b_1 \hat{\pi}_t^{tox} + b_2 vol_{it} + b_3 \hat{\varphi}_t + b_4 \hat{\sigma}_t^{tox} + b_5 trsize_{it} + b_6 nrorders_{it} + b_7 illiq_{it}^{EBS} + b_8 VPIN_{it}$$

$$+ b_9 macro_t^{US} + b_{10} macro_t^{UK} + b_{11} macro_t^{EMU} + u_{it}, \quad for \quad i \in \{GU, EU, EG\}.$$

We instrument $\hat{\pi}_t^{tox}$ with the introduction of Reuters D-3000 AutoQuote (see the text). The first stage regression of the IV is:

$$\hat{\pi}_t^{tox} = \omega_i + \xi_t + a_1 AD_t + a_2 vol_{it} + a_3 \hat{\varphi}_t + a_4 \hat{\sigma}_t^{tox} + a_5 trsize_{it} + a_6 nrorders_{it} + a_7 illiq_{it}^{EBS} + a_8 VPIN_{it}$$

$$+ a_{10} macro_t^{US} + a_{11} macro_t^{UK} + a_{12} macro_t^{EMU} + u_{it} \quad for \quad i \in \{GU, EU, EG\}.$$

Indexes GU, EU, and EG refer to the GBP/USD, EUR/USD, and EUR/GBP currency pairs, respectively. $\hat{\pi}_t^{tox}$ is the number of toxic arbitrage opportunities that terminate with a trade on day t divided by the total number of toxic 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; $\hat{\sigma}_t^{tox}$ is the average size of arbitrage opportunities in day t (in basis points); $spread_{it}$ is the average quoted bid-ask spread (in basis points) in currency pair i on day t; $espread_{it}$ (in basis points) 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 computed using EBS data; vol_{it} is the realized volatility (in percentage) of 5-minutes returns for currency pair i in 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; $VPIN_{it}$ 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.

	spread		espread		slope	
	1^{st} stage	2^{nd} stage	1^{st} stage	2^{nd} stage	1^{st} stage	2^{nd} stage
$\hat{\pi}^{tox}$		8.007 (3.74)		3.359 (3.51)		4.692 (3.78)
AD	0.038 (-3.90)		0.040 (3.94)		0.038 (3.91)	
vol	-0.001 (-0.02)	0.321 (3.16)	-0.001 (-0.04)	0.419 (9.14)	-0.001 (-0.03)	0.185 (3.13)
$\hat{\varphi}$	-0.015 (-0.42)	0.744 (2.44)	-0.015 (-0.42)	0.517 (3.78)	-0.014 (-0.40)	0.470 (2.67)
$\hat{\sigma}^{tox}$	-0.021 (-3.53)	0.309 (4.76)	-0.021 (-3.57)	0.210 (7.24)	-0.020 (-3.51)	0.164 (4.35)
$trsize$	0.044 (0.82)	-0.447 (-1.01)	0.042 (0.77)	-0.195 (-0.97)	0.046 (0.85)	-0.412 (-1.62)
$nrorders$	-0.003 (-3.71)	0.001 (0.24)	-0.003 (-3.68)	-0.005 (-1.93)	-0.003 (-3.72)	-0.001 (-0.28)
$illiq^{EBS}$	0.002 (0.55)	0.025 (0.93)	0.001 (0.81)	-0.001 (-0.23)	0.001 (0.59)	0.001 (0.21)
$VPIN$	-0.039 (-0.45)	1.142 (1.56)	-0.039 (-0.44)	0.329 (1.00)	-0.041 (-0.46)	0.307 (0.73)
$macro^{US}$	-0.002 (-1.75)	0.017 (1.92)	-0.002 (-1.76)	0.008 (1.98)	-0.002 (-1.76)	0.010 (1.88)
$macro^{UK}$	-0.012 (-2.16)	0.095 (1.74)	-0.012 (-2.17)	0.038 (1.56)	-0.012 (-2.15)	0.053 (1.69)
$macro^{EMU}$	-0.010 (-1.80)	0.081 (1.55)	-0.010 (-1.78)	0.036 (1.53)	-0.010 (-1.80)	0.044 (1.45)
$Adj.R^2$	3.14%	33.88%	3.16%	62.41%	3.14%	23.97%
$Fstat$	15.2		15.2		15.3	
Currency pair FE	YES		YES		YES	
Month dummies	YES		YES		YES	

References

Akram, F., Rime, D. and 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. and Vega, C.: 2003, Micro effects of macro announcements: Real-time price discovery in foreign exchange, *American Economic Review* **93**, 38–62.

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

Bjønnes, G., Osler, C. and Rime, D.: 2011, Sources of information advantage in the foreign exchange market, *Working paper* .

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

Bound, J., Jaeger, D. A. and Baker, R. M.: 1995, Problems with instrumental variables estimation when the correlation between the instruments and the endogenous explanatory variable is weak, *Journal of the American Statistical Association* **90**, 443–450.

Brogaard, J., Hendershott, T. and Riordan, R.: 2013, High frequency trading and price discovery, *Review of Financial Studies* . Forthcoming.

Budish, E., Cramton, E. and 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. and Vega, C.: 2014, Rise of the machine: Algorithmic trading in the foreign exchange market, *Journal of Finance* . Forthcoming.

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

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

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

Duffie, D.: 2010, Presidential address: Asset price dynamics with slow moving capital, *Journal of Finance* **65**(4), 1237–1267.

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

Evans, M. and Lyons, R.: 2008, How is macro news transmitted to exchange rates?, *Journal of Financial Economics* **88**(1), 26–50.

Fenn, D., Howison, S., Mcdonald, M., Williams, S. and Johnson, N.: 2009, The mirage of triangular arbitrage in the spot foreign exchange market, *International Journal of Theoretical and Applied Finance* **12**(8), 1105–1123.

Fong, W., Valente, G. and 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. and Rosu, J.: 2012, News trading and speed, *Working paper, HEC Paris* .

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

Garvey, R. and Wu, F.: 2010, Speed, distance, and electronic trading: New evidence on why location matters, *Journal of Financial Markets* **13**(4), 367–396.

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

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

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

Grossman, S. and Miller, M.: 1988, Liquidity and market structure, *Journal of Finance* **43**(3), 617–637.

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

Hendershott, T. and Menkveld, A.: 2014, Price pressures, *Journal of Financial Economics* . forthcoming.

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

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

Hoffman, P.: 2012, A dynamic limit order market with fast and slow traders, *Working paper* .

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

Holden, C., Jacobsen, S. and Subrahmanyam, A.: 2014, The empirical analysis of liquidity, *Working Paper, University of Indiana* .

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

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

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

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

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

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

Lyons, R. K.: 1995, Tests of microstructural hypotheses in the foreign exchange market, *Journal of Financial Economics* **39**, 321–351.

Mancini-Griffoli, T. and 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. and Young, M.: 2008, Exploitable arbitrage opportunities exist in the foreign exchange market, *American Finance Association Annual Meeting, New Orleans.*

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

Moallemi, C. and Saglam, M.: 2013, The cost of latency in high frequency trading, *Working Paper, Columbia University* .

Naranjo, A. and Nimalendran, M.: 2000, Government intervention and adverse selection costs in foreign exchange markets, *Review of Financial Studies* **13**(2), 453–477.

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

Pagnotta, E. and Phillipon, T.: 2013, Competing on speed, *Working Paper, NYU Stern School of Business* .

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

Pierron, A.: 2007, Electronic platforms in foreign exchange trading, *Technical report*, Celent, Paris.

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

Shive, S. and Schultz, P.: 2010, Mispricing of dual class shares: Profit opportunities, arbitrage, and trading, *Journal of Financial Economics* **98**(4), 524–549.

Shkilko, A. V., Van Ness, B. F. and Van Ness, R. A.: 2008, Locked and crossed markets on NASDAQ and the NYSE, *Journal of Financial Markets* **11**, 308–337.

Sussman, A., Tabb, L. and 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.