

Toxic Arbitrage*

Thierry Foucault[†]

Roman Kozhan[‡]

Wing Wah Tham[§]

March, 2014

Abstract

High frequency arbitrage opportunities arise when the price of one asset follows, with a lag, changes in the value of another related asset due to information arrival. These opportunities are toxic because they expose liquidity suppliers to the risk of being picked off by arbitrageurs. Hence, more frequent toxic arbitrage opportunities and a faster arbitrageurs' response to these opportunities impair liquidity. We find support for these predictions using high frequency triangular arbitrage opportunities in the FX market. In our sample, a 1% increase in the likelihood that a toxic arbitrage terminates with an arbitrageur's trade (rather than a quote update) raises bid-ask spreads by about 4%.

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

JEL Classification: D50, F31, G10

*We thank Mark Van Achter, Alain Chaboud, Geir Bjørnnes, Jean-Edouard Colliard, Pierre Colin-Dufresnes, Arie Gozluklu, Terry Hendershott, Frank de Jong, Albert P. Kyle, Olga Lebedeva, Bruce Lehmann, Katya Malinova, Michael Moore, Angelo Ranaldo, Dagfinn Rime, Marco Pagano, Andreas Park, Vikas Raman, Elvira Sojli, Clara Vega, Kumar Venkataraman. We are also grateful to seminar participants at the Norwegian School of Business, the 9th Central Bank Workshop on the Market Microstructure of Financial Markets, the BIRS workshop on modeling high frequency trading activity, the 6th Erasmus Liquidity Conference and the Conference on Liquidity and Arbitrage Trading in Geneva, EIEF and Consob.

[†]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 expensive asset and selling the cheap 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 can persist over time (see Gromb and Vayanos (2010)). However, these frictions are less likely to play out for high frequency arbitrage opportunities, that is, those lasting *fractions of a second*.¹ For such opportunities, attention costs and technological constraints on traders’ speed are the main barrier to the Law of One Price. This barrier is falling as high frequency arbitrageurs invest massively in fast trading technologies to exploit ever faster very short-lived arbitrage opportunities. This increase in arbitrageurs’ speed is a source of concern for regulators and market participants (see U.S. Securities and Exchange Commission (2010), Section B, p.51 or Halder (2011)).

Can faster arbitrage harm market quality? Why? Our paper addresses these questions. At first glance, regulators’ concerns seem misplaced. High frequency arbitrageurs should strengthen price efficiency by accelerating the speed at which arbitrage opportunities vanish. In addition, by responding faster to transient supply and demand shocks (“price pressures”), fast arbitrageurs should enhance market liquidity (see, for instance, Holden (1995), Gromb and Vayanos (2002), and Gromb and Vayanos (2010)).²

However, high frequency arbitrage opportunities do not arise only because of transient supply and demand shocks. Asynchronous adjustments in asset prices following new information can also cause temporary LOP violations. 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”). Through this channel, fast arbitrage might reduce, rather than increase, liquidity because

¹Such opportunities are very frequent. They arise both from the proliferation of derivatives securities (e.g., ETFs) and market fragmentation (the fact that the same asset trades in multiple different platforms).

²For instance, Gromb and Vayanos (2002) write (on p.362): *“In our model, arbitrage activity benefits all investors. This is because through their trading, arbitrageurs bring prices closer to fundamentals and supply liquidity to the market.”*

market makers require a compensation for the risk of being picked off (Copeland and Galai (1983)).

Consider, for instance, two “market makers” (or limit order books) A and B trading the same asset and suppose that good news regarding the asset arrives. One market maker, say A, instantaneously adjusts his bid and ask quotes to reflect the news while B is slower in adjusting his quotes. 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 on this trade, 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 they raise adverse selection costs for market makers.⁴ Through this channel, fast arbitrage might indeed impair market liquidity. In this paper, we provide evidence on this channel. Specifically, we develop a measure of market makers’ exposure to picking off risk by arbitrageurs and we use this measure to test whether an increase in arbitrageurs’ speed increases adverse selection costs borne by market makers.

To guide our empirical analysis, we rely on a simple model of cross-market arbitrage. Our model has three main features: (a) traders’ speed of reaction to arbitrage opportunities is endogenous and negatively related to market illiquidity, (b) market illiquidity increases with arbitrageurs’ speed, and (c) the duration of arbitrage opportunities is endogenous and positively related to illiquidity. In this model, as in the data, an arbitrage opportunity terminates either when one arbitrageur exploits the opportunity or market makers update their quotes. The model predicts that illiquidity should increase with (i) the fraction of all arbitrage opportunities that are toxic and (ii) the probability that a toxic arbitrage opportunity terminates with an arbitrageur’s trade, because this probability increases with arbitrageurs’ relative speed. We call this probability $PTAT^c$ (“probability of a trade conditional on a toxic arbitrage”). This probability measures market makers’ effective exposure to toxic arbitrage (i.e., the likelihood

³In equity markets, participants refer to cases in which the bid price of an asset exceeds the offer price for the same asset as a “crossed market”. 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. Storckenmaier and Wagener (2011) consider a sample of transactions for FTSE 100 stocks traded on the LSE, Chi-X, BATS and Turquoise. They find that in April/May 2009 (2010), quotes on these platforms were crossed for 16 minutes per trading day with an average duration of 19.8 seconds.

⁴Our 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.”

that their stale quotes are indeed picked off by arbitrageurs).

To test these two predictions, we use data on triangular arbitrage opportunities for three currency pairs (dollar-euro, dollar-pound, and pound-euro).⁵ Our model and methodology should apply to any type of high frequency arbitrage opportunities. We focus on triangular arbitrage opportunities for several reasons. The first is practical. For our tests, we need to measure very accurately when an arbitrage begins, when it terminates, how it terminates (with a trade or a quote update), and to track prices after the arbitrage terminates (to identify arbitrage opportunities that are toxic; see below). This requires a level of precision for the data that is not easily available to researchers. Our data has the required granularity: it comprises all orders and trades from January 2003 to December 2004 in Reuters D-3000 (one of the two major platforms used by banks' currency trading desks at the time of our sample) with a time stamp accuracy of 10 milliseconds. Secondly, we can exploit a technological change in Reuters D-3000 as an instrument to test whether arbitrageurs' relative speed affects $PTAT^c$ and, through this channel, illiquidity (more on this below). Last, triangular arbitrage opportunities are very similar to other high frequency arbitrage opportunities: (i) they are frequent (we observe more than 37,000 in our sample), (ii) very short-lived (they last for less than one second on average), and (iii) deliver a small profit per opportunity (1 to 2 bps).⁶

As other arbitrage opportunities, triangular arbitrage opportunities arise for two reasons: (i) asynchronous price adjustments of different exchange rates to new information or (ii) price pressure effects in one exchange rate. By definition, price pressure effects are followed by reversals in the exchange rate initiating the arbitrage opportunity, whereas asynchronous price adjustments are eventually followed by permanent shifts in exchange rates.⁷ Thus, we use price patterns following the occurrence of an arbitrage opportunity to sort arbitrage opportunities in our sample into two groups: toxic (due to asynchronous price adjustments) and non-toxic

⁵One 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.

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

⁷As an illustration, suppose that market makers in the dollar-euro market 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 market makers in, say the dollar-pound market, are slow to adjust their quotes to reflect this information, a triangular arbitrage opportunity appears: one can indirectly buy dollars in euros at a price less than the current bid price in the dollar/euro market. When the arbitrage opportunity vanishes, exchange rates do not revert to their position before the arbitrage opportunity because the initial shock was a shock to fundamentals. Alternatively, if market makers in the euro-dollar market receives a string of buy orders of euros, they accumulate a large short position in euro and mark up the value of the euro against the dollar. This price pressure effect disappears as market makers start receiving sell orders of euros and reduce their inventory risk (see Ho and Stoll (1981), Grossman and Miller (1988), or Hendershott and Menkveld (2011) for models of price pressure effects in market maker markets).

(due to price pressure effects). Using a conservative classification scheme, we obtain 15,908 toxic arbitrage opportunities (about 32 per day and 41% of all arbitrage opportunities in the sample). Furthermore, the daily average value of $PTAT^c$ is equal to 0.74, which means that about 3/4 of all toxic arbitrage opportunities terminate with an arbitrageur's trade in our sample.

As arbitrageurs' speed and illiquidity are jointly determined in equilibrium, identifying the effect of arbitrageurs' speed on illiquidity is challenging. Our model however shows that exogenous shocks to the cost of speed for arbitrageurs can be used for identification. To this end, we use a technological shock. Until July 2003, traders had to manually submit their orders to Reuters D-3000. This was slowing down the speed at which arbitrageurs could exploit triangular arbitrage opportunities. In July 2003, Reuters introduced the "Autoquote API" service (API means "Application Programming Interface"). With this technology traders could directly feed their algorithms with Reuters D-3000 data and send orders automatically to this market. This technological change on Reuters D-3000 relaxed speed limits for arbitrageurs. Consistently, the introduction of "Autoquote API" is associated with shorter toxic triangular arbitrage opportunities (by 0.06 second, a 6.8% reduction in the average duration of these opportunities) and an increase of about 4% in $PTAT^c$, i.e., the likelihood that a toxic arbitrage opportunity terminates with a trade.

Using "Autoquote API" as an instrument for $PTAT^c$, we find that a 1% increase in dealers' exposure to arbitrage-induced picking off risk, $PTAT^c$, 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). We find a similar effect when we use effective spreads or the slope of limit order books (a measure of market depth) as measures of market illiquidity. Furthermore, as predicted, an increase in the fraction of toxic arbitrage opportunities in a given day is positively associated with illiquidity. Given the volume of trade in currency markets, the economic size of the effect of arbitrageur's speed ($PTAT^c$) on illiquidity is significant. In fact, using the volume of trade for the currency pairs in our sample, we estimate that a 1% increase in $PTAT^c$ increases the daily total cost of trading for the currency pairs in our sample by about \$161,000.

Hendershott et al. (2011) use the implementation of "Autoquote" on the NYSE as an instrument to study the effect of algorithmic trading and find a positive effect of algorithmic trading on market liquidity. Using a similar technological change, we reach a somewhat opposite conclusion. One possible reason is that the effect of relaxing speed limits in financial markets depends on whether this enables liquidity providers or arbitrageurs to react faster to new information.

In the former case, our model predicts that liquidity should improve (as found by Hendershott et al. (2011)) whereas in the second case, liquidity should decrease, as we find.

Chaboud et al. (2014) study the effect of algorithmic trading on price efficiency in the foreign exchange market. They find that algorithmic trading reduces the frequency of triangular arbitrage opportunities measured at a one second frequency. Consistent with this finding, our model predicts that an increase in arbitrageurs' speed should reduce the duration of arbitrage opportunities and therefore the frequency with which such opportunities are observed when data are sampled at fixed time intervals (e.g., every second). In contrast to Chaboud et al. (2014), our focus is on whether faster arbitrage increases adverse selection cost and we provide evidence that this is the case. Thus, both papers complement each other by considering different effects of high frequency arbitrage.

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 relationship: the effect of arbitrageurs' speed on liquidity. Kumar and Seppi (1994) emphasize the connection between cross-asset arbitrageurs and informed traders as we do in this paper. To our knowledge, our paper is first to test whether faster reaction of arbitrageurs to asynchronous price adjustments are a source of adverse selection and therefore illiquidity.⁸ More generally our paper is related to papers on speed in securities markets.⁹ Consistent with theory (e.g., Biais et al. (2011), Jovanovic and Menkveld (2012), or Foucault et al. (2012)), our empirical findings suggest that asymmetries in speeds among traders can be a cause of market illiquidity.

Arbitrage opportunities in the foreign exchange market (either violations of covered interest parity or triangular arbitrage) are well documented.¹⁰ 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, not because arbitrageurs are privately

⁸Roll 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 trades by arbitrageurs but do not specifically show that these trades explain the relation.

⁹See Hendershott and Moulton (2011), Garvey and Wu (2010), Hoffman (2012), or Pagnotta and Phillipon (2013)).

¹⁰See, for instance, Peel and Taylor (2002), Akram et al. (2008), Fong et al. (2008), Fenn et al. (2009), Mancini-Griffoli and Rinaldo (2011), Marshall et al. (2008), Kozhan and Tham (2012), Ito et al. (2013), or Chaboud et al. (2014))

informed but because they react faster to publicly available information (the occurrence of an arbitrage opportunity).

In the next section, we develop a simple model of cross-market arbitrage and show that illiquidity increases with the likelihood that a toxic arbitrage opportunity terminates with an arbitrageur’s trade (rather than a quote update). In Section 3, we describe our data, explain how we classify arbitrage opportunities in toxic and non-toxic arbitrage opportunities and test whether an increase in arbitrageurs’ speed has a positive effect on market illiquidity. We conclude in Section 4.

2 Hypotheses Development

2.1 Cross-Market Arbitrage, Speed, and Liquidity

In this section, we present the model that guides our empirical analysis.¹¹ 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 2.2.

Consider two risky assets X and Y that trade over an infinite number of trading rounds $t \in \{1, 2, 3, \dots, \infty\}$. Fundamental values of asset X and Y at the beginning of trading round t are denoted v_t^X and v_t^Y , respectively. They are linked by a deterministic relationship so that changes in v_t^X and v_t^Y are perfectly correlated. Specifically:

$$v_t^X = \sigma \times v_t^Y. \quad (1)$$

Hence, when the value of asset Y increases by 1, the value of asset X increases by σ .

There is one market maker in each asset and one arbitrageur. At the beginning of trading round t , the market maker in asset j posts an ask price, a_{tj} and a bid price b_{tj} for Q shares of asset $j \in \{X, Y\}$ with:

$$a_{tj} = v_{tj} + \frac{S_{tj}}{2}, \quad (2)$$

and

$$b_{tj} = v_{tj} - \frac{S_{tj}}{2}. \quad (3)$$

Thus, S_{tj} is the bid-ask spread of asset j in round t .

¹¹We build upon Foucault et al. (2003) and Foucault et al. (2013), who also consider models in which traders’ speed of reaction to market events are endogenous but without cross-market arbitrage.

In each trading round, two events can happen: (i) a permanent jump in v_{tY} , the fair value of asset Y , with probability α , or (ii) with probability $(1 - \alpha)$, the submission of a buy or a sell market order in asset j by one liquidity trader. In this case, the liquidity trader is equally likely to trade asset X or asset Y . Jumps in the value of asset Y are equal to $+1/2$ or $-1/2$ with equal probabilities.¹²

If a liquidity trader submits a market order, a trade takes place and the trading game proceeds to the next trading round. If there is an up (down) jump in v_{tY} then the fair value of asset X increases (decreases) by $\sigma/2$ (see eq.(1)). Thus, if $S_{tX} < \sigma$, there is an arbitrage opportunity. For instance, suppose that v_{tY} jumps upward. An arbitrageur can buy one share of asset X at price a_{tX} and sell σ shares of asset Y at price $b_{tY} = v_{tY} + 1/2 - S_{tY}/2$. With this trade, the arbitrageur locks in a profit equal to:

$$ArbProfit = \sigma \times b_{tY} - a_{tX} = (\sigma - S_{tX} - \sigma S_{tY})/2. \quad (4)$$

As the market maker in asset X sells an asset worth $v_t^X + \sigma/2$ at price a_{tX} , he loses $(\sigma - S_{tX})/2$.

Thus, the market maker in asset X loses money when he trades with an arbitrageur, when $S_{tX} < \sigma$ and the arbitrageur exploits a breakdown of the no-arbitrage link between X and Y , caused by a permanent jump in the value of asset Y (as assumed for the moment). We refer to these arbitrage opportunities as “*toxic*”. In the data, arbitrage opportunities can also stem from transient price pressures. These opportunities are non-toxic (they generate a profit for market makers). We incorporate them in the model in Section 2.2.

Hence, when the fair value of asset Y changes, the market maker in asset X and the arbitrageur race to react first to the jump. The market maker wants to update his quotes (to avoid a loss) before the arbitrageur exploits the arbitrage opportunity. We model this race as follows. When a jump in the value of asset Y occurs, it takes a time D_t^a for the arbitrageur to react to the jump and exploit the opportunity. This time is exponentially distributed with parameter γ_t . Similarly, it takes a time D_t^m for the market maker in asset X to react to the jump and update his quotes. This time is exponentially distributed with parameter λ_t . Hence, the average reaction times to a toxic arbitrage opportunity for the market maker and the arbitrageur are $E(D_t^m)^{-1} = \lambda_t$ and $E(D_t^a)^{-1} = \gamma_t$, respectively. Thus, the higher are λ_t and γ_t , the faster are

¹²In practice, such jumps are due to revisions in forecasts of asset cash-flows, due to news arrivals for instance. 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)).

the traders. We therefore refer to λ_t as the speed of the market maker and to γ_t the speed of arbitrageur, respectively.

Conditional on a jump in the value of asset Y , a trading round terminates when the arbitrage opportunity disappears, either because the arbitrageur exploits the opportunity or because the market maker in asset X updates his quotes. The likelihood that an arbitrage opportunity terminates with a trade from an arbitrageur, rather than a quote update, is:

$$p_t = \Pr(D_t^a < D_t^m) = \frac{\gamma_t}{\lambda_t + \gamma_t}, \quad (5)$$

because D_t^m and D_t^a are exponentially distributed. Naturally, p_t is higher when the speed ratio, $\frac{\gamma_t}{\lambda_t}$, is higher, that is, when the arbitrageur becomes relatively faster than the market maker.

In reality, traders call their response time to market information (e.g., a jump in the price of one asset) *latency* (see Moallemi and Saglam (2013) or Pagnotta and Phillipon (2013)). Minimizing latencies is costly as it requires technological investments (e.g., in hardware and software, datafeed, dedicated communication lines etc.) and attention (e.g., computer capacity).¹³ Hence, we assume that if the market maker operates at speed λ then it bears a cost $\Psi^m(\lambda) = \frac{c^m \lambda_t}{2}$ per trading round. Similarly, if the arbitrageur operates at speed γ then it bears a cost $\Psi^a(\gamma) = \frac{c^a \gamma_t}{2}$ per trading round. Differences in marginal costs of speed between the arbitrageur and the market maker might stem from differences in opportunity costs of paying more attention to the price of asset Y for each type of trading activity. Alternatively, these costs could be identical but, for the same investment in speed, market design may enable one type of trader to react more quickly to new information (jumps in the value of asset Y).¹⁴ For instance, suppose an investment of $\hat{\gamma}$ in speed for the arbitrageur produces an actual speed of only $\gamma = \kappa \hat{\gamma}$ with $\kappa < 1$. This is equivalent to $c^a > c^m$ and an increase in κ is formally equivalent to a reduction in $\frac{c^a}{c^m}$.¹⁵

Liquidity and Toxic Arbitrage. The speed ratio, $\frac{\gamma_t}{\lambda_t}$, is a determinant of market liquidity.

¹³Attention 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.

¹⁴For 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).

¹⁵This follows directly from the first order conditions that characterize traders’ optimal choices of speed. See (9) and (10).

Indeed, the expected profit of the market maker in asset X in a given trading round is:¹⁶

$$\Pi^m(S_{tX}; \lambda_t, \gamma_t) = -\frac{\alpha p_t}{2} Q(\sigma - S_{tX}) + \left(\frac{1-\alpha}{4}\right) S_{tX} Q - \frac{c_t^m \lambda_t}{2}. \quad (6)$$

The expected profit of market maker X decreases in p_t and therefore in the arbitrageur's speed relative to the market maker's speed, $\frac{\gamma_t}{\lambda_t}$. Hence, other things equal, to break-even, the market maker in asset X must charge a larger bid-ask spread when p_t is higher. Specifically, the competitive bid-ask spread, S_{tX}^c in asset X (that is, the spread such that $\Pi^m(S_{tX}^c; \lambda_t, \gamma_t) = 0$) is:

$$S_{tX}^c = \frac{2(\alpha p_t \sigma + c_t^m Q^{-1} \lambda_t)}{2\alpha p_t + (1-\alpha)}. \quad (7)$$

Thus, other things equal, S_{tX}^c increases in p_t and therefore γ . This yields our first testable implication.

Implication 1. *Consider a pair of assets X and Y linked by a no arbitrage relationship. If arbitrageurs react faster (γ_t is higher) to a permanent jump in the value of Y breaking down the no arbitrage relationship then the bid-ask spread in asset X increases.*

The effect of the market maker's speed (λ_t) on the bid-ask spread is less clear cut. On the one hand, a higher λ_t lowers the likelihood of being picked off for the market maker. On the other hand, it increases his cost of speed. The bid-ask spread decreases only when the first effect dominates, which depends on parameter values.

Consider the liquidity of asset Y now. By assumption, jumps in the value of asset Y always lead those in asset X . Thus, the market maker in asset Y is never exposed to toxic arbitrage trades and as a result the competitive bid-ask spread in asset Y is always zero. Hence, in the rest of this section, we focus on the case in which $S_{tY} = 0$. We discuss the case in which jumps can occur in both assets in Section 2.2.

Speed, Liquidity, and Arbitrage Duration. Testing Implication 1 is difficult because γ_t is not exogenous. In particular, an arbitrageur should devote fewer (cognitive or technological) resources to react fast to arbitrage opportunities in pair (X, Y) when (i) bid-ask spreads for this pair are large or (ii) market makers are very fast because, in either case, expected arbitrage profits are lower. To analyze these effects, we now consider the determination of their speeds by market maker X and Y for a given value of S_{tX} . The market maker in asset X chooses λ_t

¹⁶If a liquidity-motivated trade happens in trading round, it is equally likely to happen in asset Y or X . Hence, the market maker's expected profit on liquidity-motivated trades per round is $(\frac{1-\alpha}{2}) \frac{S_{tX}}{2} Q$, which is the second term in 6.

to maximize his expected profit in (6). The arbitrageur chooses γ_t to maximize her expected profit:

$$\Pi^a(S_{tX}; \lambda_t, \gamma_t) = \alpha p_t Q \frac{(\sigma - S_{tX})}{2} - \frac{c^a \gamma_t}{2}. \quad (8)$$

For a given bid-ask spread in asset X , equilibrium speeds, γ_t^* and λ_t^* are such that (i) γ_t^* maximizes (8) when $\lambda_t = \lambda_t^*$ and (ii) λ_t^* maximizes (6) when $\gamma_t = \gamma_t^*$. Thus, when $\lambda_t^* > 0$ or $\gamma_t^* > 0$, equilibrium speeds solve the following first order conditions:

$$\frac{\partial \Pi^m}{\partial \lambda_t} = \alpha Q \frac{(\sigma - S_{tX})}{2} \frac{\gamma_t^*}{(\lambda_t^* + \gamma_t^*)^2} - \frac{c^m}{2} = 0, \quad (9)$$

$$\frac{\partial \Pi^a}{\partial \gamma_t} = \alpha Q \frac{(\sigma - S_{tX})}{2} \frac{\lambda_t^*}{(\lambda_t^* + \gamma_t^*)^2} - \frac{c^a}{2} = 0. \quad (10)$$

Let $r = (\frac{c^m}{c^a})$ and suppose that $S_{tX} \leq \sigma$ (this will be the case in equilibrium; see below). Solving (9) and (10) for γ_t^* and λ_t^* , we obtain the unique equilibrium values of traders' speeds:¹⁷

$$\lambda_t^*(S_t; c^a, c^m) = \frac{\alpha Q (\sigma - S_{tX}) r}{c^m (1 + r)^2}, \quad (11)$$

$$\gamma_t^*(S_t; c^a, c^m) = \frac{\alpha Q (\sigma - S_{tX}) r}{c^a (1 + r)^2}. \quad (12)$$

Thus, as expected, the arbitrageur's optimal speed decreases in the bid-ask spread and her marginal cost of speed, c^a . It also decreases in the market maker's marginal cost of speed, c^m . Indeed, an increase in this cost implies that the market maker optimally chooses a lower speed, for a given bid-ask spread. This in turn enables the arbitrageur to reduce her speed because what matters for arbitrage profits is not being fast in the absolute but being fast *relative* to the market maker. Thus, the market maker's speed and the arbitrageur's speed are complements: the arbitrageur has an incentive to be slower if she expects the market maker to be slower and vice versa.¹⁸ In fact (11) and (12) imply that:

$$\gamma_t^*(S_t; c^a, c^m) = \lambda_t^*(S_t; c^a, c^m) r, \quad (13)$$

¹⁷Clearly, traders' expected profits are concave in their speed. Hence, solving the first order conditions is sufficient to obtain equilibrium speeds.

¹⁸This implies that investment in speeds are similar to an arms race in which each party reacts to another party's investment by investing more. This escalation eventually leaves unchanged each party's relative position (see equation (13)). Hence, there is over-investment in speed from the market maker and the arbitrageur's viewpoint: both would be better off committing to lower speeds than equilibrium speeds. This property of the model is in line with the way in which academics, regulators, and practitioners portray investment in latency in reality (see, for instance, "Stop the high frequency trader arms race," Financial Times, December 2012, by Larry Harris. Biais et al. (2011) study welfare consequences of this arms race. Budish et al. (2014) propose to replace continuous trading with batch auctions as a way to stop the arms race.

which means that, in equilibrium, the arbitrageur’s speed relative to the market maker’s speed (the speed ratio) is $r = c^m/c^a$. Thus, shifts in bid-ask spreads affect absolute speeds, leaving unchanged relative speeds, which are determined only by marginal costs of speeds.

As speed is costly, equilibrium speeds are never infinite, even though they can be very high when c^m and c^a are small. Thus, in equilibrium, arbitrage opportunities do not immediately vanish.¹⁹ Let D_t be the duration of an arbitrage opportunity (the “time-to-efficiency.”) In trading round t , the expected time-to-efficiency is:

$$E(D_t) = E(\text{Min}\{D_t^a, D_t^m\}) = \frac{1}{\lambda_t^*(S_{tX}) + \gamma_t^*(S_{tX})} = (\lambda_t^*(S_{tX})(1+r))^{-1}. \quad (14)$$

When the bid-ask spread in asset X increases, the arbitrageur chooses a lower speed of reaction to arbitrage opportunities because these are less profitable. In turn, the market maker chooses a lower speed as well because he is less exposed to being picked off by the arbitrageur. Hence, we have the following second testable implication.

Implication 2: *Consider a pair of assets X and Y linked by a no arbitrage relationship. An increase in bid-ask spread in, say, asset X results in a longer time-to-efficiency when an arbitrage opportunity arises in this pair.*

A few studies look at the effect of illiquidity on the duration of arbitrage opportunities. Deville and Riva (2007) find that deviations from put-call parity last longer for less liquid options, in line with Implication 2. Chordia et al. (2008) 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 and that return predictability has declined with tick size reductions in U.S. equity markets. This finding is also consistent with longer time-to-efficiency in less liquid markets.

Illiquidity and time-to-efficiency are determined simultaneously. Large bid-ask spreads lead to longer times-to-efficiency (Implication 2). Conversely, a faster arbitrageur (and therefore a lower time-to-efficiency) leads to a larger bid-ask spread (Implication 1). Thus, to test Implication 2, one needs a source of variation in bid-ask spreads that does not simultaneously affect traders’ speeds. Testing Implication 2 is not our goal, however. Instead, our main objective is to test Implication 3 below. However, looking at the duration of arbitrage opportunities helps to validate our interpretation of the empirical findings (see Section 3.5.2).

¹⁹Equilibrium speeds are also never zero. This is not necessarily the case when there are more than two arbitrageurs. See Section 2.2.

Equilibrium. In equilibrium, the competitive bid-ask spread and traders' speeds must be consistent with each other. That is, the competitive bid-ask spread, S_{tX}^c , λ^* , and γ^* must solve (7), (12) and (13). Figure I shows how the equilibrium is determined. The dashed line is the market maker's competitive spread as a function of the arbitrageur's speed (eq.(7)) and the plain line is the arbitrageur's speed as a function of the market maker's spread (eq. (12)), when the market maker's speed is such that her relative speed is equal to $1/r$ (as required by (13)). The equilibrium bid-ask spread and the arbitrageur's speed is at the intersection of the two lines. As the figure shows, when the arbitrageur's cost of speed (c^a) becomes smaller, the arbitrageur's speed becomes larger and, for this reason, the competitive bid-ask spread increases. This illustrates how shifts in traders' marginal costs of speed can be used to identify the effect of arbitrageurs' speed on illiquidity.

Insert Figure I about here.

To show this more formally, we solve for the competitive bid-ask spread as a function of the exogenous parameters. The speed ratio is equal to r in equilibrium (see (13)). Thus, using (5), we obtain:

$$p^* = \frac{\gamma_t^*}{\lambda_t^* + \gamma_t^*} = \frac{r}{1+r}. \quad (15)$$

In equilibrium, the likelihood that a toxic arbitrage terminates with an arbitrageur's trade is determined by the ratio of the market maker's cost of speed to the arbitrageur's cost of speed ($r = \frac{c^m}{c^a}$) and is increasing in this ratio.

Substituting λ_t^* and p^* by their expressions in (11) and (15) in (7), we can solve (7) for the equilibrium bid-ask spread as a function of r . We obtain:

$$S_{tX}^* = \frac{\alpha r(3+2r)\sigma}{\alpha r(3+2r) + (1-\alpha)(1+r)^2}. \quad (16)$$

The equilibrium competitive spread is less than σ , as conjectured previously. Furthermore it increases in r , that is, it gets larger when the market maker's cost of speed increases relative to the arbitrageur's cost of speed. The reason is that the arbitrageur's relative speed gets larger in equilibrium when r increases. Thus, a toxic arbitrage opportunity is more likely to terminate with a trade rather than with an update of his quotes update by the market maker.

We cannot directly observe the speed ratio, r . However, in equilibrium, p^* can serve as a

proxy for r since $r = \frac{p^*}{1-p^*}$ (see (15)). Hence, we can rewrite (16) as:

$$S_t^* = \frac{p_t^*(3 - p_t^*)\alpha\sigma_t}{\alpha p_t^*(3 - p_t^*) + (1 - \alpha)}, \quad (17)$$

This yields the following implication.

Implication 3: *Consider a pair of asset X and Y linked by a no arbitrage relationship. A reduction in arbitrageurs' cost of speed (c^a) relative to market makers' cost of speed (c^m) triggers an increase in p_t^* – the probability of an arbitrageur's trade, conditional on the occurrence of a toxic arbitrage – and for this reason an increase in bid-ask spreads.*

This implication is the main focus of our paper. To test it, we exploit variations in traders' costs of speed. According to the model, these variations should trigger a variation in p^* and, through this channel only, they should affect bid-ask spreads as (17) shows.

There are at least two sources of time-series variations in marginal costs of speed for a given pair (X, Y) . First, these costs are in part determined by the opportunity costs of attention to arbitrage opportunities. These are likely to vary over time. For instance, a human market maker's attention to current quotes will be reduced if he has to process other types of information (e.g., newswires, analysts reports etc.).²⁰ This suggests that opportunity costs of speed for human market makers are larger in days that are rich in attention-catching events. Increased trading activity should also reduce the amount of attention that a market maker can devote to quotes in other assets, including Y .²¹ For machines, the capacity allocated to track quotes in a given pair of assets is lost to seize opportunities in other pairs. Fluctuations in the profitability of these opportunities will therefore also trigger variations in c^a and c^m , and therefore r . Thus, for a given pair of assets X and Y , we expect to find a positive serial correlation (e.g., at the daily frequency) between bid-ask spreads in this pair and the likelihood that a toxic arbitrage opportunity terminates with a trade for this pair.

Second, changes in market structures can affect traders' speeds, holding their investment in speeds constant. For this reason, as explained previously, they are equivalent to variations in

²⁰Duffie (2010) also argues that it is sub-optimal for investors to continually focus attention on trading decisions. As an anecdotal evidence, he mentions the sharp drop in trading volume on the NYSE that occurred at the exact time Tigers Wood gave a televised speech to apologize for his marital infidelity. Schmidt (2013) shows that sport events reduce investors' attention to trading and firms' specific news.

²¹In line with these hypotheses, Chakrabarty and Moulton (2012) find that stocks assigned to NYSE specialists have lower liquidity when there is an earning announcement in the set of stocks handled by the specialists and show that a large part of this effect is due to attention constraints. Furthermore Coughenour and Corwin (2008) find that NYSE specialists allocate more attention to their most active stocks in periods of increased activity, resulting in less liquidity for their remaining assigned stocks.

$r = c^m/c^a$. In particular, greater automation of an arbitrageur’s reaction to arbitrage opportunities has the same effect as a decrease in c^a (and therefore an increase in p^*): it enables the arbitrageur to react faster to arbitrage opportunities, holding its investment in speed constant. Such technological changes offer a stronger way to test whether variations in arbitrageurs’ relative speed affect bid-ask spreads. Insofar as these changes only affect liquidity through traders’ speeds, they provide good instruments to identify the effect of arbitrageurs’ relative speed on liquidity. We use this approach to test Implication 3 (see Section 3.5).

Consider a technological change that reduces c^a and c^m but for which the effect on c^a is bigger. In this case, the bid-ask spread increases in equilibrium (Implication 3). The effect on time-to-efficiency is ambiguous, however. As c^a and c^m are lower, both the arbitrageur and the market maker react faster to arbitrage opportunities, which works to reduce the duration of arbitrage opportunities. However, the bid-ask spread is larger in equilibrium, which tends to increase the duration of arbitrage opportunities (Implication 2). Nevertheless, one can show that the first (direct) effect always dominates the second. Hence, in equilibrium, the average duration of arbitrage opportunities always decreases following a drop in marginal costs of speed for traders, whether this drop is relatively bigger for arbitrageurs or not. This is our next implication.

Implication 4: *Consider a decrease in both c^a and/or c^m . Whether this change reduces or increases r , it should trigger a reduction in time-to-efficiency (see the appendix for a formal proof).*

Chaboud et al. (2014) find that an increase in measures of algorithmic trading leads to fewer triangular arbitrage opportunities per second in the FX market. This is consistent with Implication 4. Indeed, an increase in speed, be it for market makers or arbitrageurs, reduces the duration of arbitrage opportunities and therefore the likelihood of observing an arbitrage opportunity when these opportunities are sampled at fixed interval of times.

2.2 Extensions

We now outline 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 (in particular Implication 3) still hold in each case.

Non-Toxic Arbitrage Opportunities. In the baseline model, jumps in the value of asset Y are permanent. In reality, as discussed in the introduction, arbitrage opportunities can also

arise when supply or demand shocks push the price of one asset away from its fundamental value temporarily.²² To account for this possibility, suppose that the price of asset Y at date t is:

$$price_{tY} = v_{tY} + U_t \times (1 - J_t) \times \delta_t, \quad (18)$$

with

$$v_{t+1Y} = v_{tY} + U_t \times J_t \times \epsilon_t, \quad (19)$$

where $U_t = 1$ with probability α , $U_t = 0$ with probability $(1 - \alpha)$, $J_t = 1$ with probability φ , and $J_t = 0$ with probability $(1 - \varphi)$. Furthermore, jumps ϵ_t and δ_t take the values $+1/2$ or $-1/2$ with equal probabilities. Thus, there is a jump (of size $1/2$) in the price of asset Y at date t with probability α and this jump is permanent (affects the fundamental value of Y) with probability φ . The frequency of permanent jumps in the value of asset Y is $\alpha\varphi$ whereas the frequency of transient jumps is $\alpha(1 - \varphi)$. The baseline case analyzed in the previous section (with only permanent jumps) corresponds to $\varphi = 1$.

A jump in asset Y yields the same profit for the arbitrageur, whether it is transient or permanent (and in fact the arbitrageur does not need to understand the nature of the jump to benefit from it). In contrast, if the market maker trades with the arbitrageur when the latter reacts to a transient jump in asset Y , he earns half the spread because, conditional on a transient jump in asset Y , the value of X is unchanged. Hence, from the market maker's viewpoint, arbitrage trades triggered by transient jumps in the value of Y are *non-toxic*. Thus, in this case, the market maker has no incentive to update his quotes to close the arbitrage opportunity and therefore the trading round terminates with an arbitrageur's trade with probability one.²³

The expected profit of the market maker and the arbitrageur are therefore:

$$\Pi^m(S_{tX}; \lambda_t, \gamma_t) = -\frac{\varphi\alpha p_t}{2} Q(\sigma - S_{tX}) + \left(\frac{1 - \alpha(2\varphi - 1)}{4} \right) S_{tX} Q - \frac{c_t^m \lambda_t}{2}. \quad (20)$$

and

$$\Pi^a(S_{tX}; \lambda_t, \gamma_t) = \alpha\varphi p_t Q \frac{(\sigma - S_{tX})}{2} + \alpha(1 - \varphi) Q \frac{(\sigma - S_{tX})}{2} - \frac{c^a \gamma_t}{2}. \quad (21)$$

²²In their survey, Gromb and Vayanos (2010) write: "*Limits of arbitrage are usually viewed as one of two building blocks needed to explain anomalies. The other building block are demand shocks experienced other than arbitrageurs. Anomalies are commonly interpreted as arising because demand shocks push prices away from fundamental values and arbitrageurs are unable to correct the discrepancies.*".

²³For simplicity, we assume that the market maker can perfectly distinguish fundamental jumps and non-fundamental jumps in asset Y . This is not key. For instance, the market maker could receive an imperfect signal about the type of the jump in asset Y .

The second term in (21) is the arbitrageur's expected profit when there is a transient jump in the price of asset Y . It does not depend on p_t because the market maker in asset X leaves his quotes unchanged in this case, whether or not he becomes aware of the arbitrage opportunity before the arbitrageur does. Correspondingly, the second term in (20) is the market maker's expected profit when she is counterpart to a non-toxic arbitrage trade. From the market maker's viewpoint, this trade is uninformed because it is uncorrelated with changes in v_{tX} . For this reason, the market maker's expected profit is greater when non-toxic arbitrage trades account for a larger fraction of all arbitrage opportunities (φ decreases).

Equilibrium speeds are identical to those obtained in the baseline model, except that α is replaced by $\alpha\varphi$ in (11) and (12). Holding the bid-ask spread in asset X constant, traders choose a lower speed when arbitrage opportunities are more likely to stem from transient jumps in asset Y (φ lower). The reason is that the market maker is then less exposed to toxic arbitrage trades and therefore has less incentives to invest in speed. Accordingly the arbitrageur can also cut on his investment in speed.

As in the baseline model, the likelihood that a toxic arbitrage opportunity terminates with a trade by the arbitrageur is $p^* = r/(1+r)$. Importantly, p^* is distinct from the unconditional likelihood, l^* that an arbitrage opportunity of any type terminates with an arbitrageur's trade, which is:

$$l^*(\varphi, r) = (1 - \varphi) + \varphi p^* = (1 - \varphi) + \varphi p^* = \frac{1 + r - \varphi}{(1 + r)} > p^*.$$

Unconditionally, the likelihood that an arbitrage opportunity terminates by a trade from an arbitrageur increases with the frequency of non-toxic arbitrage opportunities because the arbitrageur can always complete her trade when it is non-toxic. As shown below, this is consistent with the data.

Proceeding as in the baseline model, we can compute the equilibrium spread as a function of p^* and we obtain:

$$S_{tX}^* = \frac{\varphi p^* (3 - p^*) \alpha \sigma_t}{\alpha p^* (3 - p^*) + (1 - \alpha(2\varphi - 1))}. \quad (22)$$

The competitive bid-ask spread in asset X still increases with p^* and therefore our main implication (Implication 3) is unchanged. Furthermore, the bid-ask spread declines when a larger fraction of arbitrage opportunities are non-toxic (φ is smaller). This is another testable implication.

Implication 5. *Consider a pair of asset X and Y linked by a no arbitrage relationship. An*

increase in the fraction (φ) of toxic arbitrage opportunity due to jumps in the value of asset Y (resp, X) should trigger an increase in the bid-ask spread of asset X (resp. Y).

Equation (22) highlights the importance of using the likelihood that an arbitrage opportunity terminates with an arbitrageur's trade *conditional* on the opportunity being toxic, p^* , rather than the *unconditional* likelihood l^* for testing Implication 3. Indeed, an increase in l^* can be due to a decrease in φ and/or an increase in p^* . Thus, it is an imperfect proxy for p^* . Furthermore, the correlation between l^* and the bid-ask spread is ambiguous because an increase in l^* due to a decrease in φ is associated with a lower bid-ask spread whereas an increase in l^* due to an increase in p^* is associated with a larger bid-ask spread.²⁴

Jumps in asset X and Y . In the baseline model, asset Y leads asset X : jumps in the fundamental value of asset Y always precede adjustments in the value of asset X . Consider a more general formulation in which, in each trading round, there is a jump in the value of asset Y with probability $\alpha\theta_Y$ and a jump in the value of asset X with probability $\alpha(1 - \theta_Y)$ (in the baseline model, $\theta_Y = 1$). Market-makers in each asset are now 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_{tX} and S_{tY} (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.

One way to simplify the analysis is to assume that the arbitrageur engages in intertemporal arbitrage rather than cross-asset arbitrage.²⁵ For instance, if there is an increase in the fundamental value of Y , the arbitrageur buys asset X at price $a_{tX} = v_{tX} + \frac{S_X}{2}$ and then resell it at the market maker's bid price once the latter has adjusted his quotes, i.e., at $v_{tX} + \sigma/2 - \frac{S_X}{2}$.²⁶ In this way the arbitrageur locks in a profit equal to $(\sigma - 2S_X)/2$. Arbitrageurs in our data seem sometimes to follow this strategy (maybe because establishing the short and long legs of the arbitrage portfolio takes too much time or is too costly).²⁷ In this case, one can easily solve

²⁴The model implies that non-toxic arbitrage trades should improve liquidity because they help the market maker in asset X to recoup losses on toxic arbitrage trades. This effect is distinct from the positive effect of arbitrage on market liquidity highlighted by the literature on limits to arbitrage. In this literature, arbitrageurs improve liquidity by responding to demand shocks pushing prices away from fundamentals in the asset affected by demand shocks (Y in our model), not assets that arbitrageurs use to hedge their position (X in our model).

²⁵Gromb and Vayanos (2010) show that asset pricing implications of cross-asset and intertemporal arbitrage are very similar.

²⁶In contrast, in the case of a non-toxic arbitrage, the arbitrageur would intervene only in asset Y . For instance, if there is a positive non-fundamental jump in Y , the arbitrageur sell asset Y and wait for the price of asset Y to revert.

²⁷Shive and Schultz (2010) also make a similar observation.

for traders' equilibrium decisions (because the two markets can be treated separately). The implications of the model remain robust in this case as well.

Competition among arbitrageurs. The baseline model features a single arbitrageur. The case with $M > 1$ arbitrageurs is straightforward to analyze and delivers identical implications.²⁸ Not surprisingly, as M increases, the equilibrium bid-ask spread becomes larger because the likelihood that one arbitrageur reacts first to the arbitrage opportunity becomes higher. Moreover, if $r \leq \frac{M}{(M-1)}$ 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.²⁹ 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. Finally, 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. As a result, the number of prop trading firms who can either act as market makers or arbitrageurs becomes infinite ($M \rightarrow \infty$). At the beginning of each trading round, each firm decides whether to be a market maker (post quotes) or an arbitrageur. As the number of trading firms is very large (formally M is infinite), arbitrageurs' expected profit is zero in equilibrium.³⁰ Market-makers' expected profit is also zero when the bid-ask spread is competitive. Thus, all firms are indifferent between both roles and, in each trading round, one can be randomly selected to be a market maker.

3 Empirical analysis

3.1 Data

To test Implication 3, we use tick-by-tick data from Reuters trading system Dealing 3000 (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,

²⁸In this case, an arbitrageur who chooses a speed γ_i exploits the arbitrage opportunity with probability $\frac{\gamma_i}{\lambda + \sum_{j=1}^N \gamma_j}$.

²⁹This 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 .

³⁰When M goes to infinite, each arbitrageur's speed (γ^*) goes to zero but arbitrageurs' aggregate speed ($M\gamma^*$) remains strictly positive in equilibrium.

USD/GBP and EUR/GBP respectively). The sample period is from January 2, 2003 to December 30, 2004. 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.

The foreign exchange (FX) market operates around the clock, all year long. However, trading activity in this market considerably slows down during week ends and certain holidays. Hence, as is standard (see, for instance, Andersen et al. (2003)), we exclude the following days from our sample: week-ends, 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. Our final sample contains 498 days.

The FX market is a two-tier market.³¹ In the first tier, market makers from banks' trading desks ("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 which market makers trade together, to share risks associated with their clients' trades and take speculative positions. 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. Dealers can trade bilaterally (by calling each other), through voice brokers, or electronic broker systems (Reuters D-3000 and EBS). In the last decade, the market share of electronic broker systems has considerably increased and was already large at the time of our sample (see Pierron (2007) and King et al. (2012)).

Our tests use data from Reuters D-3000. This trading system is an electronic limit order book market, similar to that used in major world equity markets. On this system, market makers can post quotes (by submitting limit orders) or hit quotes posted by other market makers (by submitting market orders).³² The minimum order size is for one million in each currency pair. Our data contain all orders (limit and market) submitted to Reuters D-3000 over our sample period. They have several attractive features for our tests.

First, the three currency pairs are linked together by so called triangular arbitrage relationships (see Section 3.2.1). Deviations from these relationships are very easy to identify and,

³¹See Pierron (2007), Osler (2008), King and Rime (2010), and King et al. (2012) for excellent descriptions about participants, market structure and recent developments in foreign exchange markets and surveys of the market microstructure literature on these markets. King et al. (2012) note that the frontiers between the two-tiers of the FX market have been breaking down in recent years.

³²Dealers 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.

importantly, they are frequent (see Section 3.2.2) and do not vanish immediately. There is no risk, no short-selling constraints, and no minimal funding constraints involved in exploiting triangular arbitrage opportunities. Hence, standard frictions cannot explain why triangular arbitrage opportunities are not eliminated immediately. The most likely explanation is that, as in our model, traders lack full attention to the trading process and that technological constraints limits the speed at which traders can react to arbitrage opportunities. Last, triangular arbitrage opportunities are representative of arbitrage opportunities that high frequency traders seek to exploit in today's markets: they are very short-lived, small, and quasi riskless. This is exactly the type of opportunities that our model formalizes.

Second, our dataset 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 and a sell order. As each order has a unique identifier, we can track the order over its life. Thus, we can reconstruct the entire limit order book of each currency at any point in time. In this way, we can use the slope of the book as a measure of market illiquidity, in addition to standard measures such as bid-ask spreads. Furthermore, and more importantly for our purpose, we can identify whether an arbitrage opportunity terminates with a trade (the submission of a market order) or quote updates in limit order books for the three currencies in our sample. We can therefore accurately measure the frequency with which a toxic arbitrage opportunity terminates with a trade, that is, p^* , in the model. This is key for testing Implication 3.

Third, 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 to check whether the arbitrage opportunity is due to a permanent or a transitory jump in the value of the currency pair triggering the arbitrage opportunity. In this way, we can 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 as we exclusively focus on triangular arbitrage opportunities *within* Reuters D-3000. When an arbitrageur exploits a toxic

arbitrage opportunity on Reuters D-3000, it inflicts a loss on market makers with stale quotes in this system. 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 or 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 highly correlated (see Table V) because they are affected by common factors. We will therefore use liquidity measures from EBS for the currency pairs in our sample to control for omitted variables that create systematic time-series variations in liquidity in these pairs. Our EBS data, acquired from ICAP, consists of the 10 best quotes on each side of the limit order books for the three currency pairs in our sample, from January 2, 2003 to December 30, 2004.

Data from EBS are similar to those obtained for Reuters D-3000 with one important difference. The time-stamps of quotes, trades etc. is accurate only up to the second. For instance, all trades occurring in the same second during the trading day receive the same time stamp.³³ 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.³⁴ 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 classify triangular arbitrage opportunities into two groups: toxic and non-toxic opportunities. This classification is an important step for our tests.

3.2.1 Triangular arbitrage opportunities

Triangular arbitrage opportunities involve three currencies, say, i , j , and k . 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 . Now consider a trader who wants to buy one unit of currency j with currency i . He can do so directly by paying $A_t^{i/j}$ units of currency i . Alternatively, he can first buy $A_t^{k/j}$ units of currency k with

³³Representatives from EBS told us that they do not provide data at a more granular level.

³⁴For 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 p^* using EBS data.

currency i , at cost $A_t^{i/k} \times A_t^{k/j}$, and then buys one unit of currency j at $A_t^{k/j}$ in the market for currency k vs. j . The cost of this alternative strategy is $\hat{A}_t^{i/j} \equiv A_t^{i/k} \times A_t^{k/j}$. Similarly, a trader with one unit of currency j can obtain $B_t^{i/j}$ 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 in exchange of $B_t^{k/j}$ 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 prices for currency j in the i vs. j market.

A triangular arbitrage opportunity exists when

$$\hat{A}_t^{i/j} < B_t^{i/j} \quad \text{or,} \quad (23)$$

$$\hat{B}_t^{i/j} > A_t^{i/j}. \quad (24)$$

In the first case, one can secure a risk free profit equal to $\Pi_t^B = B_t^{i/j} - \hat{A}_t^{i/j}$ 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 $\hat{A}_t^{i/j}$ with two transactions in other currency pairs. In the second case, one can secure a risk free profit equal to $\Pi_t^A = \hat{B}_t^{i/j} - A_t^{i/j}$ 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 $\hat{B}_t^{i/j}$. These two arbitrage opportunities cannot occur simultaneously because (23) and (24) cannot both be true at the same time; see Kozhan and Tham (2012).

In reality, traders on Reuters D-3000 bear brokerage and membership fees paid to Reuters D-3000. These costs may vary across traders and some are fixed (e.g., the subscription fee to the Reuters D-3000 platform). To account for these costs, we say that a triangular arbitrage opportunity exists at time t in our data if and only if one of the following inequalities is satisfied:

$$\frac{B_t^{i/j} - \hat{A}_t^{i/j}}{\hat{A}_t^{i/j}} - z > 0, \quad (25)$$

$$\frac{\hat{B}_t^{i/j} - A_t^{i/j}}{\hat{B}_t^{i/j}} - z > 0, \quad (26)$$

where z is a measure of the cost of exploiting an arbitrage opportunity (expressed as a fraction of synthetic quotes) over and above the bid-ask spread. These costs are difficult to estimate and vary across traders. We set $z = 1$ basis point (bps), which is quite conservative.

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 (see Figure II). 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.

Insert Figure II about here.

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 (25) and (26) 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 (25) and (26).³⁵

³⁵The arrival of a market order cannot create an arbitrage opportunity. For instance, suppose that $A_t^{i/j} > \hat{B}_t^{i/j}$ and $\hat{A}_t^{i/j} > B_t^{i/j}$, so that, at time t , there is no arbitrage opportunity. 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.

2. If an arbitrage opportunity exists, we deduce that the limit order arrival created the arbitrage opportunity and we record the time t at which the order arrived as being the time at which an arbitrage opportunity begins.
3. We then record the first time $t + \tau$ at which a trade or the update of limit orders yield new quotes such that the triangular arbitrage opportunity disappears. Hence, τ is the duration of the arbitrage opportunity.

3.2.2 Classifying arbitrage opportunities.

As previously explained, triangular arbitrage opportunities can arise for two reasons: (i) asynchronous price adjustments to the arrival of fundamental information (a permanent jump in the fundamental value of one currency pair followed, with a delay, by an adjustment in the rates of other currency pairs) and (ii) price pressure effects (a transitory deviation of the rate for one currency pair from its fundamental value). The first type of arbitrage opportunity is toxic (it exposes slow market makers to the risk of a loss if they trade with an arbitrageur) whereas the second type is non-toxic (and is even a source of profits for market makers who are counterparties to arbitrageurs' trades; see Section 2.2). For our tests, we must measure the frequency with which *toxic* arbitrage opportunities terminate with a trade from an arbitrageur. Thus, we need to classify all triangular arbitrage opportunities into two subgroups: toxic and non-toxic.

We proceed as follows. Consider again the example described after Table I. If the depreciation of the euro against the dollar (asset Y in the model) is permanent (e.g., due to macro-economic news), the rates for other currencies (asset X in the model) must permanently adjust at levels where the arbitrage opportunity disappears. Alternatively, the depreciation of the euro against the dollar might just be transitory. For instance, to quickly reduce a long position in euros, market makers in the $\$/\text{€}$ might be willing to sell euros at a bargain. To do so, they can post a low ask price, $A^{\$/\text{€}}$, relative to $\hat{B}^{\$/\text{€}}$ (which reflects the fundamental value of the $\$/\text{€}$).³⁶ In this case, the rates of other currencies should not adjust and the $\$/\text{€}$ ask

³⁶This behavior is indeed consistent with models of optimal inventory management for dealers (e.g., Amihud and Mendelson (1980)). Bjønnes and Rime (2005) find that currency dealers in their sample often use limit orders on EBS and Reuters to unwind their inventories (even though their quotes across all venues seem relatively insensitive to inventories). One may wonder why market makers do not unload their risk by trading at the synthetic quote when $A^{\$/\text{€}} < \hat{B}^{\$/\text{€}}$. For instance, market makers with a long position in euros would be willing to sell at $\hat{B}^{\$/\text{€}}$ if they are willing to sell at a lower price, $A^{\$/\text{€}}$. One possible reason is that offering a good quote for a large quantity is a more efficient and quicker way to unwind a risky position than hitting other quotes. In line with this conjecture, we find that, on average, the limit order initiating a non-toxic arbitrage has a larger size than the amount posted at indirect quotes (2.412 million vs. 1.765 million of the base currency).

price will revert to its pre arbitrage level after market makers in the \$/€ market have unwound a sufficiently large fraction of their euro position.

Thus, 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.³⁷ If instead this price change persists after the arbitrage opportunity terminates, we consider that the arbitrage opportunity is due to an asynchronous price adjustments in the rates of the three currencies. More specifically, for each triangular arbitrage opportunity in our sample, we compare the exchange rate for the three currency pairs just *before* the onset of the arbitrage opportunity (time t) and just *after* the end (time $t + \tau$) of this opportunity. 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 constitute our sample of toxic arbitrage opportunities.

Insert Figure III about here.

Figure III illustrates the 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. Thus, we classify these arbitrage opportunities as toxic. In contrast, in Panel B, only the quotes of the currency pair at the origin of the arbitrage opportunity change during the arbitrage opportunity. Moreover, they 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.

Using this methodology, we identify 37,689 triangular arbitrage opportunities in our sample, of which 15,908 are classified as toxic. Per day, there are on average 32 (standard deviation:

³⁷Shive 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).

20.83) toxic triangular arbitrage opportunities and 45 non-toxic arbitrage opportunities. On average, the daily fraction of toxic arbitrage opportunities is $\hat{\varphi} = 41.5\%$ (s.d.= 10%).

Insert Figure V about here

Panel A of Figure V 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 being very hot (e.g., in May or June 2003) and some days being quieter. The numbers of toxic and non-toxic arbitrage opportunities move in locksteps, which suggests that periods of intense activity in the market are also periods with frequent information arrivals.

Panel B of Figure V shows 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 toxic and non-toxic 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, in defining our variables and in our tests, we only use observations from 7:00 to 17:00 GMT in each day.

Intuitively, market makers in one currency pair are less exposed to a toxic arbitrage opportunity if their quotes lead other currency pairs in terms of price discovery. In other words, market makers that tend to adjust their quotes first in response to a change in fundamentals (market makers in asset Y in the baseline version of the model) are less exposed than other market makers. In our sample, we expect the EUR/USD pair to lead others in terms of price discovery because this is by far the most traded currency pair worldwide (see Bank of International Settlements (2010)). Table II supports this conjecture. It shows that permanent changes in EUR/USD rates are at the origin of about 51% toxic arbitrage opportunities in our sample, vs., 28.68% and 20.32% for GBP/USD and EUR/GBP, respectively.

Insert Table II about here

To exploit arbitrage opportunities, arbitrageurs should use buy market orders in one currency pair and sell market orders in others. For instance, consider the toxic arbitrage opportunity featured on the left hand side of Panel A in Figure III. This arbitrage is due to a permanent shift of the quotes in the EUR/GBP pair. To exploit this opportunity, arbitrageurs should place

buy market orders in the two other pairs (EUR/USD and GBP/USD) and sell market orders in the EUR/GBP pair. This behavior should generate positive order imbalances (differences between buys and sales with market orders) in the former pairs and negative order imbalances in the latter over the duration of the arbitrage opportunity.

If arbitrageurs just follow textbook cross-market arbitrage strategies, these imbalances should on average be of equal size in absolute value (long positions should be equal to short positions). In reality, however, we expect systematic differences in the size of these imbalances between initiating and non-initiating currency pairs in toxic and non-toxic arbitrage opportunities. Indeed, as explained in Section 2.2, some arbitrageurs are likely to engage in inter-temporal arbitrage rather than cross-market arbitrage, both to save time and to save on transaction costs. In toxic arbitrage trades, these arbitrageurs should then only trade in non-initiating currencies (and make a profit when their rates eventually adjust). In contrast, in non-toxic arbitrage trades, they should only trade in the initiating currency (and make a profit when its rate reverts). Thus, if our classification of arbitrage opportunities is correct (and if some arbitrageurs engage in intertemporal arbitrage), we should observe, right after the onset of an arbitrage opportunity, larger order imbalances (in absolute value) in a given currency pair when this pair initiates a non-toxic arbitrage opportunity than when it initiates a toxic arbitrage trade.³⁸ Furthermore, we should observe the exact opposite pattern in currency pairs that do not initiate the arbitrage.

Table III shows that this is the case, lending thereby support for our classification. For instance, consider the EUR/USD pair. When this pair “initiates” an arbitrage opportunity (i.e., a change in quotes for this pair triggers the occurrence of an opportunity), the absolute difference between purchases and sales (measured in million of the base currency) in this pair, one second after the beginning of the opportunity, is significantly higher (by 0.18 million of euros with a t-stat of 8.78) when the arbitrage opportunity is *non-toxic*. In contrast, in this case, the absolute difference between purchases and sales in non-initiating pairs, is significantly lower (e.g., -0.06 with a t-stat of -7.02) than when the arbitrage opportunity is toxic.

Insert Table III about here

³⁸The sign of the order imbalance in a given currency pair around triangular arbitrage can be either positive or negative depending on whether arbitrageurs should buy or sell this pair. This is the reason why we use absolute order imbalances.

3.3 Variables of interest

In this section, we describe the key variables used in our empirical analysis. Our main prediction is that an increase in arbitrageurs' relative speed of reaction to toxic arbitrage opportunities makes the market less liquid. As explained in Section 2.1, we can test this prediction by studying the effect of p_t^* , the likelihood that a toxic arbitrage opportunity terminates with a toxic trade, on illiquidity. As a proxy for this conditional probability, we use the daily fraction of toxic arbitrage opportunities that terminate with a trade (the submission of market orders). We denote this fraction by $PTAT_t^c$ (conditional probability of a toxic arbitrage trade). By definition, in day t :

$$PTAT_t^c = \frac{\text{No. of toxic arbitrage opport. that terminate with a trade on day } t}{\text{No. of toxic arbitrage opportunities in day } t}. \quad (27)$$

Similarly, we also compute the frequency with which non-toxic arbitrage opportunities terminate with a trade on each day. We denote this frequency by $PNTAT_t^c$. That is:

$$PNTAT_t^c = \frac{\text{No. of non-toxic arbitrage opport. that terminate with a trade on day } t}{\text{No. of non-toxic arbitrage opportunities in day } t}. \quad (28)$$

Illiquidity should increase with (i) the likelihood that an arbitrage opportunity is toxic, φ , and (ii) the magnitude of price jumps, σ , when an arbitrage opportunity arises (see (22)). Thus, in our tests, we use daily proxies for φ and σ as controls. As a proxy for φ on day t , we use the fraction of all arbitrage opportunities that are toxic on this day:

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

To obtain a proxy for σ , we proceed as follows. For a given currency pair (say i/j), we define two different midquotes: $f_{i/j} = \frac{A^{i/j} + B^{i/j}}{2}$ and $\hat{f}_{i/j} = \frac{\hat{A}^{i/j} + \hat{B}^{i/j}}{2}$. The first mid-quote is based on actual quotes while the second is based on synthetic quotes. We then use the average *percentage* difference between these two mid-quotes *at the time of arbitrage opportunities* on day t , for all currency pairs, as proxy for σ , that is, the size of arbitrage opportunities on day t when they occur. That is, our proxy for σ on day t is:

$$\hat{\sigma}_t = \frac{|f_{i/j} - \hat{f}_{i/j}|}{\hat{f}_{i/j}}.$$

The realizations of this proxy always exceed the quoted spread by at least one basis point

because we only retain arbitrage opportunities such that (26) and (25) are satisfied.³⁹

Variables $PTAT^c$, $\hat{\varphi}_t$, $\hat{\sigma}_t$ are not currency specific because a triangular arbitrage opportunity involves the three currency pairs in our sample. In addition, we also use currency-specific controls known to be correlated with measures of market illiquidity: the average daily trade size in each currency (denoted $trsize_{it}$ in currency i on day t); the daily realized volatility, i.e., the sum of squared five minutes mid-quote returns in each currency (denoted vol_{it}); and the daily number of orders (entry of new limit and market orders as well as limit order updates) denoted $nrorders_{it}$. This variable measures the level of activity on Reuters D-3000 on each day.

We use various standard proxies for market illiquidity in each currency pair i : (i) average daily *percentage* quoted bid-ask spreads ($spread_{i,t}$ for currency pair i), that is, the absolute quoted spread divided by the mid-quote; (ii) average daily effective spreads ($espread_{i,t}$), i.e., twice the average absolute difference between each transaction price and the mid-quote at the time of the transaction; and (iii) average daily slopes of the limit order book ($slope_{it}$). In limit order markets, a buy (sell) market order can walk up (down) the limit order book if the quantity available at the best ask (bid) price is insufficient for full execution of the order. Hence, buy (sell) market orders execute at a larger markup (discount) when quantities offered at each price in the limit order book are smaller; that is, if the elasticity of supply and demand schedules implied by limit orders is lower. The variable $slope_{it}$ is a measure of this elasticity. Specifically, 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 divided by 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, for each type of arbitrage opportunity, we also measure the average time it takes for the arbitrage opportunity to disappear as a proxy for $E(D_t)$ in the model. We denote this time by TTE_t (Time-To-Efficiency) and by TTE_t^{nontox} for toxic and non toxic arbitrage opportunities, respectively, on day t . We do not use these durations as controls in our tests because, as explained previously, they are jointly determined with bid-ask spreads. However, in Section 3.5.2, we check whether the technological shock on traders speed in our sample (see

³⁹Actually it is easily shown that (26) and (25) imply that $\frac{f_{i/j} - \hat{f}_{i/j}}{\hat{f}_{i/j}} > z + (s_{ij} + \hat{s}_{ij})/2$ where s_{ij} and \hat{s}_{ij} are the percentage bid-ask spreads for the currency pair i/j and the synthetic pair i/j .

below) triggered a reduction in time-to-efficiency as predicted by Implication 4.

3.4 Summary statistics

Table IV 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.

Insert Table IV here

The average size of a toxic arbitrage opportunity, $\hat{\sigma}$, is 3.535 bps with a standard deviation of 0.757. The average daily arbitrage profit (expressed in percentage term, as in (25) and (26)) after accounting for both implicit and explicit transaction costs (we set $z = 1$ bps for explicit transaction costs) on a toxic arbitrage opportunity is 1.427 bps with a standard deviation of 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.⁴⁰

The likelihood that a toxic arbitrage opportunity terminates with a trade ($PTAT^c$) is 74.1% on average with a standard deviation of 0.110. The likelihood that a non-toxic arbitrage opportunity ($PNTAT^c$) terminates with a trade is higher on average (0.807), as predicted by the model (see Section 2.2).

Panel C of Table IV 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 (from about 1 bps or less to at most 5 bps). All three illiquidity measures indicate that the most illiquid currency pair is GBP/USD. For instance, on Reuters, the average effective spread for this pair is 2.073 bps vs. 0.96 bps for EUR/GBP, which the most

⁴⁰Another 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).

liquid pair on Reuters. Furthermore, 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 bps on average vs. 2.5 bps for EBS. In contrast, EBS is more liquid for the EUR/USD. 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 V about here

Table V reports the unconditional correlation of the variables used in our tests. The duration of toxic arbitrage opportunities (TTE) and the various measures of market illiquidity are positively correlated. That is, toxic arbitrage opportunities last longer on average in days in which the market for the three currency pairs is more illiquid. This is consistent with our Implication 2. 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.

In general, the correlation between $PTAT^c$ (our proxy for p_t^*) and measures of illiquidity is positive (as predicted by Implication 1) but not significantly different from zero. As previously explained, this may stem from the fact that $PTAT^c$ 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).

As implied by (22), measures of illiquidity for the three currency pairs in our sample are positively and significantly correlated with the fraction of toxic arbitrage opportunities ($\hat{\varphi}$) and the size of arbitrage opportunities ($\hat{\sigma}$).

The correlation between $PNTAT^c$ 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 (as emphasized in the literature on limits to arbitrage). Furthermore, market makers in non-initiating currencies earn the bid-ask spread on trades with arbitrageurs in non-toxic arbitrage (as, from their standpoint, these trades are uninformed).

Lastly, the correlation between $PNTAT^c$ and $PTAT^c$ is slightly positive and statistically significant. This low correlation indicates that variations in $PNTAT^c$ and $PTAT^c$ contain different information and are not driven by the same factors. This is consistent with the idea that market makers in non-initiating pairs behave differently in toxic and non-toxic arbitrage opportunities. They should actively try to 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.

3.5 Tests

3.5.1 Are faster arbitrageurs a source of illiquidity?

Our main goal is to test whether an increase in arbitrageurs' relative speed is a source of market illiquidity because it raises market makers' exposure to the risk of being picked off when a toxic arbitrage trade occurs (Implication 3). As explained in Section 2.1, we can test whether this channel plays a role by studying the effect of p^* (proxied by $PTAT_t^c$) on measures of illiquidity. According to Implication 3, this effect should be positive if fast arbitrageurs expose market makers to adverse selection. Another implication of the model is that an increase in the fraction of toxic arbitrage opportunities should have a positive effect on illiquidity (Implication 5).

To test these two implications, we first estimate the following OLS panel regression:

$$illiq_{it} = \omega_i + \xi_t + b_1 PTAT_t^c + b_2 vol_{it} + b_3 \hat{\varphi}_t + b_4 \hat{\sigma}_t + b_5 trsize_{it} + b_6 nrorders_{it} + b_7 illiq_{it}^{EBS} + \varepsilon_{it}, \quad (29)$$

where $illiq_{it}$ is one of our three proxies for illiquidity and ω_i and ξ_t are, respectively, currency and time fixed-effects (dummies for each month in our sample). For each illiquidity measure, we also use its EBS counterpart ($illiq_{it}^{EBS}$) as control. In this way, we 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). Other controls are the daily realized volatility (5 minutes aggregation vol_{it}); the daily average trade size ($trsize_{it}$); the daily number of orders ($nrorders_{it}$); the daily average size of toxic arbitrage opportunities ($\hat{\sigma}_t$); and the fraction of arbitrage opportunities that are toxic ($\hat{\varphi}_t$). We estimate equation (29) using standard errors robust to heteroscedasticity and time series autocorrelation. Panel A of Table VI reports the results.

Consistent with Implication 5, there is a significant and positive association between the

fraction of toxic arbitrage opportunities and all measures of illiquidity. Moreover, an increase in the size of arbitrage opportunities is positively associated with illiquidity, as (22) implies. Finally, the relationship between $PTAT_t^c$ and illiquidity is positive as well. However, it significant (at the 10% level) only when illiquidity is measured with the effective spread. As explained previously, omitted variables might affect both time-series variations in our illiquidity measures and arbitrageurs' relative speed of reaction to toxic arbitrage opportunities (proxied by $PTAT_t^c$). In presence of such variables, the error term in (29) and $PTAT_t^c$ are correlated. This correlation can bias the estimate of the effect of $PTAT_t^c$ on illiquidity (i.e., coefficient b_1).

Insert Table VI about here

To address this issue, we use a technological shock that affects the speed at which traders can react to arbitrage opportunities on Reuters D-3000, without directly affecting the bid-ask spread. In July 2003, Reuters D-3000 introduced a new service, named "Autoquote (API)", enabling traders to (a) consolidate limit order book data for multiple currencies, (b) input these data automatically in their pricing engines, and (c) automate their order entry. "API" stands for "Application Programming Interface." As noted by Pierron (2007), APIs allow traders to branch directly their algorithms onto trading systems.⁴¹ He further notes: "*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 pool of liquidity in the FX market.*"⁴² Initially, Reuters provided this service only to a limited number of clients because of capacity constraints as users of APIs consume more bandwidth than manual users. This does not in itself invalidate our tests because our implications hold even if there is only one fast arbitrageur (as in the baseline model). However, a low usage of Autoquote makes it more difficult for us to detect its effect.

Insert Figure IV about here.

To test whether Autoquote did in fact trigger a significant increase in automated trading activity, we use the ratio of the daily number of orders in an asset to the daily number of trades

⁴¹Hendershott et al. (2011) use the introduction of AutoQuote on the NYSE, which also took place in 2003, as an instrument for the level of algorithmic trading activity in NYSE stocks. They write: "*Autoquote was an important innovation for algorithmic traders because an automated quote update could provide more immediate feedback about the potential terms of trade. This speedup of a few seconds would provide critical new information to algorithms*" (Hendershott et al. (2011), p.13).

⁴²In minutes of the Canadian Foreign Exchange Committee, Jack Linker (an executive from Reuters) noted that Reuters Autoquote (API) application was quickly adopted by traders and partly responsible for the growth in trading volume on Reuters D-3000 trading platform; see CFEC (2006). The electronic broker EBS launched a similar service, "EBS Spot Ai", in 2004; see King and Rime (2010) and ?.

in this asset as a proxy for algorithmic trading activity. As explained by Hendershott et al. (2011), a higher value of the order-to-trade ratio is associated with more automated trading. Figure IV presents a time series of this ratio for the three currencies in our sample. Dashed lines 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 in July 2003. Hence, the introduction of Autoquote on Reuters D-3000 coincides with a clear increase in automated trading on this platform.

Autoquote enables traders to react faster to changes in limit order books and therefore arbitrage opportunities. In principle, it might either increase or reduce the speed ratio, $r_t = \frac{c^m}{c^a}$. In the first case, arbitrageurs should become relatively slower ($PTAT^c$ should decrease) whereas in the second case they should become relatively faster ($PTAT^c$ should increase). *In either case*, the introduction of Autoquote on Reuters D-3000 can be used as an instrument to identify the effect of arbitrageurs' relative speed ($PTAT_t^c$) on illiquidity because it is a source of variation for this speed that should not directly affect bid-ask spreads (as argued by Hendershott et al.(2011)).

Accordingly, we re-estimate our baseline regression (29) using an IV approach. In the first-stage, we regress our proxy for arbitrageurs' relative speed, $PTAT_t^c$, on a dummy variable AD_t , which is equal to one after the introduction of Autoquote on Reuters D-3000 (July 2003) and control variables used in (29):

$$PTAT_t^c = \omega_i + \xi_t + a_1 AD_t + a_2 vol_{it} + a_3 \hat{\varphi} + a_4 \hat{\sigma}_t + a_5 trsize_{it} + a_6 nrorders_{it} + a_7 illiq_{it}^{EBS} + u_{it}. \quad (30)$$

The estimates of the first and second stages of the IV regressions are reported in Panel B of Table VI.⁴³ The first stage of the IV regression shows that the introduction of Autoquote on Reuters D-3000 had a significant positive effect on the likelihood that an arbitrage opportunity terminates with a trade rather than a quote update. The F statistics is around 16 and rejects the null that our instrument has no explanatory power.⁴⁴ Overall, the first stage of the IV supports our conjecture that Autoquote increased arbitrageurs' relative speed of reaction to toxic arbitrage opportunities. The coefficient 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 increases by about 4% after July 2003.

⁴³The 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.

⁴⁴Bound et al. (1995) (p.446) mention that "F statistics close to 1 should be cause for concern".

Turning to the second stage of the IV, we find that the effect of $PTAT_t^c$ on illiquidity, b_1 , is now positive and statistically significant at the 1% level for all measures of illiquidity. The magnitude of the effect is always much higher than that measured by the OLS regression. For instance, a 1% increase in $PTAT_t^c$ raises the quoted bid-ask spread by 0.079 bps according to the IV estimates whereas the effect estimated with OLS is not significantly different from zero.

The effect of $PTAT_t^c$ on illiquidity might look small in economic terms. However, bid-ask spreads for currencies in our sample are very tight: 0.079 bps represents 4% of the average bid-ask spread for the currencies in our sample (about 2bps). Hence the effect of Autoquote on $PTAT_t^c$ (4% on average) raised bid-ask spreads by about 16% for currency pairs in our sample. The NYSE “Hybrid Market” reform, which also triggered an increase in execution speed for NYSE market orders, had an effect of a similar magnitude on bid-ask spreads (see Hendershott and Moulton (2011)).

Furthermore, the average trade size in the markets 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), \$1.831 million in EUR/GBP (dollar value of €1.548), or \$1.959 million on average in all three markets (see Table IV). Thus, an increase of 0.07934 bps in the quoted spread raises trading costs by at least \$16 on the average trade (about \$2 million). By way of comparison, Lyons (1995) estimates that the market maker studied in his paper (active in the DM/USD pair) requires, for a \$10 million trade, a compensation of about 0.28 bps for adverse selection costs. This amounts to \$56 on a two million trade.

Lastly, on average, in our sample, 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 adverse selection in toxic arbitrage opportunities is at least $7.934 \times (\$2.390 \times 4,692 + \$1.655 \times 2,365 + \$1.831 \times 2,841) = \$88,971 + \$31,054 + \$41,272 = \$161,297$ for the three markets in total or about \$40 million per year.⁴⁵

For all illiquidity measures, we again find a positive and statistically significant relation between the fraction of toxic arbitrage opportunities ($\hat{\varphi}$) and illiquidity. Furthermore, 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.

As a robustness check, we have also run the tests using hourly estimates of each variable

⁴⁵As 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.

in our regressions rather than daily estimates. Qualitatively, the findings are very similar and, in economic terms, they are stronger. For instance, an increase by 1% in $PTAT_t^c$ triggers an increase of 0.08 bps for the effective spread when we estimate (29) at the hourly frequency with the IV approach vs. 0.034 bps at the daily frequency. For brevity, we do not report the estimates obtained with tests at the hourly frequency. They are available upon request.

3.5.2 Time-to-efficiency

On average, the increase in $PTAT_t^c$ associated with the introduction of Autoquote triggers an increase in illiquidity for the currencies in our sample. This is consistent with Implication 3 of our model. As this increase is due to an increase in arbitrageurs' relative speed, the model also implies that, despite the increase in illiquidity, we should also observe a reduction in the average duration of arbitrage opportunities (Implication 4).

We test this prediction by estimating the effect of Autoquote on the average duration of arbitrage opportunities. Specifically, we estimate the following regression:

$$\log(TTE_t) = \omega_i + \xi_t + a_1 AD_t + a_2 vol_{it} + a_3 \hat{\varphi}_t + a_4 \hat{\sigma}_t + a_5 trsize_{it} + a_6 nrorders_{it} + u_{it}, \quad (31)$$

where TTE_t is the average duration of arbitrage opportunities. Estimates are reported in Table VII. In Column 1 we only consider toxic arbitrage opportunities whereas in Column 2 we consider all arbitrage opportunities. Autoquote is associated with a decrease in the duration of toxic arbitrage opportunities by about 6.8%.⁴⁶ This result is robust to inclusion or exclusion of control variables. Similar estimates are obtained when we use the average duration of all arbitrage opportunities. Thus, the introduction of Autoquote on Reuters simultaneously triggers an increase in illiquidity *and* pricing efficiency. At first glance, this seems paradoxical and inconsistent with Implication 2. Yet, this is a possibility because illiquidity and price efficiency are simultaneously determined. In particular, a reduction in arbitrageurs' cost of speed can simultaneously increase illiquidity (Implication 3) and improve price efficiency (Implication 4).

Interestingly, we also find that, on average, the duration of arbitrage opportunities is shorter when the fraction of toxic arbitrage opportunities ($\hat{\varphi}$) increases. This finding is also consistent with the model because an increase in $\hat{\varphi}$ should lead both arbitrageurs and market makers to

⁴⁶As 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 opportunities.

react faster to arbitrage opportunities according to the model (see the discussion in Section 2.2).

3.5.3 Heterogeneous effects across currencies

As observed in Table II, toxic arbitrage opportunities start most frequently with an update in the EUR/USD pair. This means that this pair frequently plays the role of asset Y in the baseline model whereas 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 fraction of toxic arbitrage opportunities ($\hat{\varphi}$) than the other pairs and (ii) arbitrageurs' relative speed of reaction ($PTAT^c$).

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 VIII present the results for each currency pair in our sample. Realizations of control variables in the first stage and the second stage are specific to each currency pair (e.g., the average trade size in a given day is specific to each currency). Yet, consistent with our earlier results, we find that Autoquote is associated with a significant increase in $PTAT^c$, which is roughly the same across all currencies, as one would expect.

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 $PTAT^c$ 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 $PTAT^c$ on other measures of illiquidity seems sometimes stronger than in other currency pairs.

Insert Tables VIII about here

3.6 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 measure of market makers' exposure to arbitrageurs' picking off risk ($PTAT^c$) 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 $PTAT^c$ proxies for informational asymmetries associated with macro economic announcements, we would expect the effect of $PTAT^c$ on illiquidity to be weaker when we control for surprises in macro-economic announcements. To test whether this is the case, we use data from Money Market Survey (MMS), provided by InformaGM, to construct macroeconomic announcement surprises in the different geographical areas (EMU, U.K., and U.S.) relevant for our currency pairs.

The MMS data provide median forecasts of all macro-economic announcements by market participants (collected on the Thursday prior to the announcement week) and their actual realization on the day of the announcement. Announcement surprises are measured as the realized announced value minus the median forecast. Following Andersen et al. (2003), we standardize announcements surprises by their standard deviation. Specifically, the surprise $N_{k\tau}$ of announcement type k (e.g., 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).⁴⁷ 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

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

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 .⁴⁸ We find a significant a positive correlation among the VPIN measures for the three currency pairs in our sample. In contrast, correlation between $PTAT^c$ 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 $PTAT^c$ and $VPIN$ do not capture the same exposure to adverse selection for market makers.

Table IX 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).⁴⁹ We also find a positive and marginally significant relationship between VPIN and quoted or effective bid-ask spreads. However, and more importantly, the effect of $PTAT^c$ on illiquidity remains positive and significant for all measures of illiquidity. Furthermore, estimates of the effect of these variables are very similar to those reported in Table VI. Hence, $PTAT^c$ contain information about market makers’ exposure to adverse selection,

⁴⁷VPIN is an alternative to the PIN measure that has been extensively used in finance; see Easley et al. (2012).

⁴⁸Computation 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).

⁴⁹Informational asymmetries created by surprises in macroeconomic announcements quickly fade away. Hence, it is difficult to detect their effect using daily measures of market illiquidity.

that is not captured by VPIN and not associated with informational asymmetries around macro-economic announcements.

Insert Table IX about here

4 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 pricing efficiency. This conclusion does not necessarily hold for liquidity. Indeed, a fast response of arbitrageurs to opportunities created by asynchronous price adjustments (toxic arbitrage opportunities) is a source of illiquidity because it increases picking off risk for liquidity suppliers. Through this channel, faster arbitrage can impair market liquidity. We provide evidence for this channel using a sample of high frequency triangular arbitrage opportunities. Specifically we find that daily bid-ask spreads for the currency pairs in our sample are larger when the frequency of toxic arbitrage opportunities is higher. Moreover, an increase in picking off risk by arbitrageurs raises measures of market illiquidity.

Thus, setting speed limits in financial markets might help to strike a balance between pricing efficiency and liquidity. For instance, in 2013, EBS introduced a “latency floor” in the Australian dollar/U.S. dollar and the Swiss franc/U.S. dollar and consider extending this initiative to other currencies.⁵⁰

⁵⁰See “EBS to consider latency floor extensions within weeks” in FX.Week, December 2013 available at <http://www.fxweek.com/>.

Appendix

Derivation of Implication 4

Using equations (11), (17) and (14), we obtain:

$$E(D_t) = (\lambda_t^*(S_{tX}^*)(1+r))^{-1} = c^m(Q\alpha(1-\alpha)\sigma f(p^*))^{(-1)},$$

where $f(p^*) = \frac{p^*}{\alpha(3-p^*)p^*+(1-\alpha)}$. We have

$$\frac{\partial f}{\partial p^*} = \frac{(p^*)^2}{(\alpha(3-p^*)p^*+(1-\alpha))^2} > 0.$$

Now consider the following scenario (scenario 1): a decrease in c^a and c^m such that r increases (i.e., the decrease in c^a is larger than the decrease in c^m). In this case, $f(p^*)$ gets larger because p^* increases with r . Thus, as $E(D_t)$ is inversely related to $f(p^*)$ and increases with c^m , we deduce that $E(D_t)$ decreases in scenario 1, despite the fact that as r is larger, the bid-ask spread is higher in equilibrium.

Using the fact that $r = c^m/c^a$, we can also rewrite $E(D_t)$ using (4) as:

$$E(D_t) = c^a(Q\alpha(1-\alpha)\sigma g(p^*))^{(-1)},$$

where $g(p^*) = \frac{(1-p^*)}{\alpha(3-p^*)p^*+(1-\alpha)}$. We have:

$$\frac{\partial g}{\partial p^*} = -\frac{(1-\alpha) + \alpha(3-2p^* + (p^*)^2)}{(\alpha(3-p^*)p^*+(1-\alpha))^2} < 0.$$

Now, consider scenario 2: a decrease in c^a and c^m such that r decreases (i.e., the decrease in c^a is smaller than the decrease in c^m). In this case, $g(p^*)$ gets larger because p^* increases with r . Thus, as $E(D_t)$ is inversely related to $g(p^*)$ and increases with c^a , we deduce that $E(D_t)$ decreases in scenario 2, as well.

Figures

Figure I: Equilibrium Arbitrageur's Speed and Illiquidity

This figure shows (a) the competitive bid-ask spread as a function of the arbitrageur's speed, γ_t (dashed line) and (b) the arbitrageur's optimal speed as a function of the bid-ask spread in asset X, S_{tX} (dashed line). For each values of S_{tX} and γ_t , the market maker's speed, λ_t , is fixed at its equilibrium level (that is, $\lambda_t = \gamma_t/r$; see equation (13)). Thus, the equilibrium value of the bid-ask spread and the arbitrageur's speed is at the intersection of the dashed line and the plain line. The two equilibria shown in the figure are obtained for two different values of the arbitrageur's marginal cost of speed: $c^a = 0.5$ (low cost) and $c^a = 1.5$ (high cost). Other parameters are $c^m = 1$, $\alpha = 0.2$, $\sigma = 1$, and $Q = 1$.

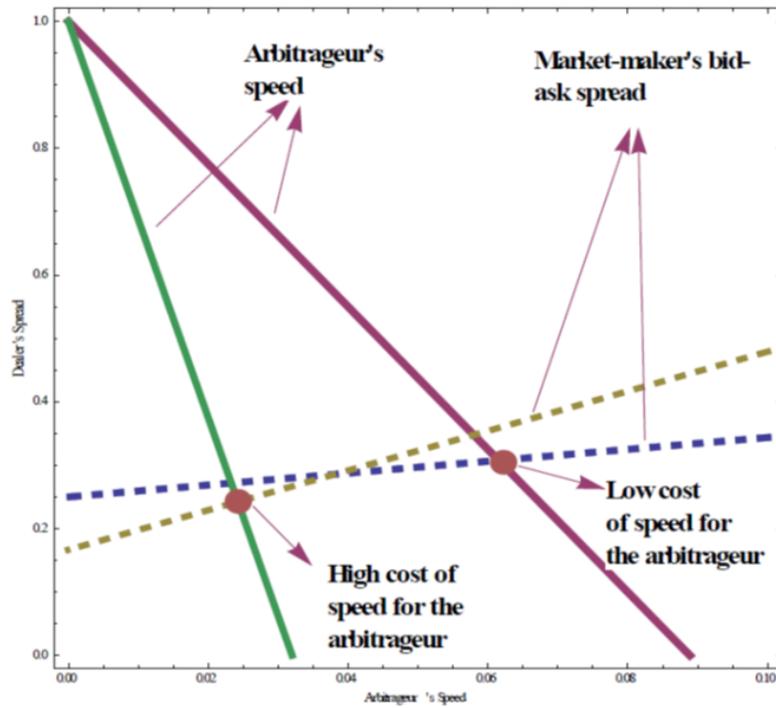
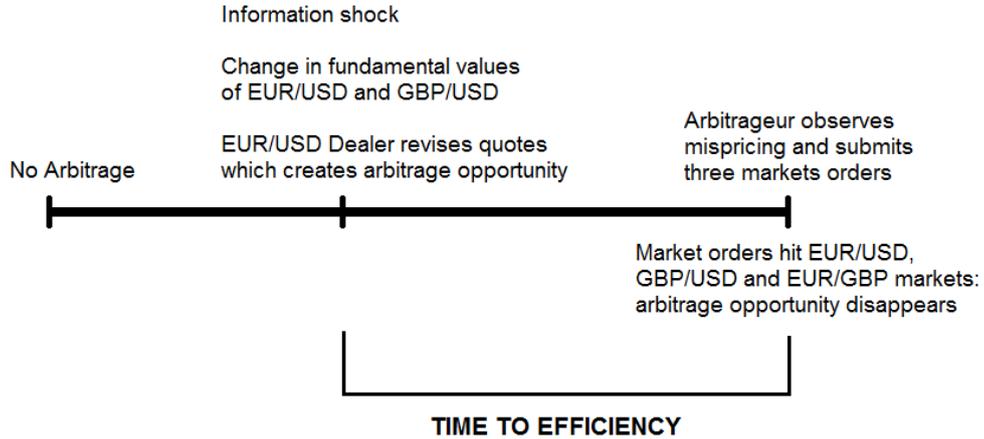


Figure II: Arbitrage Opportunities Can Terminate in Two Ways

This figure shows the two ways in which a triangular arbitrage opportunity can terminate. In Scenario 1, an arbitrageur is first to detect the arbitrage opportunity and submit three market orders (one in each currency pair) to exploit the opportunity, as explained in the text. The arbitrage opportunity then terminates with the submission of market orders. In Scenario 2, market makers update their quotes before an arbitrageur exploits the opportunity. The arbitrage opportunity terminates with a quote update in at least one limit order book of the currency pairs in our sample.

Scenario 1



Scenario 2

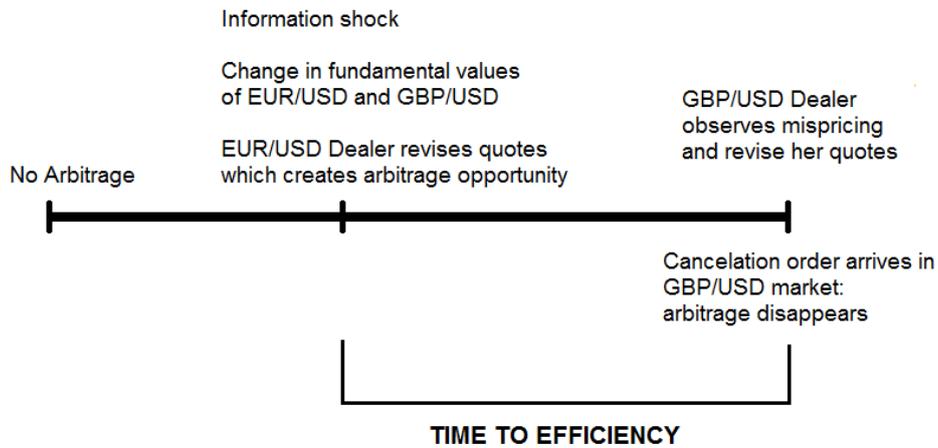
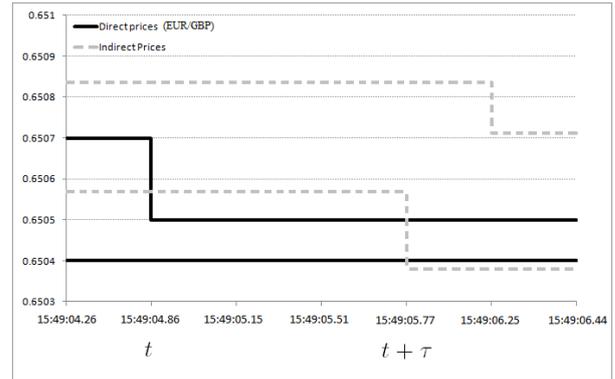
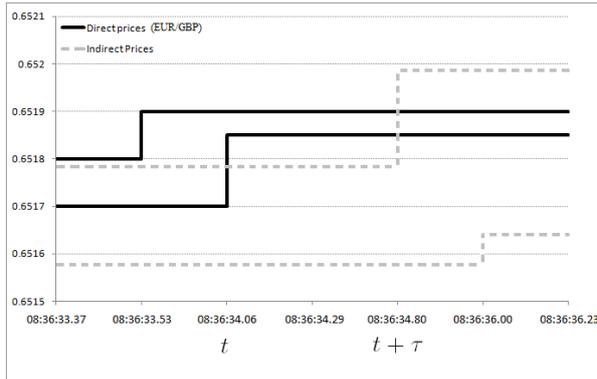


Figure III: Toxic vs. Non-Toxic Arbitrage Opportunities

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

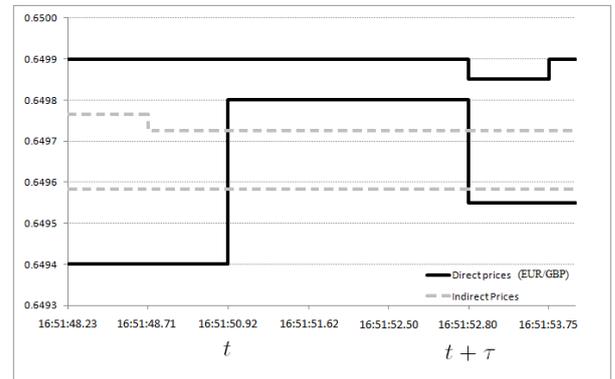
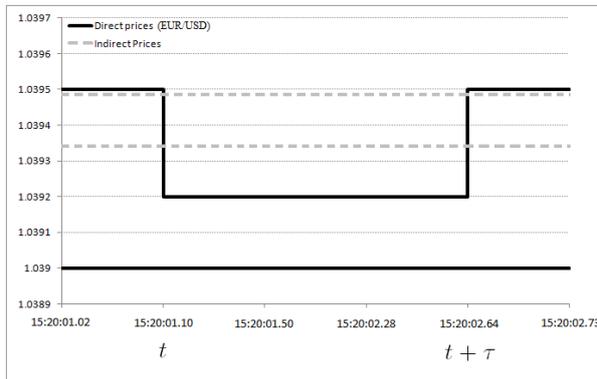


Figure IV: Autoquote and Message to Trade Ratio.

This figure presents a time series of the ratio of the daily number of orders to the daily number of trades in Reuters D-3000 for the three currency pairs in our sample. The dashed lines indicate the average levels of this ratio before and after July, 1st 2003.

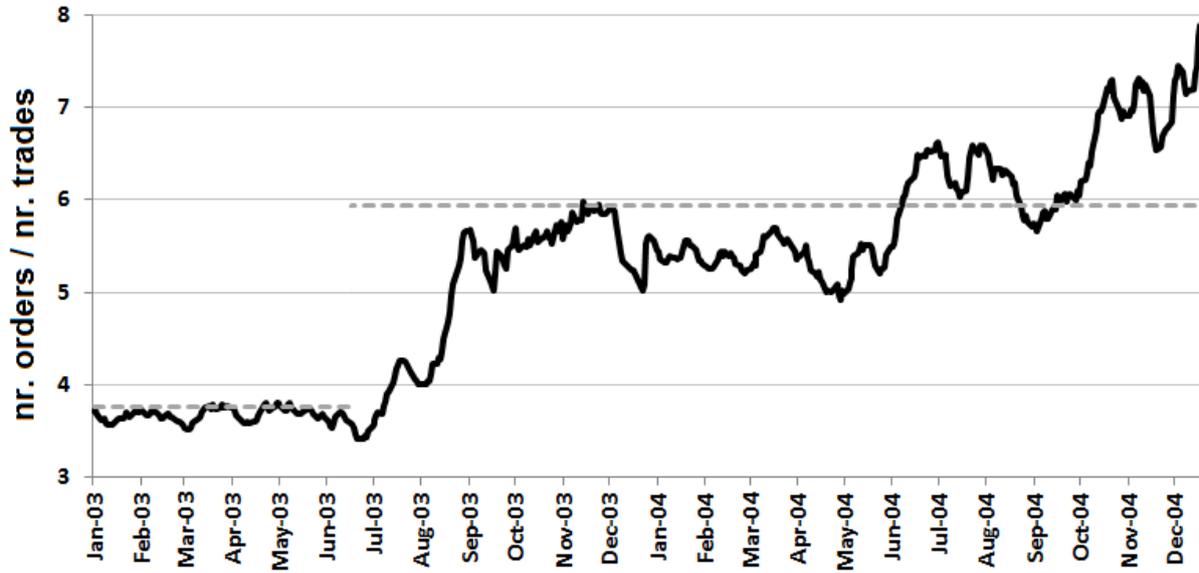
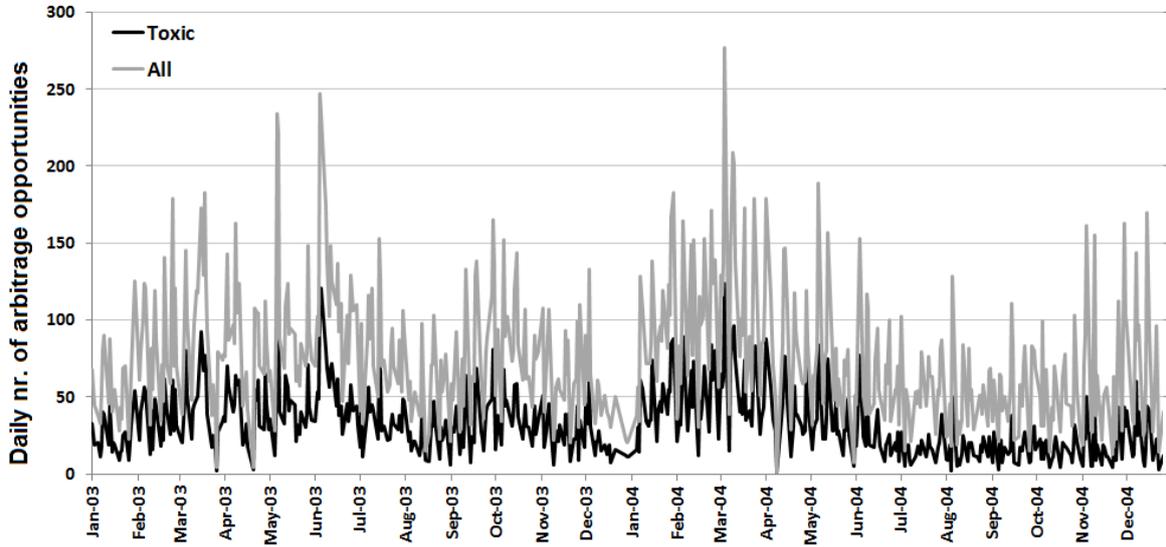


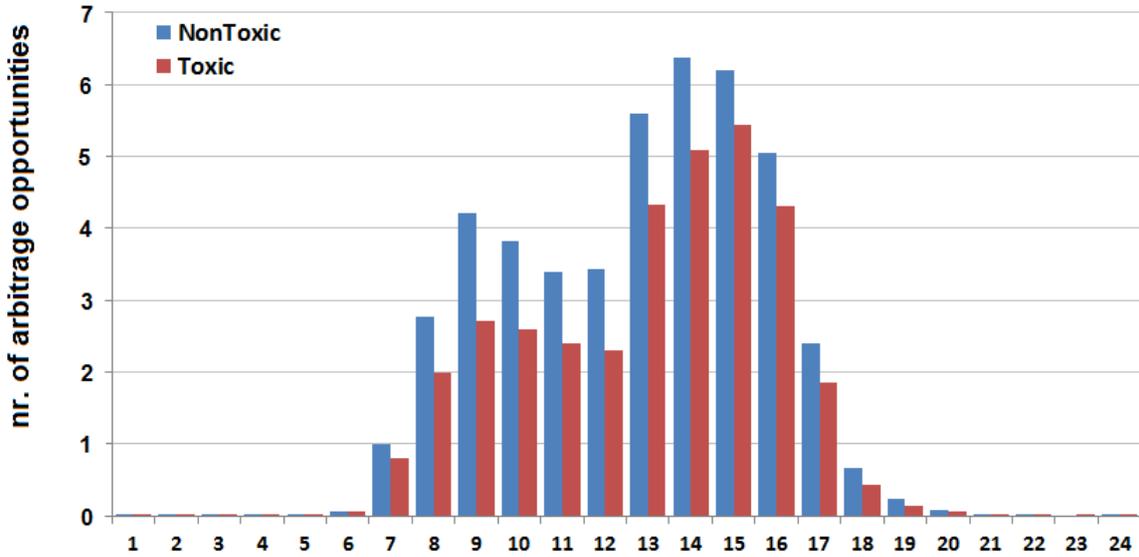
Figure V: Number of Arbitrage Opportunities

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



Tables

Table II: Where Do Toxic Arbitrage Opportunities Begin?

This table reports the fraction of all toxic arbitrage opportunities that start in a given currency pair. The pair in which the arbitrage starts is the *initiating pair*. Other pairs are called *non-initiating pairs*.

	<i>Initiating Pair</i>	<i>Non-Initiating Pair</i>
GBP/USD	28.68%	71.32%
EUR/USD	50.98%	49.02%
EUR/GBP	20.32%	79.68%

Table III: Trading Activity around Toxic and Non-Toxic Arbitrage Opportunities

For each currency pair $i \in \{GU, EU, EG\}$, we compute the absolute difference between the amount of base currency purchased and sold (order imbalance, measured in millions of base currency) by traders submitting market orders when the pair initiates the arbitrage (“initiating currency pair”) and when it does not (“non-initiating currency pair”) in toxic ($|OI_i^T|$) and non-toxic ($|OI_i^N|$) arbitrage opportunities. We then report the average difference between $|OI_i^N|$ and $|OI_i^T|$, 1 second, 5 seconds, and 10 second after the arbitrage begins when the currency pair initiates the arbitrage opportunity and when it does not (t-stats are reported in parentheses). Indexes GU , EU , and EG refer to the GBP/USD, EUR/USD, and EUR/GBP currency pairs, respectively. The sample period is from January 2, 2003 to December 30, 2004.

		<i>Initiating Currency Pair</i>			<i>Non-Initiating Currency Pair</i>		
		GBP/USD	EUR/USD	EUR/GBP	GBP/USD	EUR/USD	EUR/GBP
1 second	$ OI_{GU}^N - OI_{GU}^T $	0.056 (2.40)				-0.159 (-9.11)	-0.168 (-7.09)
	$ OI_{EU}^N - OI_{EU}^T $		0.181 (8.78)		-0.067 (-7.02)		-0.070 (-3.49)
	$ OI_{EG}^N - OI_{EG}^T $			0.152 (4.74)	-0.134 (-6.62)	-0.117 (-6.58)	
5 seconds	$ OI_{GU}^N - OI_{GU}^T $	0.010 (0.26)				-0.166 (-4.26)	-0.219 (-4.80)
	$ OI_{EU}^N - OI_{EU}^T $		0.204 (6.96)		-0.044 (-2.07)		-0.079 (-2.77)
	$ OI_{EG}^N - OI_{EG}^T $			0.119 (2.13)	-0.167 (-3.82)	-0.091 (-2.39)	
10 seconds	$ OI_{GU}^N - OI_{GU}^T $	-0.271 (-5.08)				-0.386 (-8.21)	-0.430 (-7.57)
	$ OI_{EU}^N - OI_{EU}^T $		0.113 (3.44)		-0.096 (-3.18)		-0.128 (-3.00)
	$ OI_{EG}^N - OI_{EG}^T $			0.026 (0.34)	-0.134 (-2.24)	-0.114 (-2.19)	

Table IV: 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 (TTE_t^{nontox}) denotes the duration in seconds of toxic (non-toxic) arbitrage opportunities on day t ; $nrarb_t$ ($nrarb_t^{nontox}$) is the number of toxic (non-toxic) arbitrage opportunities in day t ; $PTAT_t^c$ ($PNTAT_t^c$) 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; $\hat{\sigma}_t$ ($\hat{\sigma}_t^{nontox}$) is the average size of toxic (non-toxic) arbitrage opportunities in day t (in basis points); $profit_t$ ($profit_t^{nontox}$) is the average profit in bps 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
<i>Panel A: Toxic arbitrage</i>								
TTE_t	0.894	0.301	0.262	0.725	0.847	1.006	4.060	
$nrarb_t$	32.01	20.83	1	17	28	43	124	
$PTAT_t^c$	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	
$\hat{\sigma}_t$	3.535	0.757	2.224	3.112	3.439	3.843	13.61	
$profit_t$	1.427	0.277	1.115	1.336	1.401	1.470	6.668	
<i>Panel B: Non-toxic arbitrage</i>								
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
$PNTAT_t^c$	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
<i>Panel C: Illiquidity Measures</i>								
$spread_{GUt}$	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 IV continued.

Variable	Mean	Std.Dev.	Min	Q1	Median	Q3	Max
<i>Panel D: Common variables</i>							
<i>vol_{GU}</i>	0.683	0.268	0.117	0.532	0.622	0.753	2.456
<i>vol_{EU}</i>	0.827	0.386	0.258	0.616	0.744	0.920	4.363
<i>vol_{EG}</i>	0.387	0.094	0.203	0.325	0.381	0.440	1.256
<i>nrorders_{GU}</i>	17.65	6.091	0.576	12.51	17.62	22.45	32.22
<i>nrorders_{EU}</i>	19.05	6.831	0.188	14.78	18.15	22.88	44.07
<i>nrorders_{EG}</i>	14.77	5.810	0.307	9.326	16.07	19.41	28.93
<i>trsize_{GU}</i>	1.386	0.043	1.247	1.357	1.382	1.415	1.509
<i>trsize_{EU}</i>	1.401	0.056	1.000	1.365	1.396	1.434	1.605
<i>trsize_{EG}</i>	1.548	0.076	1.294	1.497	1.541	1.591	1.853
<i>nrtr_{GU}</i>	4.692	1.505	0.175	3.634	4.523	5.639	9.611
<i>nrtr_{EU}</i>	2.365	0.707	0.027	1.859	2.377	2.870	4.103
<i>nrtr_{EG}</i>	2.841	0.811	0.068	2.301	2.761	3.318	6.329
Nr. days				498			

Table V: 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 denotes the duration in seconds of toxic (non-toxic) arbitrage opportunities on day t ; $PTAT_t^c$ ($PNTAT_t^c$) 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 ; $\hat{\varphi}_t$ is the number of toxic arbitrage opportunities in day t divided by the number of arbitrage opportunities in that day; $\hat{\sigma}_t$ 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 for illiquidity measures computed using EBS data. The last row of the table 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	$\hat{\sigma}_t$	$\hat{\varphi}_t$	$spread_{GUt}$	$spread_{EUt}$	$spread_{EGt}$	$slope_{GUt}$	$slope_{EUt}$	$slope_{EGt}$	$espread_{GUt}$	$espread_{EUt}$	$espread_{EGt}$	$PNTAT_t^c$
$PTAT_t^c$	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	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
$\hat{\sigma}_t$		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	0.955	0.680	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
$illiq_{EBS}$				0.273	0.735	0.414	0.174	0.814	0.233	0.314	0.594	0.521	

Table VI: Effect of $PTAT^c$ on Market Illiquidity

In Panel A, we report OLS estimates of the following regression for $i \in \{GU, EU, EG\}$:

$illiq_{it} = \omega_i + \xi_t + b_1 PTAT_t^c + b_2 vol_{it} + b_3 \hat{\varphi}_t + b_4 \hat{\sigma}_t + b_5 trsize_{it} + b_6 nrorders_{it} + b_7 illiq_{it}^{EBS} + \varepsilon_{it}$, where indexes GU , EU , and EG refer to the GBP/USD, EUR/USD, and EUR/GBP currency pairs, respectively. In panel B, we report IV estimates of the previous regression in which we instrument $PTAT_t^c$ with the introduction of AutoQuote on Reuters D-3000 (see the text). The first stage regression of the IV is: $PTAT_t^c = \omega_i + \xi_t + a_1 AD_t + a_2 vol_{it} + a_3 \hat{\varphi}_t + a_4 \hat{\sigma}_t + a_5 trsize_{it} + a_6 nrorders_{it} + a_7 illiq_{it}^{EBS} + u_{it}$, where AD_t is a dummy variable equal to one after July 2003 and zero before. $PTAT_t^c$ 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; $\hat{\sigma}_t$ 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 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 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: OLS

	<i>spread</i>	<i>espread</i>	<i>slope</i>
$PTAT^c$	-0.017 (-0.27)	0.075 (1.26)	0.004 (0.11)
vol	0.299 (8.00)	0.371 (7.56)	0.177 (9.45)
$\hat{\varphi}$	0.622 (7.47)	0.486 (7.38)	0.407 (9.33)
$\hat{\sigma}$	0.157 (10.1)	0.188 (5.32)	0.074 (10.1)
$trsize$	-0.341 (-3.37)	-0.275 (-2.19)	-0.385 (-7.23)
$nrorders$	-0.011 (-8.48)	-0.009 (-9.15)	-0.007 (-9.99)
$illiq^{EBS}$	0.026 (3.47)	-0.004 (-1.79)	0.002 (2.15)
$Adj.R^2$	86.53%	87.83%	82.89%
Currency pair FE	YES	YES	YES
Month dummies	YES	YES	YES

Panel B: Instrumental Variable

	<i>spread</i>		<i>espread</i>		<i>slope</i>	
	1 st stage	2 nd stage	1 st stage	2 nd stage	1 st stage	2 nd stage
$PTAT^c$		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)
$\hat{\sigma}$	-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 VII: Impact of “Autoquote” on 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; $\hat{\sigma}_t$ 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 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)
$\hat{\sigma}$	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 VIII: Currency-Level Tests

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

$illiq_{it} = \omega_i + \xi_t + b_1 PTAT_t^c + b_2 vol_{it} + b_3 \hat{\varphi}_t + b_4 \hat{\sigma}_t + b_5 trsize_{it} + b_6 nrorders_{it} + b_7 illiq_{it}^{EBS} + \varepsilon_{it}$ for $i \in \{GU, EU, EG\}$, where indexes GU , EU , and EG refer to the GBP/USD, EUR/USD, and EUR/GBP currency pairs, respectively. For each currency pair, we estimate this regression using an IV approach in which we instrument $PTAT_t^c$ with the introduction of AutoQuote on Reuters D-3000 (see the text). The first stage regression of the IV is: $PTAT_t^c = \omega_i + \xi_t + a_1 AD_t + a_2 vol_{it} + a_3 \hat{\varphi}_t + a_4 \hat{\sigma}_t + a_5 trsize_{it} + a_6 nrorders_{it} + a_7 illiq_{it}^{EBS} + u_{it}$, where AD_t is a dummy variable equal to one after July 2003 and zero before. $PTAT_t^c$ 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; $\hat{\sigma}_t$ 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 ; t-statistics are calculated using robust standard errors correcting for heteroscedasticity and serial correlation. The sample period is from January 2, 2003 to December 30, 2004.

		<i>spread</i>		<i>espread</i>		<i>slope</i>	
		1 st stage	2 nd stage	1 st stage	2 nd stage	1 st stage	2 nd stage
GBP/USD	$PTAT^c$		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)
	$\hat{\sigma}$	-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	$PTAT^c$		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)
	$\hat{\sigma}$	-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	$PTAT^c$		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)
	$\hat{\sigma}$	-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	
<i>Month dummies</i>		YES		YES		YES	

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

This table presents the estimation results of two-stage least squares instrumental variable panel regressions of the impact of $PTAT^c$ on market illiquidity. The first-stage regression is:

$$PTAT_t^c = \omega_i + \xi_t + a_1 AD_t + a_2 vol_{it} + a_3 \hat{\varphi}_t + a_4 \hat{\sigma}_t + a_5 trsize_{it} + a_6 nrorders_{it} + a_7 illiq_{it}^{EBS} + a_8 VPIN_{it} + a_{10} macro_t^{US} + a_{11} macro_t^{UK} + a_{12} macro_t^{EMU} + u_{it} \quad \text{for } i \in \{GU, EU, EG\}.$$

and the second-stage regression is:

$$illiq_{it} = \omega_i + \xi_t + b_1 \widehat{PTAT}_t^c + b_2 vol_{it} + b_3 \hat{\varphi}_t + b_4 \hat{\sigma}_t + b_5 trsize_{it} + b_6 nrorders_{it} + b_7 illiq_{it}^{EBS} + b_8 VPIN_{it} + b_9 macro_t^{US} + b_{10} macro_t^{UK} + b_{11} macro_t^{EMU} + u_{it}, \quad \text{for } i \in \{GU, EU, EG\}$$

where \widehat{PTAT}_t^c is the fitted values of $PTAT_t^c$ from the first-stage regression. Indexes GU , EU , and EG refer to the GBP/USD, EUR/USD, and EUR/GBP currency pairs, respectively. $PTAT_t^c$ 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; $\hat{\sigma}_t$ 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 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.

	<i>spread</i>		<i>espread</i>		<i>slope</i>	
	1 st stage	2 nd stage	1 st stage	2 nd stage	1 st stage	2 nd stage
$PTAT^c$		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)
$\hat{\sigma}$	-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.
- Amihud, Y. and Mendelson, H.: 1980, Dealership market: Market-making with inventory, *Journal of Financial Economics* **8**, 31–53.
- 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.
- Bank of International Settlements: 2010, Report on global foreign exchange market activity in 2010, *Triennial Central Bank Survey* p. December.
- Biais, B., Foucault, T. and Moinas, S.: 2011, Equilibrium fast trading, *Working paper, Toulouse Business School* .
- 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* .
- CFEC: 2006, Minutes of the Canadian Foreign Exchange committee meeting, *Canadian Foreign Exchange Committee* .
- 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.

- Chakrabarty, B. and Moulton, P.: 2012, Earnings announcements and attention constraints: The role of market design, *Journal of Accounting and Economics* **53**(3), 612–634.
- 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.
- Coughenour, J. and Corwin, S.: 2008, Limited attention and the allocation of effort in securities trading, *Journal of Finance* **63**, 3031–3067.
- 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., Kadan, O. and Kandel, E.: 2013, Liquidity cycles and make/take fees in electronic markets, *Journal of Finance* **68**, 299–341.
- 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.
- Halder, A.: 2011, The race to zero, *Speech at the International Economic Association Sixteenth World Congress* .
- 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.: 2011, 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*.
- 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* .
- Peel, D. and Taylor, M. P.: 2002, Covered interest rate arbitrage in the interwar period and the Keynes-Einzig conjecture, *Journal of Money, Credit and Banking* **34**(1), 51–75.
- 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.
- Schmidt, D.: 2013, Investors’ attention and stock covariation, *Working paper, HEC Paris* .
- 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.
- Storkenmaier, A. and Wagener, M.: 2011, Do we need a European “National Market System”? Competition, arbitrage, and suboptimal executions, *Working paper, Karlsruhe Institute of Technology* .
- 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.