Modelling limit order books with bilateral trading agreements

Martin D. Gould,[1,2,3,*] Mason A. Porter,[1,3] Nikolaus Hautsch,[4] Stacy Williams,[5] Mark McDonald,[5] Daniel J. Fenn,[5] and Sam D. Howison[1,2]

[1]*Oxford Centre for Industrial and Applied Mathematics, Mathematical Institute, University of Oxford, Oxford OX1 3LB, UK*
[2]*Oxford-Man Institute of Quantitative Finance, University of Oxford, Oxford, OX2 6ED, UK*
[3]*CABDyN Complexity Centre, University of Oxford, Oxford OX1 1HP, UK*
[4]*Chair of Econometrics, Humboldt-Universit ät zu Berlin, Berlin, 10178, Germany*
[5]*FX Quantitative Research, HSBC Bank, London E14 5HQ, UK*

More than half of the world's financial markets currently use a limit order book (LOB) mechanism to facilitate trade. For markets where trade is conducted through a central counterparty (CCP), trading platforms disseminate the same information to all traders in real time and all traders are able to trade with all others. By contrast, in markets that operate under bilateral trade agreements (BTAs), traders only view the LOB activity from their bilateral trading partners and cannot trade with anyone with whom they do not possess a BTA. In this paper, we examine how BTAs affect trade in the foreign exchange (FX) spot market. Using historical data from an electronic BTA LOB trading platform, we present a statistical analysis of how BTAs influence the prices paid by traders and highlight the challenges that BTAs pose for modelling. By performing model-based inference on the network of BTAs in this market, we estimate that most traders have relatively few BTA partners. We conclude with a discussion of how BTAs affect market stability.

Keywords: Complex systems; foreign-exchange market; high-frequency data; latent parameter estimation; stochastic modelling

I. INTRODUCTION

More than half of the markets in today's highly competitive and relentlessly fast-paced financial world use a *limit order book (LOB)* mechanism to facilitate trade [14]. In a traditional LOB, trade occurs via a designated *central counterparty (CCP)*. By contrast, the three most widely used LOB platforms in the foreign exchange (FX) spot market use a bilateral trade structure [13, 15]. Such *bilateral trade agreements (BTAs)* provide several benefits for traders, including explicit control over counterparty risk and reduced times to trade completion. This paper examines how BTAs affect trade via LOBs in the FX spot market. Due to space considerations, we present only a small selection of our results regarding BTA LOBs here. For a detailed discussion of our methods and other results, please refer to our series of working papers on the topic [8–10].

II. BILATERAL TRADE AGREEMENT LIMIT ORDER BOOKS

Let $\Theta = \{\theta_1, \theta_2, \ldots, \theta_N\}$ denote the set of traders in a single-asset LOB, and let $x = (p, \omega_x, \tau_x)$ denote a buy (respectively, sell) order submitted at time t_x, with price p_x and size ω_x. When a trader submits a buy (respectively, sell) order x, an LOB's trade-matching algorithm checks whether it is possible to match x to a previously submitted sell (respectively, buy) order. If so, the matching occurs immediately and x is called a market order. If not, x becomes *active*, and is called a limit order. An order remains active until either it matches to an incoming sell (respectively, buy) market order or it is cancelled. An *LOB* $\mathcal{L}(t)$ is the set of all active orders in a market at time t.

* Corresponding author: gouldm@maths.ox.ac.uk

In CCP LOBs, all traders are able to trade with all others. By contrast, in BTA LOBs, each trader $\theta_i \in \Theta$ provides the exchange with a *block list* of other traders with whom they are unwilling to trade. Traders θ_i and θ_j are able to trade with each other if and only if θ_i does not appear on θ_j's block list, and vice-versa, in which case we call them *trading partners* and write $\theta_i \leftrightarrow \theta_j$. Otherwise, we write $\theta_i \not\leftrightarrow \theta_j$. Traders in a BTA LOB are not able to view the whole of $\mathcal{L}(t)$. Instead, each trader $\theta_i \in \Theta$ views their local LOB:

$$\mathcal{L}_{\theta_i}(t) := \{x \in \mathcal{L}(t) |\, x \text{ is owned by some } \theta_j \text{ with } \theta_j \leftrightarrow \theta_i\}\,.$$

Figure 1 illustrates several definitions related to BTA LOBs. Given a pair of traders $\theta_i \neq \theta_j$ with $\theta_j \not\leftrightarrow \theta_i$, active orders owned by θ_j are not visible to θ_i and are not considered for matching against any of θ_i's orders. Each trader θ_i in a BTA LOB also views a time series of all trades that have occurred in the market, irrespective of whether the traders involved were θ_i's trading partners.

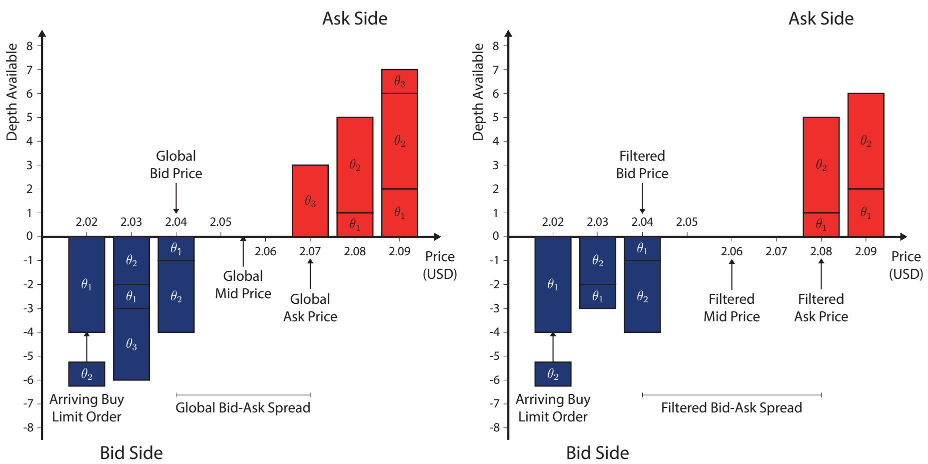

FIG. 1. Schematic of a LOB $\mathcal{L}(t)$ (left) and θ_2's corresponding local LOB $\mathcal{L}_{\theta_2}(t)$, for a BTA LOB in which $\Theta = \{\theta_1, \theta_2, \theta_3\}$, $\theta_1 \leftrightarrow \theta_2$, $\theta_1 \leftrightarrow \theta_3$ and $\theta_2 \not\leftrightarrow \theta_3$. In the diagram, each order is labelled according to its owner, but this information is not made available to traders in real BTA LOBs (nor is it recorded in the data we study).

III. SOME STATISTICAL PROPERTIES OF BTA LOBS

In this section, we present a small selection of our statistical analyses describing how BTAs affect trade. We performed our analyses using historical data from the BTA LOB trading platform Hotspot FX, which is the third most widely used multi-institution electronic trading platform in the FX spot market [13]. The data describes all arrivals and departures of limit orders and all market orders, during the "normal trading hours" of 08:00 – 17:00 BST, for 25 trading days during May 2010 and June 2010. We studied three different currency pairs, namely Euro/U.S. dollar (EUR/USD), Pounds sterling/U.S. dollar (GBP/USD), and Euro/Pounds sterling (EUR/GBP). Full details of how we pre-processed and cleaned the data are available in our working paper [10].

Most statistical studies and models of CCP LOBs define the "price" of a buy (respectively, sell) order x to be the difference between p_x and the price of the highest priority active buy (respectively, sell) order in $\mathcal{L}(t_x)$. We call such prices *global relative prices,* because they are derived from the "global" LOB $\mathcal{L}(t)$. Several empirical studies covering a wide range of different CCP LOBs have reported distributions describing order flow to take simple parametric forms when measured using global relative pricing (see, e.g., [1, 3, 12, 16]). By contrast, we found that studying order flow using global relative pricing resulted in distributions that exhibited several local minima and maxima (see Figure 2 (top)). Given the BTA structure of Hotspot FX, these findings are not surprising: In a BTA LOB, a trader $\theta_i \in \Theta$ does not know the prices of the highest priority active orders in $\mathcal{L}(t)$, so cannot use these prices when making their trading decisions. However, all traders in a BTA LOB observe the same time series of previous trades, so this provides an alternative anchor for measuring relative prices. We define the *trade relative price* of a buy (resp., sell) order x to be the difference between p_x and the price at which the most recent seller-initiated (resp., buyer-initiated) trade occurred before t_x. When measured using trade-relative pricing, we found the distributions describing order flow to exhibit simple parametric forms without multiple local extrema (see Figure 2 (bottom)).

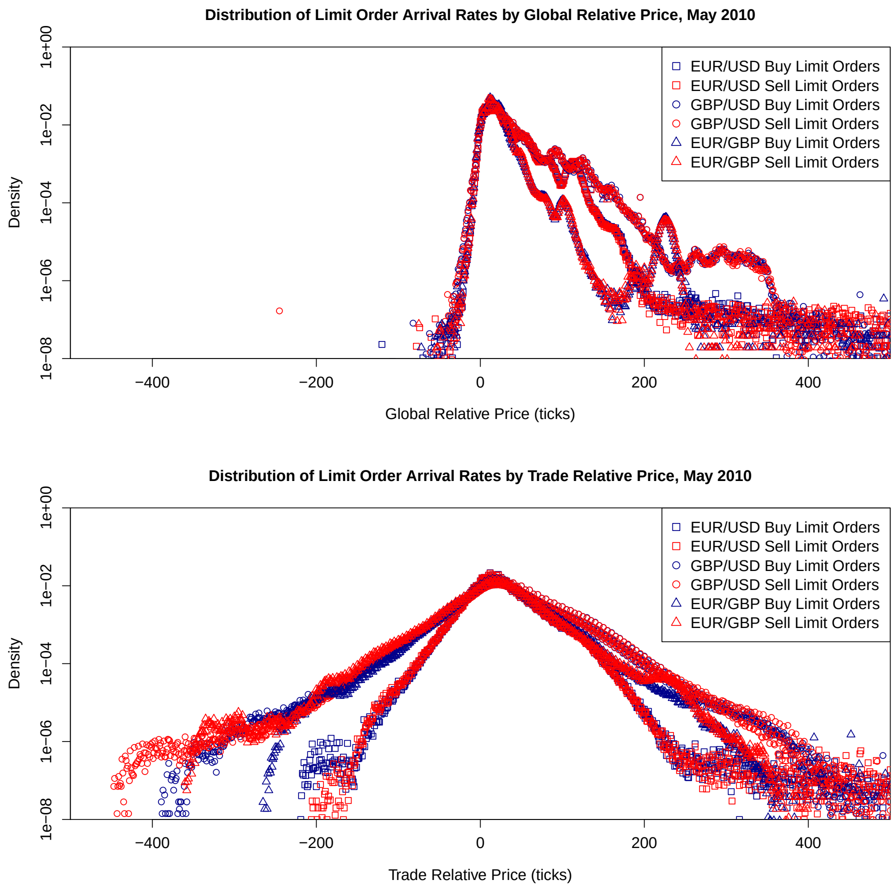

FIG. 2. Empirical densities of arrival rates of incoming orders during May 2010, by global relative price (top) and trade relative price (bottom).

For each incoming buy (resp., sell) market order x, we calculated the difference between the price at which the market order matched and the price of the highest priority active sell (resp., buy) order in $\mathcal{L}(t_x)$, and labelled this Δ_x (see Figure 3 (top)). On Hotspot FX, a single price level equals 0.1 basis points, measured in units of the base currency. More than 30% of market orders had $\Delta_x \geq 1$ and more than 5% of market orders had $\Delta_x \geq 10$. Given that the modal order size was 1 million units of the base currency in all cases, the resulting price difference for the trade is quite substantial.

Figure 3 (bottom) displays the mean fraction of active buy (resp., sell) orders that were bypassed by incoming crossing orders x (i.e., sell (resp., buy) orders whose price

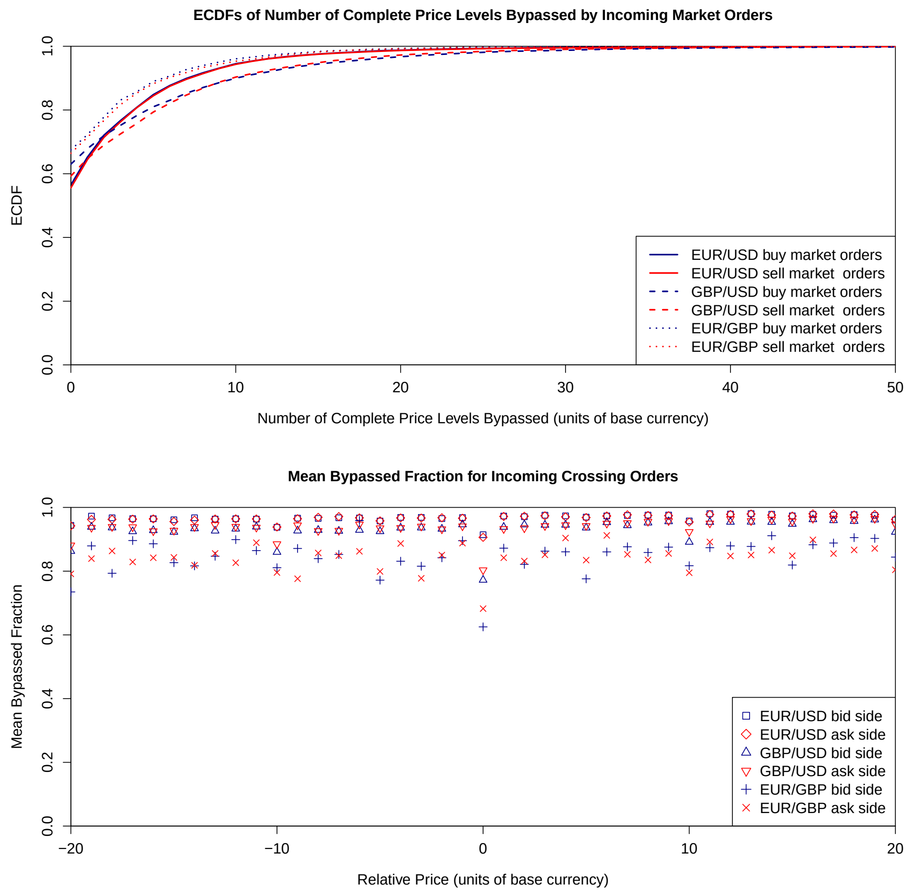

FIG. 3. ECDFs of Δ_x for incoming market orders (top) and mean fraction of active orders bypassed by incoming crossing orders (bottom).

was lower (resp., higher) than the price of the highest priority active buy (resp., sell) order in $\mathcal{L}(t_x)$). Most crossing orders x bypassed a large portion of the active orders in $\mathcal{L}(t_x)$, suggesting that most traders had very few trading partners in the market.

IV. AN AGENT-BASED MODEL OF BILATERAL TRADE

In this section, we use a simple agent-based model (ABM) populated by homogeneous agents to perform model-based inference about the structure of the BTA network on Hotspot FX. In our model, each agent possesses a private buy valuation $b_{\theta_i}(t)$ and a private sell valuation $a_{\theta_i}(t)$ of the asset being traded, each of which evolve as an asymmetric discrete time random walk on a discrete pricing grid. At each time step t, an agent $\theta_i \in \Theta$ is chosen uniformly at random. With probability $1/2$, θ_i revises $b_{\theta_i}(t)$, otherwise they revise $a_{\theta_i}(t)$. Independently, with probability $0.5 + k$, their revision of $b_{\theta_i}(t)$ or $a_{\theta_i}(t)$ is one price level towards the price at which the most recent trade before time t occurred, otherwise it is one price level away from it. A trade occurs whenever $b_{\theta_i}(t)$ equals or exceeds $a_{\theta_j}(t)$, for a pair of traders $\theta_i \neq \theta_j$ for which $\theta_i \leftrightarrow \theta_j$. Traders who are not trading partners never trade with each other, irrespective of their private valuations. Therefore, for a pair of traders $\theta_i \not\leftrightarrow \theta_j$, it is possible for $b_{\theta_i}(t) > a_{\theta_j}(t)$. Figure 4 displays a single simulation run of the ABM. The model is designed to capture 2 key features of BTA markets, while being as simple as possible. First, each trader $\theta_i \in \Theta$ may only trade with their trading partners. Second, the only common information available

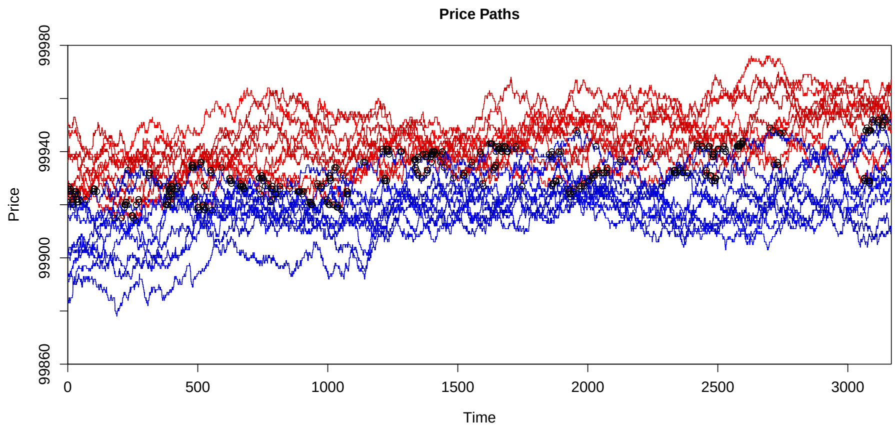

FIG. 4. A typical run of the ABM. Blue paths denote the evolution of the $b_{\theta_i}(t)$, red paths denote the evolution of the $a_{\theta_i}(t)$, and black circles denote trades.

to all traders is the time series of previously traded prices.

Besides the drift parameter k, the model's only other "parameter" is the network of BTAs. We used the method of simulated moments (MSM) [5, 11] to estimate these parameters, given the Hotspot FX data. We used 4 moment conditions, which measured the standard deviation and absolute mean price change between successive trades and the mean and standard deviation of Δ_x. We used the threshold accepting heuristic with linearly decreasing thresholds [4, 7] to search the space of BTA networks. Our results suggest that EUR/USD had the largest number of active traders, followed by GBP/USD, followed by EUR/GBP. This agrees with the ordering of the currency pairs by their total volume of trade (although we did not use this fact when performing the MSM). In each case, the mean percentage of bilateral trade partners for each individual trader was less than 20% of the total number of traders in the market. We note that with the parameters fitted by MSM, the ABM also reproduced several non-trivial properties exhibited by real price series in the Hotspot FX data, including volatility clustering, aggregational Gaussianity, positive excess kurtosis of the unconditional distribution of returns, and the absence of linear autocorrelation of price changes [2], although we leave further investigation of these properties for future research.

V. CONCLUSIONS

BTAs directly affect the evolutions of LOBs, and many market orders in BTA LOBs match at prices that are strictly worse for their owners than the best prices available in $\mathcal{L}(t)$. Because traders do not view the evolution of the full LOB $\mathcal{L}(t)$ in real time, distributions describing order flow measured in global relative pricing exhibit several local optima. It is plausible that these local optima might be caused by each trader choosing the prices for their orders based on their local LOB $\mathcal{L}_{\theta_i}(t)$. However, modelling BTA LOBs in this way is unappealing. Because each individual trader acts to satisfy their personal trading needs (which themselves change over time), the order flow generated by any single trader is erratic. However, many empirical studies have concluded that the aggregate order flow generated by the many heterogeneous agents interacting in a given LOB is predictable and stable over time (see, e.g., [6, 12]). This is akin to a central result from statistical mechanics: The aggregate motion of an ensemble of particles can be predicted very accurately, even though the motion of any single particle in the ensemble is highly unpredictable. A model in which each individual trader responds to each individual $\mathcal{L}_{\theta_i}(t)$ is akin to a model that tracks each individual particle in the ensemble. By contrast, trade relative prices are common to all traders in a given BTA LOB, and measuring prices using this framework reveals robust statistical regularities in empirical data. We note that global relative pricing could also be useful for studying CCP LOBs, particularly in markets where the state of $\mathcal{L}(t)$ changes frequently.

The results we have presented suggest that for the currency pairs we studied, traders typically have a low percentage of BTA partners. There are several other statistics that we have not reported here that also support this hypothesis. This is interesting from the perspective of market stability. The FX spot market is regarded as highly liquid, yet it seems that most traders have access to only a small fraction of the total available liquidity, and the bid-ask spreads that individual traders observe are far wider than what is apparent from studying $\mathcal{L}(t)$ alone. Although traders can observe price formation via the time series of traded prices, they cannot observe the full microscopic description of supply and demand provided by viewing the whole of $\mathcal{L}(t)$. The fewer BTAs that exist in a market, the weaker the similarity between the local LOBs and the global LOB, the higher the informational asymmetry that exists between different traders, and the closer the market comes to functioning as a dark pool of liquidity.

---

[1] J.P. Bouchaud, M. Mézard, and M. Potters. Statistical properties of stock order books: empirical results and models. *Quantitative Finance*, 2(4):251–256, 2002.

[2] R. Cont. Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative Finance*, 1(2):223–236, 2001.

[3] R. Cont, S. Stoikov, and R. Talreja. A stochastic model for order book dynamics. *Operations Research*, 58(3):549–563, 2010.

[4] G. Dueck and T. Scheuer. Threshold accepting: A general purpose optimization algorithm appearing superior to simulated annealing. *Journal of Computational Physics*, 90(1):161–175, 1990.

[5] D. Duffie and K.J. Singleton. Simulated moments estimation of Markov models of asset prices. *Econometrica: Journal of the Econometric Society*, 61(4):929–952, 1993.

[6] J.D. Farmer, P. Patelli, and I.I. Zovko. The predictive power of zero intelligence in financial markets. *Proceedings of the National Academy of Sciences of the United States of America*, 102(6):2254–2259, 2005.

[7] M. Gilli and P. Winker. A global optimization heuristic for estimating agent based models. *Computational Statistics and Data Analysis*, 42(3):299–312, 2003.

[8] M. D. Gould, M. A. Porter, S. Williams, M. McDonald, D. J. Fenn, and S. D. Howison. Limit order books. *arXiv:1012.0349*, 2013.

[9] M.D. Gould, N. Hautsch, M.A. Porter, S. Williams, M. McDonald, and S.D. Howison. An agent-based model of bilateral trading. *In preparation*, 2013.

[10] M.D. Gould, M.A. Porter, S. Williams, M. McDonald, D.J. Fenn, and S.D. Howison. Statistical properties of foreign exchange limit order books. *In preparation*, 2013.

[11] B.S. Lee and B.F. Ingram. Simulation estimation of time-series models. *Journal of Econometrics*, 47(2):197–205, 1991.

[12] S. Mike and J.D. Farmer. An empirical behavioral model of liquidity and volatility. *Journal of Economic Dynamics and Control*, 32(1):200–234, 2008.

[13] E. Pan. FX trading and technology in 2012. Technical report, StreamBase Systems, New York, 2012.

[14] I. Roşu. A dynamic model of the limit order book. *Review of Financial Studies*, 22(11):4601–4641, 2009.

[15] K. von Kleist, C. Mallo, P. Mesny, and S. Grouchko. Triennial central bank survey: Report on global foreign exchange market activity in 2010. Technical report, Bank for International Settlements, Basel, Switzerland, 2010.

[16] I. Zovko and J.D. Farmer. The power of patience: a behavioral regularity in limit order placement. *Quantitative Finance*, 2(5):387–392, 2002.