Market Microstructure in the Foreign Exchange Spot Market

Martin Gould

CFM–Imperial Institute of Quantitative Finance

5 March 2015

Imperial College London

Lessons from a Crisis

The international financial crisis underlines the vital importance of monitoring and assessing all trading counterparties

'Your card is fine. I'm just checking that your bank hasn't expired'

- Counterparty Credit Risk
- Settlement Risk
- Market Risk
- Liquidity Risk

Open-Outcry Trading

Modern Markets

The New Trading Standard

New York Stock Exchange
Hong Kong Stock Exchange
Shenzhen Stock Exchange
Australian Securities Exchange
Helsinki Stock Exchange
Euronext

Limit order books are used to match buyers and sellers in more than half of the world's financial markets

Toronto Stock Exchange
London Stock Exchange
Tokyo Stock Exchange
Vancouver Stock Exchange
Swiss Stock Exchange

Why are LOBs so Popular?

- They give market participants the freedom to evaluate their own need for immediacy
  - They allow some market participants to demand immediacy (by placing a market order) ...
  - ...while others supply it (by placing a limit order)
- They provide market participants around the world real-time access to the market
- They present a detailed and accurate picture of market state
- They radically democratize the process of trading

Why are LOBs so Popular?

- They give market participants the freedom to evaluate their own need for immediacy
  - They allow some market participants to demand immediacy (by placing a market order) . . .
  - . . . while others supply it (by placing a limit order)
- They provide market participants around the world real-time access to the market
- They present a detailed and accurate picture of market state
- They radically democratize the process of trading

The Dangers of Democracy

The recent international financial crisis underlines the vital importance of monitoring and assessing all trading counterparties

Limit Order Books

Centralized

Limit Order Books

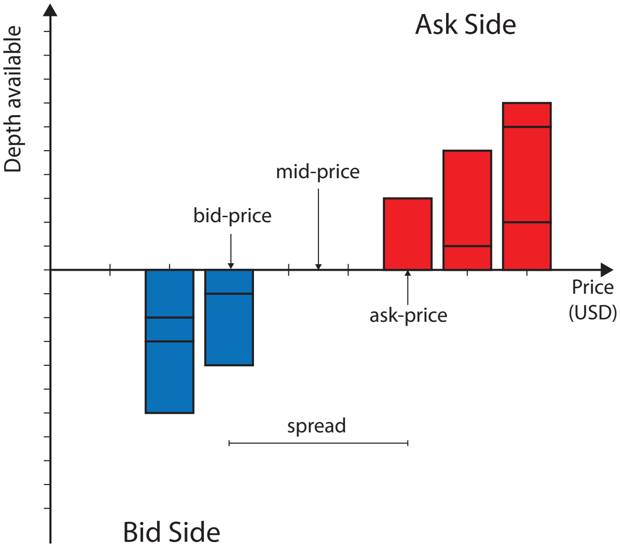

Depth available
Ask Side
mid-price
bid-price
ask-price
Price (USD)
spread
Bid Side

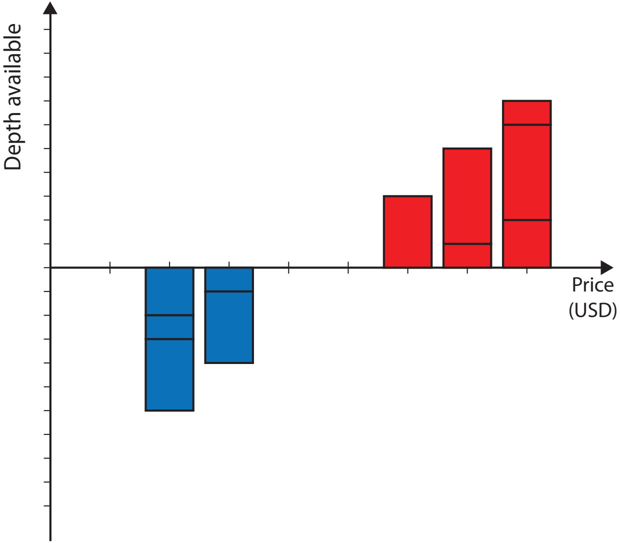

Depth available
Price
(USD)

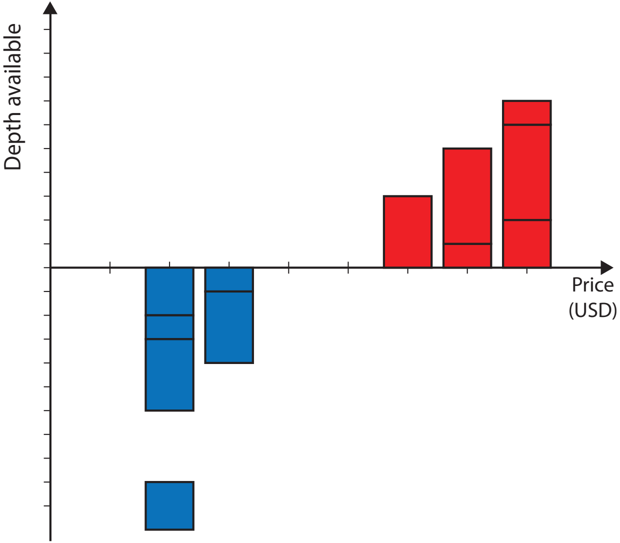

Depth available
Price
(USD)

Price Formation in a Centralized Limit Order Book

A Dynamic Example

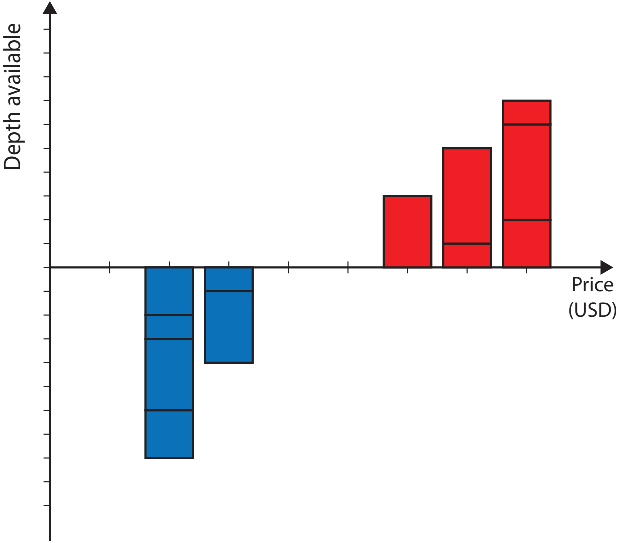

Price Formation in a Centralized Limit Order Book

A Dynamic Example

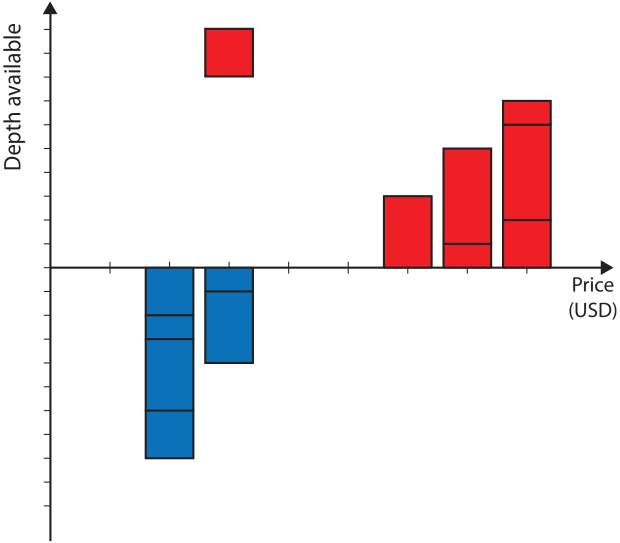

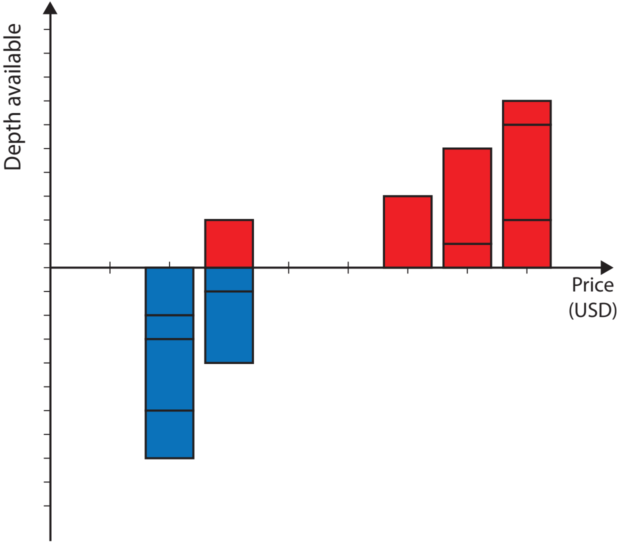
Depth available
Price
(USD)

Price Formation in a Centralized Limit Order Book
A Dynamic Example

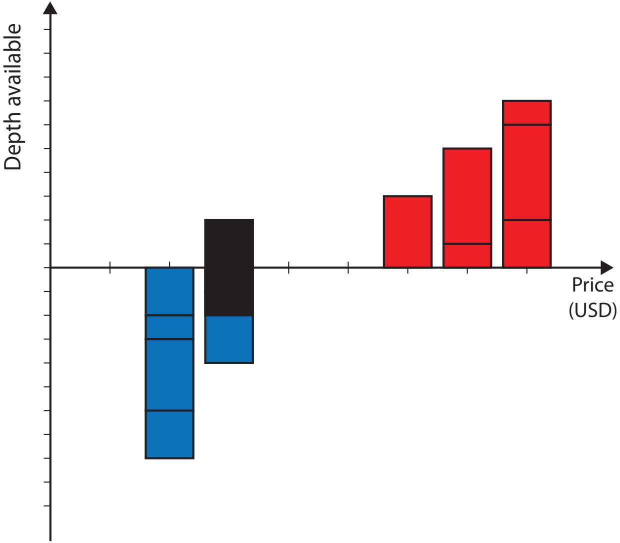

Price Formation in a Centralized Limit Order Book

A Dynamic Example

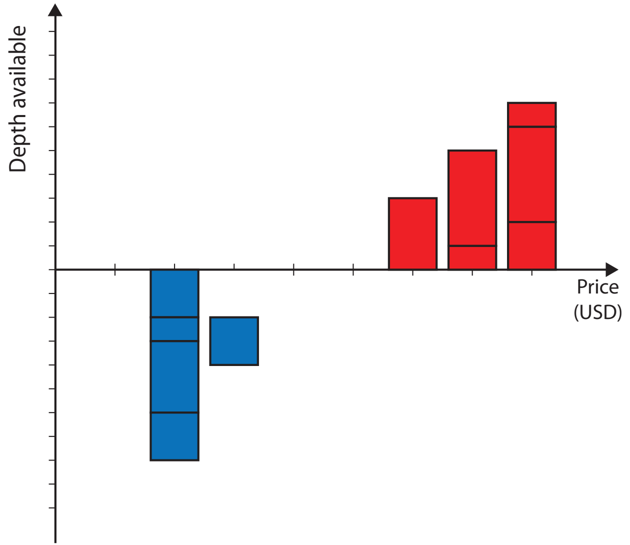

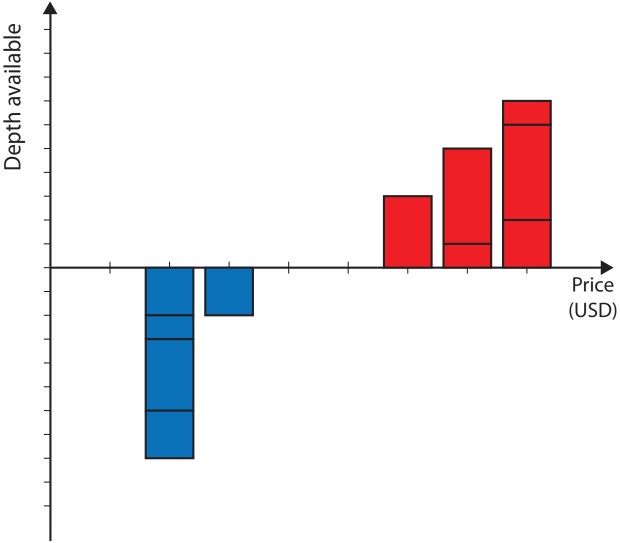

Depth available
Price
(USD)

Price Formation in a Centralized Limit Order Book

A Dynamic Example

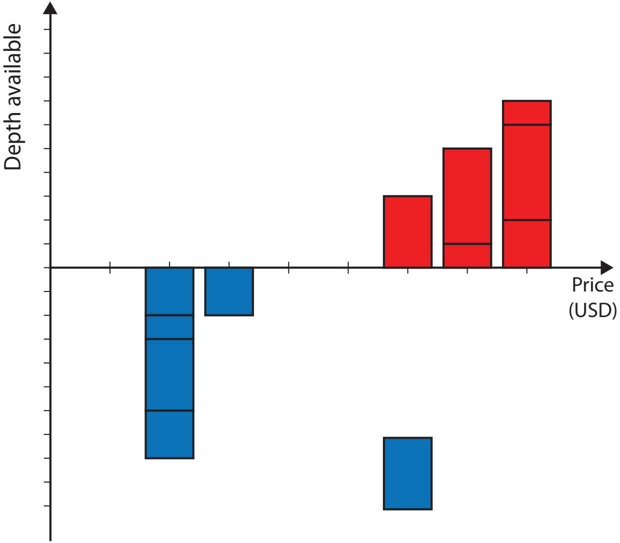

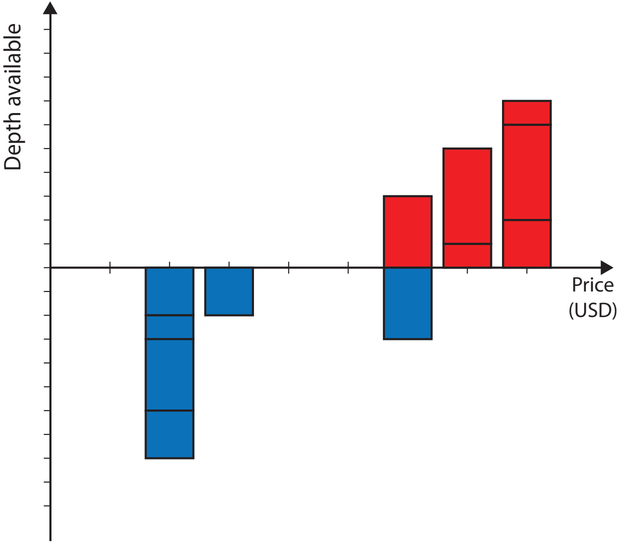

Depth available
Price (USD)

Price Formation in a Centralized Limit Order Book
A Dynamic Example

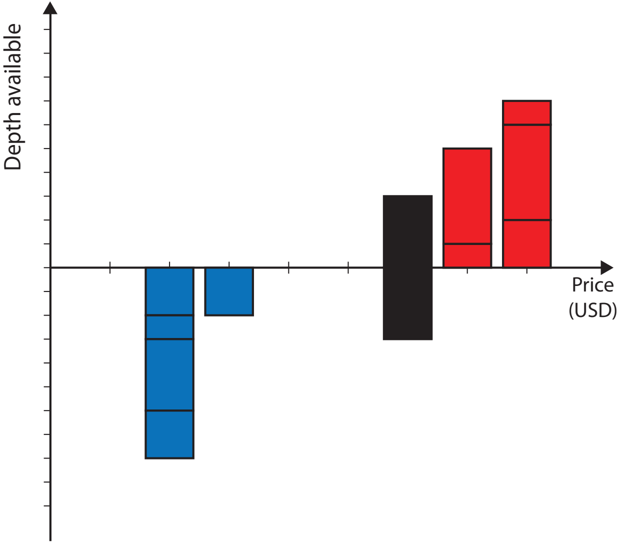

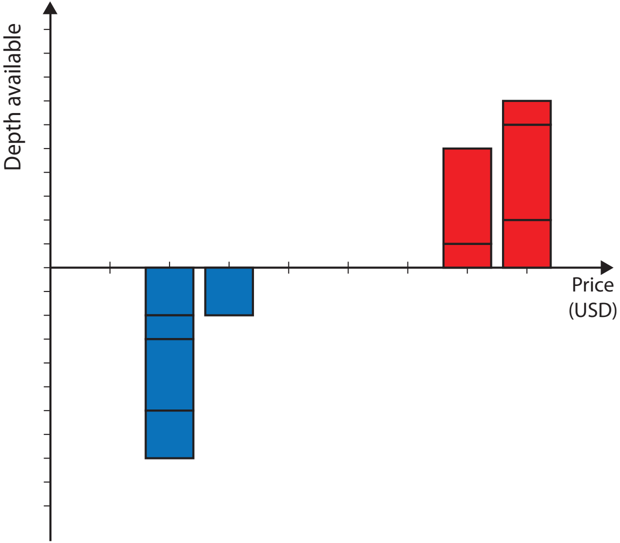
Depth available
Price
(USD)

Price Formation in a Centralized Limit Order Book
A Dynamic Example

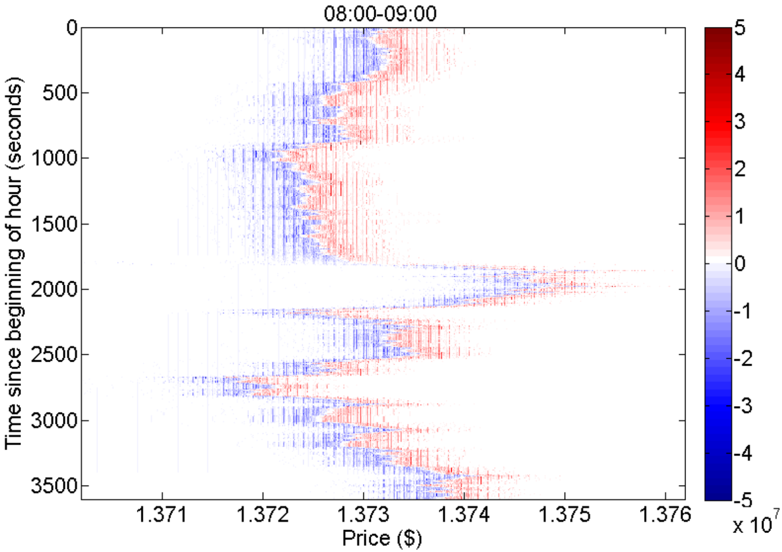

Order Anonymity

- Activity in an LOB is anonymous
- Trading platforms do not disseminate information regarding order ownership
- Institutions do not know which institutions own which limit orders
- Therefore, institutions do not know the identity of their counterparty when entering into a trade

Question

Given that any such trade incorporates counterparty risk, how should markets be organized? Does there exist a way to address counterparty risk, while still facilitating trade via an LOB?

Central Counterparties

The overwhelming majority of discussion related to this question has advocated the use of a *central counterparty (CCP)* to *novate* trades.

- Ben Bernanke: CCPs are "vital" for the healthy functioning of a financial market.
- World Federation for Exchanges: "CCP risk management capacities for mitigating and managing counterparty risk are essential contributions to the safety of the global financial marketplace"
- CAPCO: "CCPs play a critical role in post-trade management"

Central Counterparties

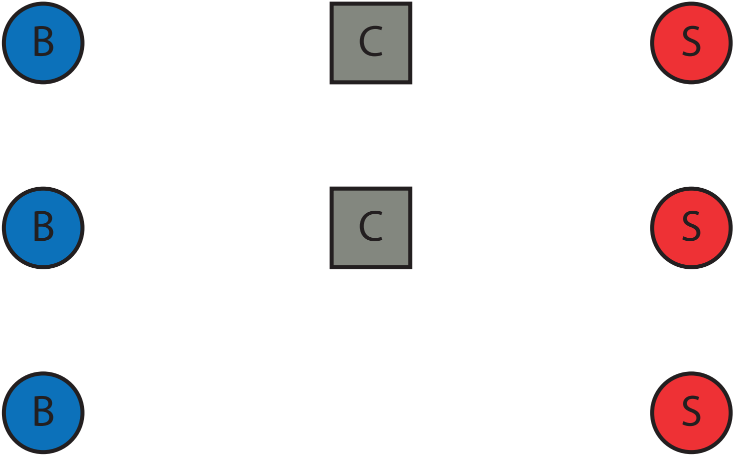

Central Counterparties

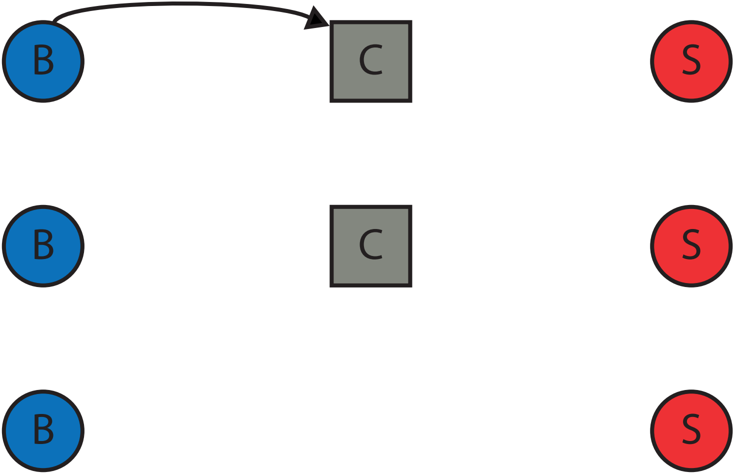

Central Counterparties

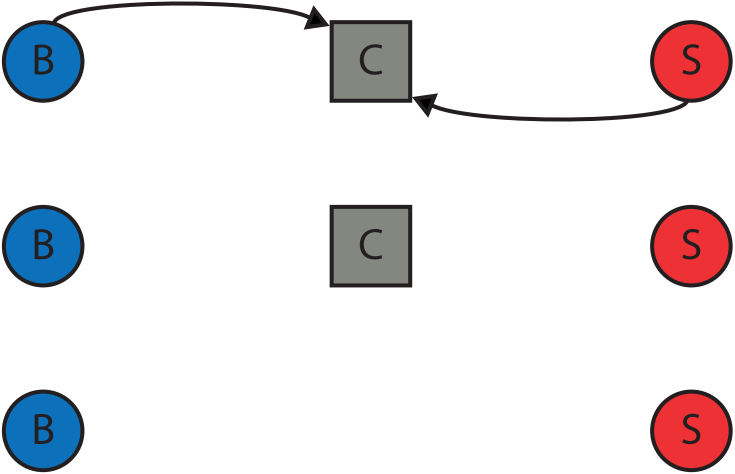

Central Counterparties

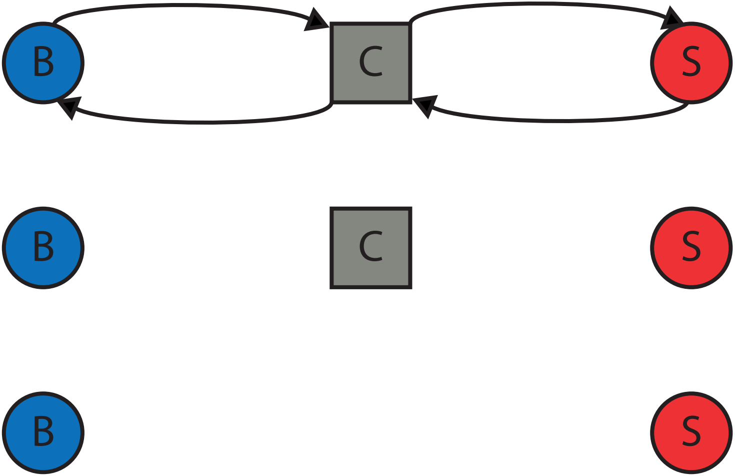

Central Counterparties

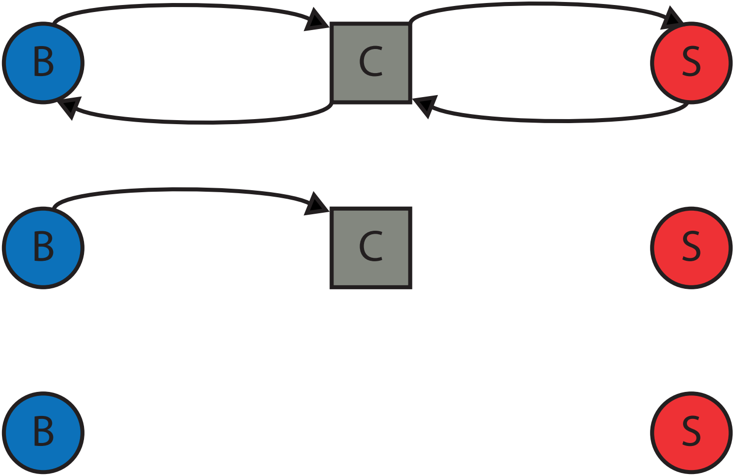

Central Counterparties

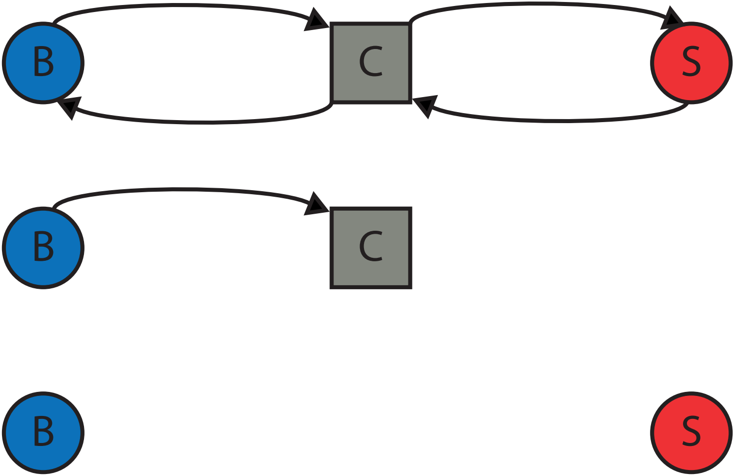

Central Counterparties

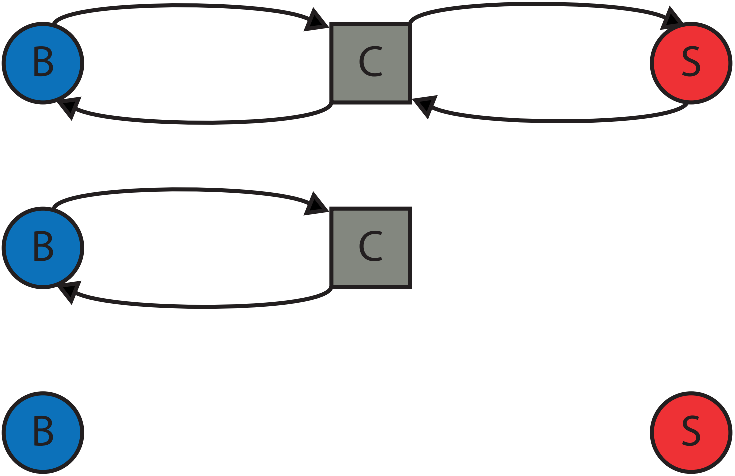

Central Counterparties

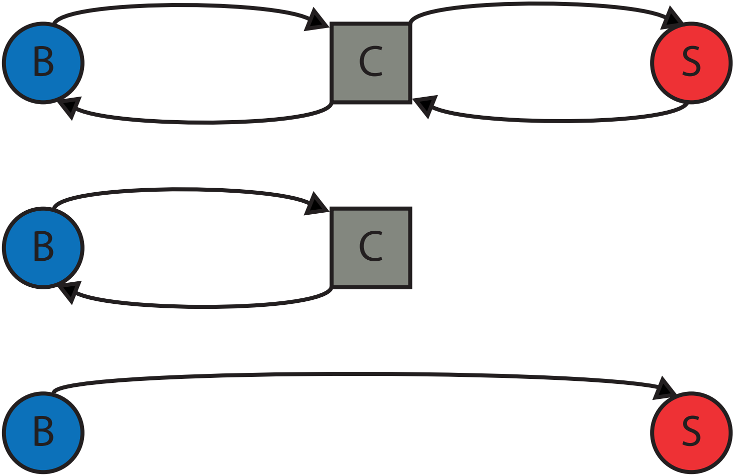

Central Counterparties

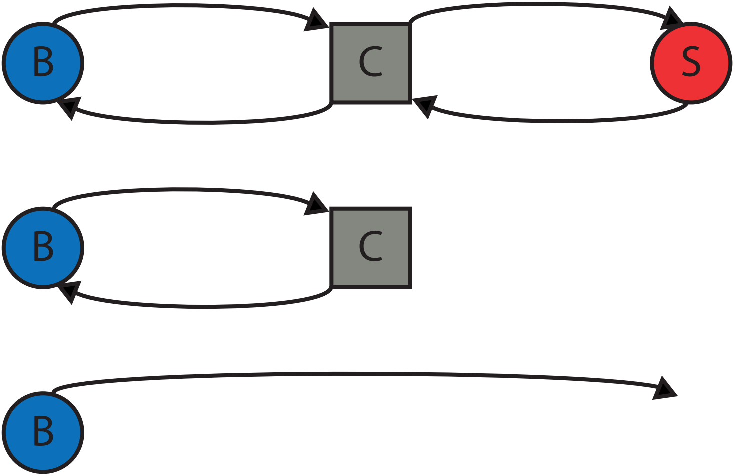

Problems with Central Counterparties

Despite the strong support for this mechanism, the use of a CCP has several important weaknesses.

- Single point of failure
- High margin requirements drain liquidity
- Moral hazard

Problems with Central Counterparties

In recent decades, several high-profile CCPs have indeed failed.

- Caisse de Liquidation in Paris in 1973
- Kuala Lumpur Commodity Clearing House in 1983
- Hong Kong Futures Guarantee Corporation in 1987
-

These failures suggest that CCPs are not a panacea for the issue of counterparty risk.

Problems with Central Counterparties

In recent decades, several high-profile CCPs have indeed failed.

- Caisse de Liquidation in Paris in 1973
- Kuala Lumpur Commodity Clearing House in 1983
- Hong Kong Futures Guarantee Corporation in 1987
- Alpari in 2015

These failures suggest that CCPs are not a panacea for the issue of counterparty risk.

Limit Order Books

Quasi-Centralized

Limit Order Books

Quasi-Centralized LOBs

Let $\Theta = \{\theta_1, \theta_2, \dots\}$ denote the set of institutions that trade a given asset on a given platform.

In a *quasi-centralized LOB (QCLOB)*, each institution θ_i assigns a *credit limit* $c_{(i,j)}$ to each other institution $\theta_j \neq \theta_i$.

Let $v_{(i,j)}$ denote the total value of outstanding trades between θ_i and θ_j.

Definition

For a pair of institutions $\theta_i, \theta_j \in \Theta$,

$$\theta_i \text{ accepts } \theta_j \Leftrightarrow v_{(i,j)} < c_{(i,j)}. \quad (1)$$

Otherwise, θ_i blocks θ_j.

Quasi-Centralized LOBs

If θ_i accepts θ_j, we write $\theta_i \rightarrow \theta_j$. If θ_i blocks θ_j, we write $\theta_i \nrightarrow \theta_j$.

If $\theta_i \rightarrow \theta_j$ AND $\theta_j \rightarrow \theta_i$, then θ_i and θ_j are *trading partners.*

If θ_i and θ_j are trading partners, we write $\theta_i \leftrightarrow \theta_j$. Otherwise, we write $\theta_i \nleftrightarrow \theta_j$.

Network Representation

For the set of institutions $\Theta = \{\theta_1, \theta_2, \theta_3\}$, if θ_2 blocks θ_3, θ_3 accepts θ_2, and θ_1 is trading partners with both θ_2 and θ_3, we can represent these trading partnerships as:

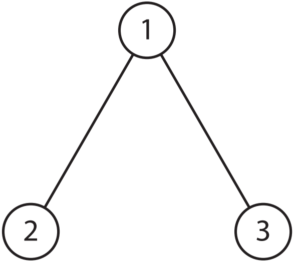

BTA LOBs

In a QCLOB, if $\theta_i \nleftrightarrow \theta_j$, then θ_i and θ_j:

- Are not able to trade with each other
- Are not able to see each other's active orders in the LOB.

Definition

Let x_i denote an order that is owned by institution $\theta_i \in \Theta$. In a QCLOB, θ_i's local LOB, denoted $\mathcal{L}_i(t)$, is the subset of orders in $\mathcal{L}(t)$ that are owned by θ_i's trading partners:

$$\mathcal{L}_i(t) := \left\{ x_j \in \mathcal{L}(t) \middle| \, \theta_i \leftrightarrow \theta_j \right\}.$$

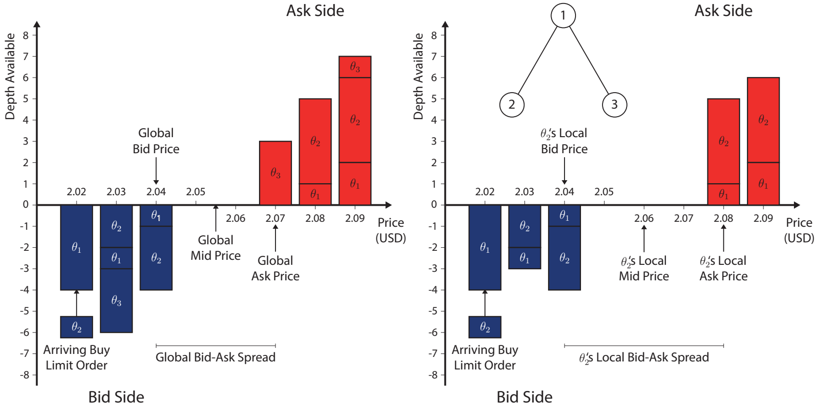

Figure: (Left) The global LOB $\mathcal{L}(t)$ and (right) θ_2's local LOB $\mathcal{L}_2(t)$

A Question

Question: In a QCLOB, how do institutions with few trading partners assess market state?

The Hotspot FX Trading Screen

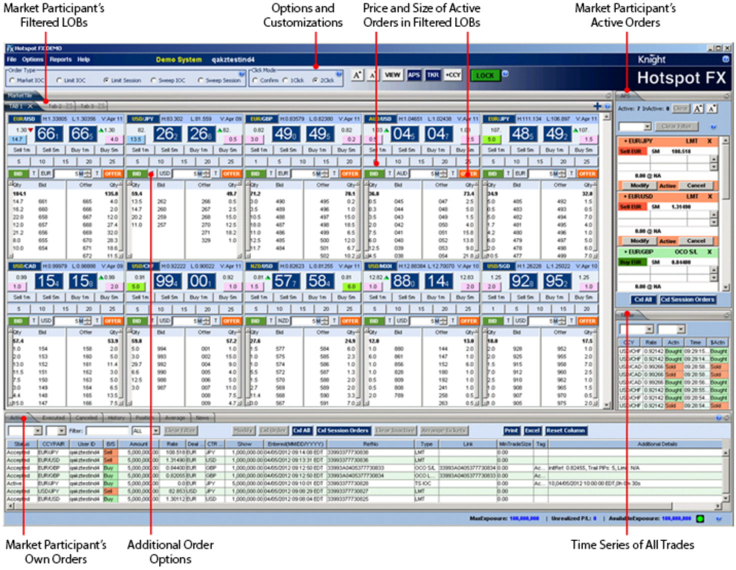

The Answer

Question: In a QCLOB, how do institutions with few trading partners assess market state?

Answer: Each institution $\theta_i \in \Theta$ views not only their local LOB $\mathcal{L}_i(t)$, but also a time series of *all* trades.

Key Research Questions

Except for the extreme case in which all institutions are trading partners with all others, different institutions in the same QCLOB have access to different trading opportunities at the same time.

Key considerations for this presentation:

- Which statistical regularities exist in such markets?
- Do these statistical regularities differ from those in centralized LOBs?
- What might they tell us about the ways in which people trade?
- Do arbitrage opportunities arise for some traders?
- How does price formation vary from that in centralized LOBs?

The Hotspot FX Data

- 3rd largest electronic trading platform in the FX spot market
- Approximately 10% of the market's total global volume
- More than 60 different currency pairs; we study:
  - EUR/USD (solid lines)
  - GBP/USD (dashed lines)
  - EUR/GBP (dotted-dashed lines)
- Data describes all order flow during "peak trading hours" of 08:00–17:00 GMT on 30 trading days during May–June 2010
- Can reconstruct the global LOB $\mathcal{L}(t)$ (at all levels) at any time during this period
- However: No information about the trade-partnership network or any $\mathcal{L}_i(t)$

Total Size of Arriving Orders

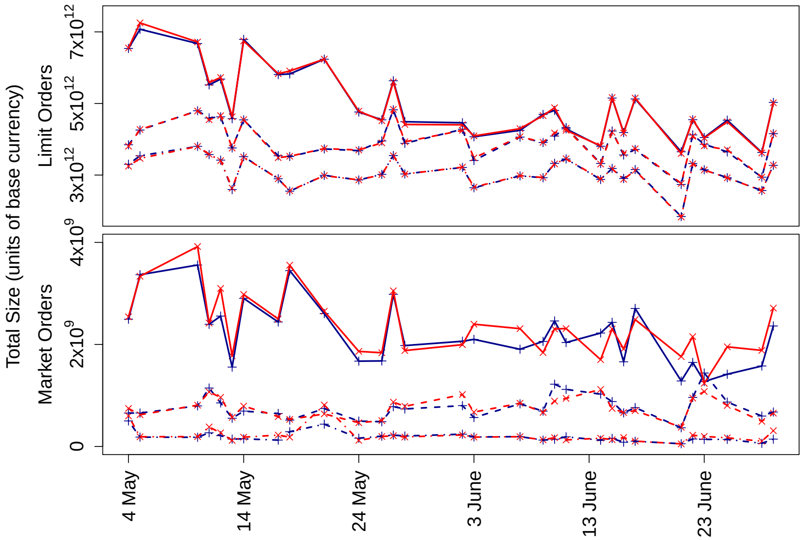

Relative Pricing

In LOBs, knowing the price of an order is of little use because typical trade prices change.

Relative pricing provides a *context* to prices:

Definition

The bid-side relative price $\delta_b(p)$ *is defined as* $\delta_b = b(t) - p$.

Definition

The ask-side relative price $\delta_a(p)$ *is defined as* $\delta_a = p - a(t)$.

Relative pricing has led to the discovery of several robust statistical properties in a wide range of centralized LOBs

- Power-law decay of arrival rates by relative price
- Power-law decay of cancellation rates by relative price
- Hump-shaped mean LOB state

Limit Order Arrivals

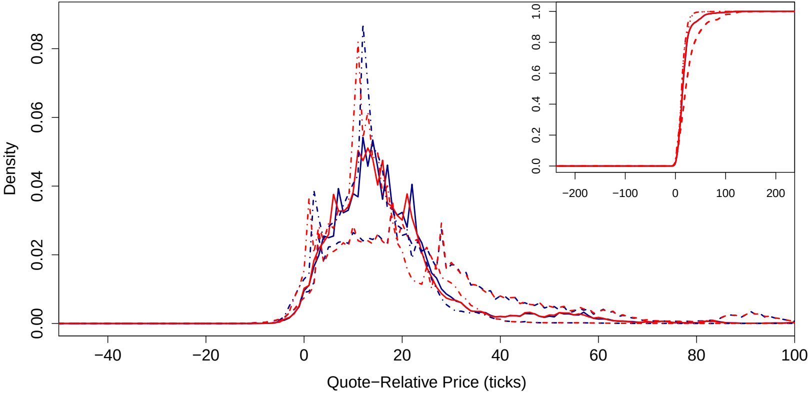

Cancellations

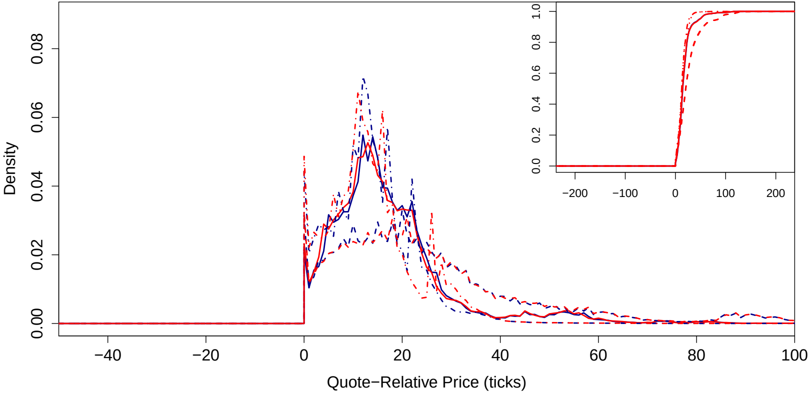

Mean Relative Depth Profile

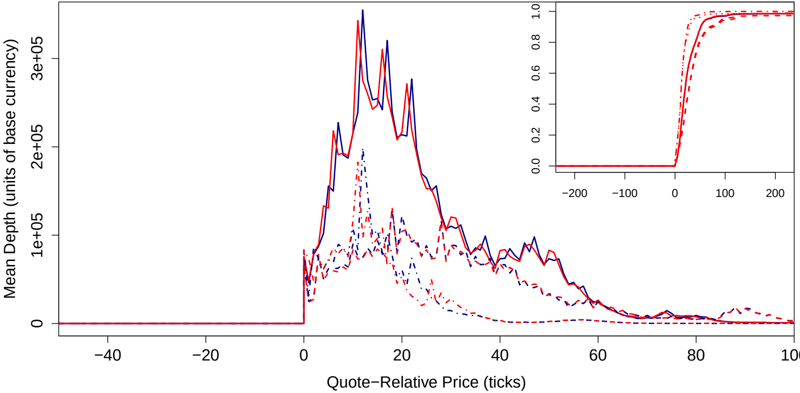

Cancellation Ratio

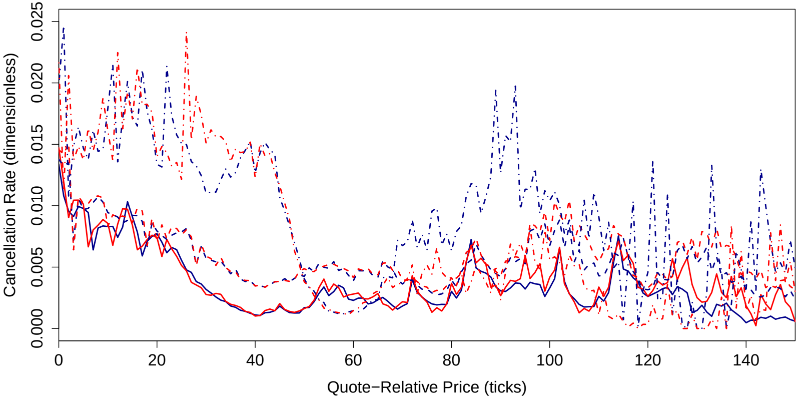

The Trouble with Relative Pricing

Clearly:

- Each market participant has different motives for trading
- The behaviour of any single market participant is erratic

Yet:

- Measured in a reference frame that aggregates activity from a large number of market participants, robust statistical properties should emerge from the ensemble
- e.g., The ideal gas law

What went wrong when studying the QCLOB data?

The Trouble with Relative Pricing

Clearly:

- Each market participant has different motives for trading
- The behaviour of any single market participant is erratic

Yet:

- Measured in a reference frame that aggregates activity from a large number of market participants, robust statistical properties should emerge from the ensemble
- e.g., The ideal gas law

What went wrong when studying the QCLOB data?

The Trouble with Relative Pricing

- In a QCLOB, no one knows the values of $b(t)$ or $a(t)$.
- Even institutions that observe orders at $b(t)$ or $a(t)$ don't know that there isn't another active order at a better price.
- If each institution makes decisions based on $b_i(t)$ and $a_i(t)$, superimposing market activity relative to $b(t)$ and $a(t)$ does not align their activity

An Alternative: Trade Relative Pricing

Recall that QCLOBs disseminate a time series of the prices of all previous trades.

Definition

At a given time t, let $B(t)$ and $A(t)$ denote the prices at which the most recent bid-side (respectively, ask-side) trade occurred.

For a given price p and a given time t:

Definition

The bid-side trade-relative price $\Delta_b(p)$ is defined as $\Delta_b = B(t) - p$.

Definition

The ask-side trade-relative price $\Delta_a(p)$ is defined as $\Delta_a = p - A(t)$.

Limit Order Arrivals

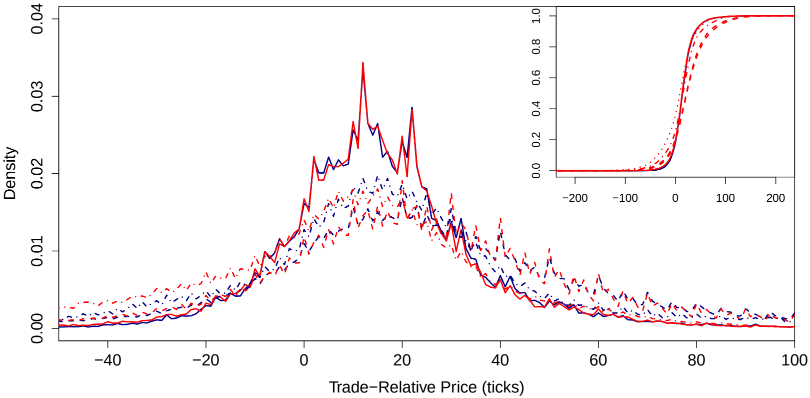

Cancellations

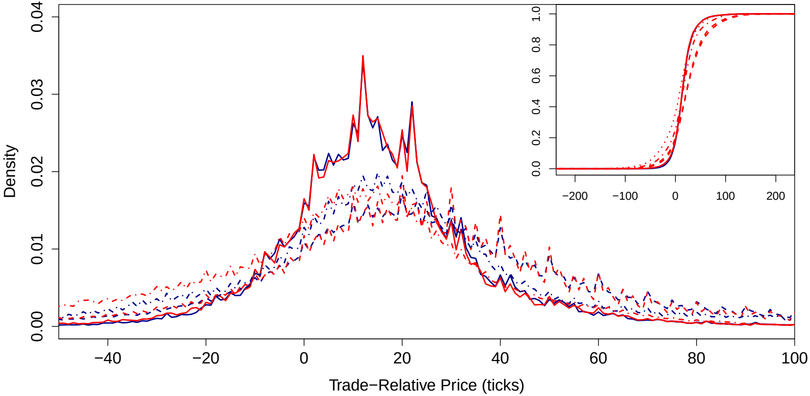

Mean Relative Depth Profile

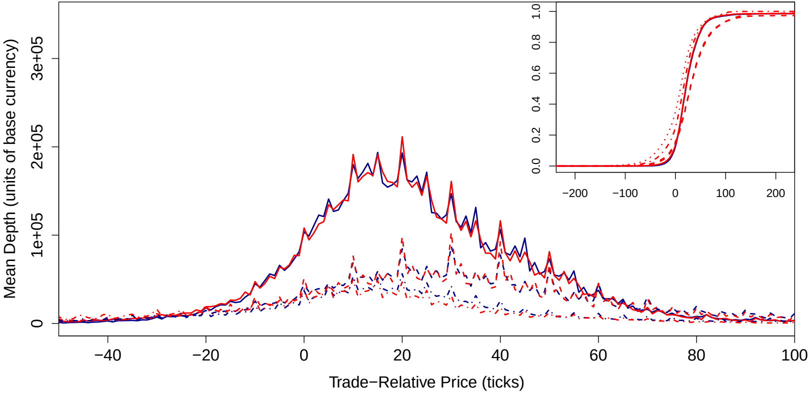

Cancellation Ratio

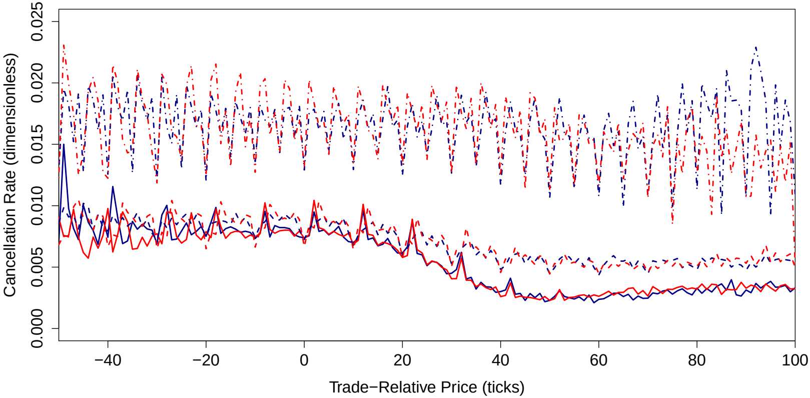

QCLOBs: A Recap

- Allowing institutions to specify credit limits gives them explicit control over their choice of trading counterparties
- Different institutions in a QCLOB have access to different trading opportunities at the same time
- Quote-relative pricing does not provide a useful framework for studying QCLOBs
- Trade-relative pricing uncovers robust statistical properties in the Hotspot FX data

Question

How does the quasi-centralized nature of trade impact price formation?

The Bypassing Price of a Trade

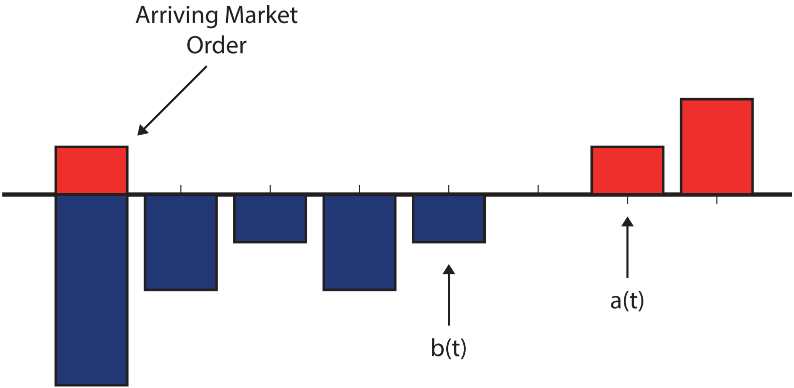

The Bypassing Price of a Trade

Consider a sequence of N^A buyer-initiated trades that occur at times $t_1^A, t_2^A, \ldots, t_{N^A}^A$.

For each such trade k, let:

- The *traded-ask price* p_k^A denote the price at which the trade occurs
- The *global-ask price* p_k^a denote the lowest price among active sell orders immediately before t_k^A
- The *bypassed-ask price*

$$p_k^{\alpha} = p_k^A - p_k^a \tag{2}$$

  denote the difference between the traded-ask price and the global-ask price

The Bypassing Price of a Trade: Summary Statistics

Statistic	EUR/USD	GBP/USD	EUR/GBP
Minimum (ticks)	0.00	0.00	0.00
Maximum (ticks)	363.00	142.00	73.00
Median (ticks)	0.00 (< 0.01)	0.00 (< 0.01)	0.00 (< 0.01)
Mean (ticks)	2.53 (0.09)	3.18 (0.12)	1.90 (0.18)
St. Dev. (ticks)	5.69 (0.97)	6.65 (0.21)	4.11 (0.56)
Percentage Zeros	55.80% (1.43%)	61.17% (1.79%)	66.35% (1.46%)

Table: Summary statistics for EUR/USD, GBP/USD, and EUR/GBP bypassing prices.

The Bypassing Price of a Trade: Survivor Functions

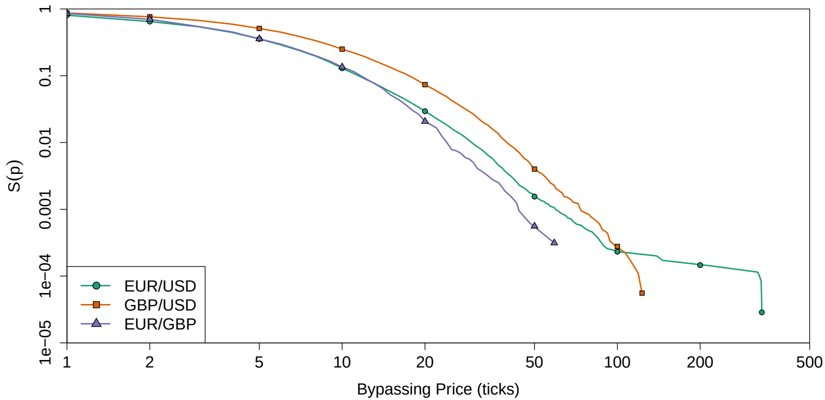

Given a pair of successive trades, the *change in traded-ask price* is

$$\Delta p_k^A = p_k^A - p_{k-1}^A, \tag{3}$$

the *change in global-ask price* is

$$\Delta p_k^a = p_k^a - p_{k-1}^a, \tag{4}$$

and the *change in bypassed-ask price* is

$$\Delta p_k^\alpha = p_k^\alpha - p_{k-1}^\alpha. \tag{5}$$

Price Changes: Survivor Functions

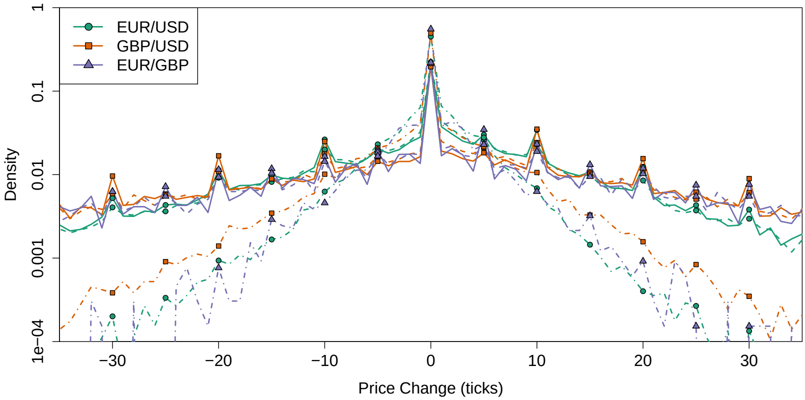

(Solid lines denote change in trade price, dashed lines denote change in quote price, dotted-dashed lines denote change in bypassing price)

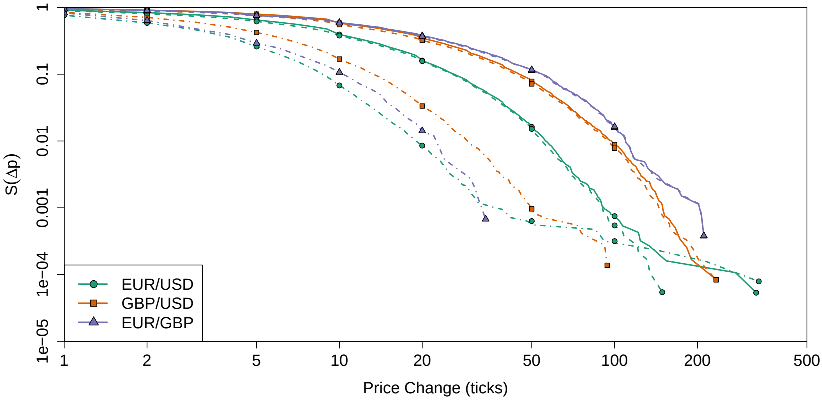

(Solid lines denote change in trade price, dashed lines denote change in quote price, dotted-dashed lines denote change in bypassing price)

Decomposing Price Changes

Observe that

$$\begin{aligned}\Delta p_k^a + \Delta p_k^\alpha &= p_k^a - p_{k-1}^a + p_k^\alpha - p_{k-1}^\alpha \\ &= p_k^a - p_{k-1}^a + (p_k^A - p_k^a) - (p_{k-1}^A - p_{k-1}^a) \\ &= p_k^A - p_{k-1}^A \\ &= \Delta p_k^A.\end{aligned}$$

Therefore, the change in traded-ask price between any consecutive pair of buyer-initiated trades can be decomposed into the sum of corresponding changes in global-ask price and bypassed-ask price,

$$\Delta p_k^A = \Delta p_k^a + \Delta p_k^\alpha. \quad (6)$$

A similar result holds for seller-initiated trades.

Decomposing Price Changes

To quantify the *relative importance of* Δp_k^a and Δp_k^α for Δp_k^A, we first perform separate ordinary least-squares (OLS) regressions

$$\hat{y}_k^1 = \mu_0 + \mu_1 \Delta p_k^a, \qquad k = 1, 2, \ldots, (N^A - 1), \tag{7}$$

$$\hat{y}_k^2 = \nu_0 + \nu_1 \Delta p_k^\alpha, \qquad k = 1, 2, \ldots, (N^A - 1), \tag{8}$$

where $\hat{y}_k^1$ and $\hat{y}_k^2$ are the OLS predictions of Δp_k^a.

Decomposing Price Changes

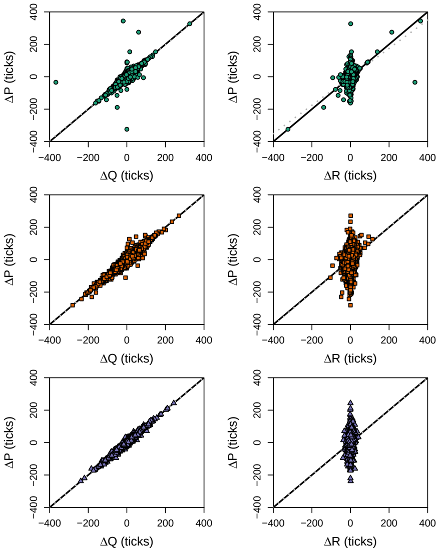

Relative Importance

In a multiple regression, the *relative importance* of an input variable measures its relative contribution to some specified aspect of the regression output.

The *dispersion importance* measures the proporion of the explained variance (i.e., R^2) of the output variable that is attributable to each input variable.

If a variable has a high dispersion importance, then its inclusion contributes a large fraction to the multiple regression's total R^2.

Decomposing Price Changes

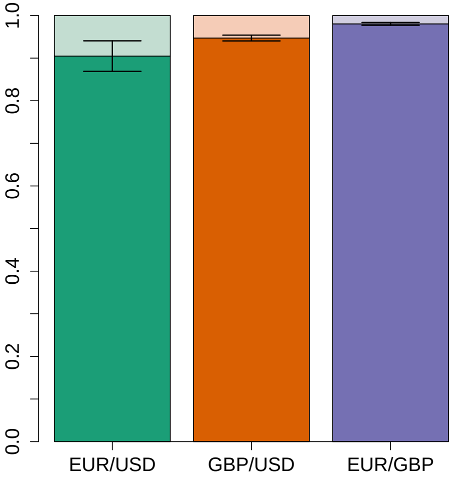

Price Formation in QCLOBs: A Recap

- In a QCLOB, different traders have access to different trading opportunities at the same time
- The bypassing price of a market order measures the cost to its owner for lacking the necessary trading partnerships to trade at the best quotes
- The bypassing prices of trades on Hotspot FX vary over about 3 orders of magnitude
- Decomposing price changes enables explicit comparisons of changes in quote price and changes in bypassing price
- The primary driver for changes in trade price is changes in quote price, not changes in bypassing price.

Conclusions

- QCLOBs provide institutions with all the key benefits of LOB trading, while also providing them with explicit choice of their trading counterparties
- Not only useful for excluding small, unknown counterparties!
- Quote-relative pricing less useful for studying QCLOBs than reported elsewhere
- Trade-relative pricing provides much clearer insight into order flow and LOB state
- Although a handful of trades have large bypassing prices, the vast majority of trades occur near to the best quotes
- Policy recommendation: Test the use of QCLOBs in other markets

And one last thing. . .

The use of QCLOBs does NOT exclude the presence (or absence!) of a trusted third-party for settling trades. . .

And one last thing...

The use of QCLOBs does NOT exclude the presence (or absence!) of a trusted third-party for settling trades...

...but this is a story for another day ☺

References

[J.P. Bouchaud, M. Mézard, and M. Potters, *"Statistical properties of stock order books: Empirical results and models,"* Quant. Finance, 2, (2002), pp. 251–256]

[R. Cont, *"Statistical modeling of high-frequency financial data,"* IEEE Signal Processing Magazine, 28, (2011), pp. 16–25]

[M. D. Gould, M. A. Porter, S. Williams, M. McDonald, D. J. Fenn, S. Howison, *"Limit order books,"* Quant. Finance, 13, (2013), pp. 1709–1742]

[M. D. Gould, M. A. Porter, S. Williams, M. McDonald, D. J. Fenn, S. Howison, *"Quasi-centralized limit order books,"* Working Paper, 2015]

[M. D. Gould, M. A. Porter, S. Williams, M. McDonald, D. J. Fenn, S. Howison, *"Price formation in quasi-centralized liquidity pools,"* Working Paper, 2015]

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

[M. Potters and J.P. Bouchaud, *"More statistical properties of order books and price impact,"* Physica A, 324, (2003), pp. 133–140]