

Quantifying the High-Frequency Trading “Arms Race”: A Simple New Methodology and Estimates

[ONLINE APPENDICES]

Matteo Aquilina, Eric Budish, and Peter O’Neill

Table of Contents

A Supporting Details for Section 3.4: Computing the Information Horizon	2
B Additional Results for Section 4: Main Results	5
C Supporting Details for Section 5.1: Sensitivity Analysis	20
D Supporting Details for Section 5.2: Additional Robustness Checks	32
E Additional Results for Section 6: Magnitudes	38
F Theory Appendix	41

A Supporting Details for Section 3.4: Computing the Information Horizon

As described in Section 3.4 of the main text, there are three elements of our Information Horizon calculation:

1. Actual Observed Latency: M1 Inbound \rightarrow M1 Outbound
2. Minimum Observed Reaction Time: M1 Outbound \rightarrow M2 Inbound
3. Upper bound on maximum possible information horizon

where M1 refers to the first message in a potential race and M2 the second message.

We can compute the Actual Observed Latency: M1 Inbound \rightarrow M1 Outbound directly in our data, for each inbound message. This is obtained by taking the difference between the inbound message’s timestamp and its outbound message’s timestamp. The median response time is 157 microseconds, and there is considerable variation: the 10th percentile is 108 microseconds and the 90th percentile is 303 microseconds.¹


To compute the Minimum Observed Reaction Time: M1 Outbound \rightarrow M2 Inbound, we start by finding all instances of the specific sequence of events where M1 outbound is a new limit order that adds liquidity at some price level, and M2 inbound is an aggressive order (i.e., take) from a different UserID at the same price level. In this sequence of events, M2 may be responding to the new liquidity at the price level by taking it. Clearly, sometimes this sequence of events will happen by chance, but sometimes this sequence of events will happen because M2 is responding to M1.² Figure A.1 reports the distribution of the difference in time between these two events.

As can be seen, this distribution spikes upwards a bit to the right of 0. We interpret the beginning of this spike as the minimum amount of time it takes the fastest market participants to respond to such an M1 with such an M2, as measured from the outbound time stamp to the inbound time stamp. Note that it need not be the case that the market participant is responding literally to the outbound message sent to the participant who sent M1; rather, the market participant is likely responding to their own receipt of information about the state of the order book from the LSE’s proprietary data feed, sent through the message server as depicted in Figure 2.1 in the main text. Using the simple statistical criterion of looking for the start of the spike by asking what is the first microsecond at which the density is more than 5 standard deviations above the distribution in the 100 microseconds leading up to time 0, we determine that the spike starts at 29 microseconds.

¹These figures are based on the M1 Inbound \rightarrow M1 Outbound response time over all messages that are the first message in a race.

²See also Dobrev and Schaumburg (2018) who study high-frequency cross-market linkages between US treasury and equity markets. Whereas our goal is to compute the information horizon, such that we can be sure that a particular message M2 is not responding to a particular message M1 within the same market, Dobrev and Schaumburg’s (2018) method can be used to study the lead-lag relationship across markets (e.g., treasuries and equities), where sometimes an event in the first market triggers a reaction in the second market, sometimes the reverse occurs where an event in the second market triggers a reaction in the first market, and sometimes events occur in both markets simultaneously, i.e., within what we term the information horizon.

Figure A.1: **Distribution of Time between M1 Outbound New Limit Order \rightarrow M2 Inbound Takes Liquidity**




Notes: Over all regular-hour messages from four high-volume symbols, BP, GLEN, HSBA, VOD, we obtain all cases where some outbound message confirms a new order added to the book and subsequently gets filled at least in part. We then obtain the first outbound message that is an execution against this new order, obtain the inbound message associated with this outbound execution message, and compute the difference in the message timestamp between the first order’s (M1) outbound message and the second order’s (M2) inbound message. Note that this difference can be negative if M2’s inbound is sent by the participant before M1’s outbound is sent by the outbound gateway. The distribution depicted is a microsecond-binned histogram truncated at -500 microseconds and +500 microseconds. As described in the text, we compute the start of the spike (29 microseconds) by computing the mean and standard deviation of the distribution in the period -100 microseconds to 0 microseconds, and then finding the first microsecond after 0 that is at least 5 standard deviations above this pre-0 mean.

We also examined the case where M1 is a partial fill, and M2 is a successful cancel (Figure A.2). In this case, the participant might be responding to their own privately-received message—so we might expect this to be faster than what we saw above for the Add-Take sequence. Here, the spike starts at around 17 microseconds. An interpretation is that the 17 microseconds is the minimum response time to a privately-observed outbound message, and the additional 12 microseconds is the minimum difference in latency between a private message sent to a particular market participant and the LSE’s broadly disseminated proprietary data feed.³ The sum of these two figures (i.e., the 29 microseconds) is more appropriate for computing the information horizon in a race, but the response time to a privately-observed book update may be of independent interest so we report the distribution here for completeness.

³A similar difference between the speed with which private messages are received versus book updates from proprietary data feeds has been a recurring source of controversy at the Chicago Mercantile Exchange. See Patterson, Strasburg and Plevin (2013) and Osipovich (2018).

Figure A.2: **Distribution of Time between M1 Outbound partial fill → M2 Inbound Successful Cancel**



Notes: Over all regular-hour messages from four high-volume symbols, BP, GLEN, HSBA, VOD, we obtain all cases where some outbound message is a partial fill and a subsequent outbound message is a successful cancel. We then obtain the inbound cancel request message associated with the outbound cancel success message, and compute the difference in the message timestamp between the partial fill outbound message (M1) and the cancel request inbound message (M2). Note that this difference can be negative if M2's inbound is sent by the participant before M1's outbound is sent by the outbound gateway. The distribution depicted is a microsecond-binned histogram truncated at -500 microseconds and +500 microseconds. As described in the text, we compute the start of the spike by computing the mean and standard deviation of the distribution in the period -100 microseconds to 0 microseconds, and then finding the first microsecond after 0 that is at least 5 standard deviations above this pre-0 mean.

Last, the upper bound on the information horizon that we utilize, 500 microseconds, was determined in consultation with supervisors at the Financial Conduct Authority. This was based on the discussions they had with fast market participants on their reaction times, differences in the speeds of competing microwave connectivity providers, the variance in arrival times across long distances (such as Chicago to London), the geographical distance between the LSE's data center and other UK exchanges' data centers, and the judgment of supervisory experts to establish an amount of time short enough for our assumption that M2 is not reacting to M1 to be reasonable. This 500 microsecond truncation of the information horizon binds in just under 4% of cases.

B Additional Results for Section 4

This appendix presents additional specifications for results presented throughout Section 4 of the main text. The material is presented in order of the corresponding parts of the main text.

B.1 Additional Results for Section 4.1

Symbol-Date Version of Table 4.1

Table 4.1 in the main text reports the number of races per day at the symbol level averaged across all dates (Panel A), and at the date level summed across all symbols (Panel B). The following table presents the number of races at the symbol-date level, i.e., without aggregating across either symbols or dates.

Table B.1: Number of Races Per Day Across Symbol-Dates

Description	Mean	sd	Min	Pct01	Pct10	Pct25	Median	Pct75	Pct90	Pct99	Max
FTSE 100	537.24	542.96	29	73	152	215	346	629	1,194	2,635	7,014
FTSE 250	70.05	103.26	0	0	0	2	35	97	182	477	1,392
Full Sample	206.03	372.02	0	0	0	11	81	231	513	1,919	7,014

Notes: Please see the notes for Table 4.1 in the main text. The table in the main text reports the distribution of the number of races at the symbol level and at the date level. This appendix table reports the distribution of the number of races detected at the symbol-date level.

Race Duration

Figure 4.1 in the main text plots the distribution of the duration of races, as measured from the time that elapses between the first success message in the race (S1) and the first fail message in the race (F1). The following table presents the percentiles of this distribution in table form, separately for FTSE 100, FTSE 250, and full sample.

Table B.2: Race Duration

Time from S1 to F1 (microseconds)

Description	Mean	sd	Pct01	Pct10	Pct25	Median	Pct75	Pct90	Pct99
FTSE 100	80.81	92.14	-9.00	3.70	12.60	48.50	123.70	207.50	402.80
FTSE 250	71.85	80.84	-4.40	4.30	12.80	37.10	111.70	185.60	338.00
Full Sample	78.65	89.63	-7.90	3.80	12.70	45.60	120.90	201.90	390.20

Notes: Please see the notes for Figure 4.1 in the main text. For each race detected by our baseline method, we compute the difference in message timestamps between the first inbound message in the race that is a success and the first inbound message in the race that is a fail. Denote these messages S1 and F1, respectively. This table reports the distribution of F1's timestamp minus S1's timestamp in microseconds, that is, by how long did the first successful message in the race beat the first failed message.

Total Time in Races

In the text of Section 4.1 we report the distribution of the number of races per day (Table 4.1) and the distribution of the duration of races (Figure 4.1). In this appendix table we report the distribution of the total time in races per day. This is reported in seconds per day at the symbol-date level.

Table B.3: **Total Time in Races Across Symbol-Dates**

Description	Mean	sd	Min	Pct01	Pct10	Pct25	Median	Pct75	Pct90	Pct99	Max
FTSE 100	0.044	0.047	0.002	0.006	0.012	0.017	0.028	0.052	0.096	0.235	0.739
FTSE 250	0.005	0.008	0.000	0.000	0.000	0.000	0.002	0.007	0.013	0.036	0.093
Full Sample	0.016	0.032	0.000	0.000	0.000	0.001	0.006	0.018	0.042	0.153	0.739

Notes: For each symbol-date in our dataset, we sum all race durations as defined in Figure 4.1 and Table B.2 and report the distribution. For example, the table indicates that in the mean FTSE 100 symbol-date, the sum of the duration of all races is 0.044 seconds.

Symbol-level Version of Table 4.2

Table 4.2 in the main text reports the percentage of volume and trades in races at the date level, i.e., averaged across all symbols in the FTSE 100, FTSE 250, and full sample respectively. In this appendix table we report the percentage of volume and trades in races at the symbol level averaged across all dates.

Table B.4: **Volume and Trades in Races**

Panel A: Percentage of volume (value-weighted) in races across symbols

Description	Mean	sd	Pct01	Pct10	Pct25	Median	Pct75	Pct90	Pct99
FTSE 100	23.48	4.90	13.08	17.84	20.07	23.30	26.34	30.62	33.75
FTSE 250	11.33	8.48	0.00	0.61	1.99	12.69	18.48	22.07	27.30
Full Sample	14.86	9.40	0.00	1.11	5.79	17.20	22.02	25.78	33.06

Panel B: Percentage of number of trades in races across symbols

Description	Mean	sd	Pct01	Pct10	Pct25	Median	Pct75	Pct90	Pct99
FTSE 100	22.19	4.56	12.54	16.69	19.58	21.79	24.78	28.44	32.09
FTSE 250	11.31	8.37	0.00	0.55	2.00	13.21	18.32	21.63	27.31
Full Sample	14.48	8.95	0.00	0.87	6.05	16.70	21.36	24.67	31.16

Notes: Please see the notes for Table 4.2 in the main text. Table 4.2 reports the distribution of percentage of volume and trades in races at the date level. This appendix table reports the same distribution but at the symbol level.

B.2 Additional Results for Section 4.2

Additional Data on Messages Per Race

Table 4.3 in the main text reports the number of participants, takes, and cancels in the T microseconds after the start of a race for values of T between 50us and 1ms. In this appendix table we break out the take messages into two types: immediate-or-cancels (IOCs) and limit orders. Recall that in many of the sensitivity analyses discussed in Section 5 of the main text we only allow for IOC take messages to count towards the 1+ fails requirement for race detection. The table shows that about 90% of take messages sent in races are IOCs as opposed to plain-vanilla limit orders.

This appendix table also reports the total number of messages and total number of firms in races. The number of firms can be lower than the number of participants in case there are multiple active trading desks within the same firm in a race, and the number of participants can in turn be lower than the number of messages in case some participants send multiple messages in a race.

Table B.5: Number of IOC / Limit Takes and Number of Messages / Firms in Races

Panel A: Number of take (IOC) messages											
Description	Mean	sd	Min	Pct01	Pct10	Pct25	Median	Pct75	Pct90	Pct99	Max
IOC takes within 50 μ s	1.56	0.99	0	0	1	1	1	2	3	5	14
IOC takes within 100 μ s	1.80	1.10	0	0	1	1	2	2	3	5	15
IOC takes within 200 μ s	2.20	1.32	0	0	1	1	2	3	4	6	17
IOC takes within 500 μ s	2.81	1.73	0	0	1	2	2	4	5	8	29
IOC takes within 1000 μ s	3.07	2.00	0	0	1	2	3	4	6	10	40

Panel B: Number of take (limit) messages											
Description	Mean	sd	Min	Pct01	Pct10	Pct25	Median	Pct75	Pct90	Pct99	Max
Limit takes within 50 μ s	0.10	0.32	0	0	0	0	0	0	0	1	5
Limit takes within 100 μ s	0.13	0.39	0	0	0	0	0	0	1	2	6
Limit takes within 200 μ s	0.17	0.45	0	0	0	0	0	0	1	2	7
Limit takes within 500 μ s	0.25	0.60	0	0	0	0	0	0	1	3	11
Limit takes within 1000 μ s	0.37	0.82	0	0	0	0	0	0	1	4	17

Panel C: Number of messages											
Description	Mean	sd	Min	Pct01	Pct10	Pct25	Median	Pct75	Pct90	Pct99	Max
Messages within 50 μ s	1.83	0.93	1	1	1	1	2	2	3	5	14
Messages within 100 μ s	2.15	1.05	1	1	1	1	2	3	3	6	15
Messages within 200 μ s	2.67	1.23	1	1	2	2	2	3	4	7	17
Messages within 500 μ s	3.46	1.72	2	2	2	2	3	4	6	9	29
Messages within 1000 μ s	3.90	2.19	2	2	2	2	3	5	7	12	41


Panel D: Number of firms											
Description	Mean	sd	Min	Pct01	Pct10	Pct25	Median	Pct75	Pct90	Pct99	Max
Firms within 50 μ s	1.55	0.69	1	1	1	1	1	2	2	4	7
Firms within 100 μ s	1.77	0.76	1	1	1	1	2	2	3	4	8
Firms within 200 μ s	2.12	0.82	1	1	1	2	2	3	3	4	8
Firms within 500 μ s	2.60	1.01	1	1	2	2	2	3	4	6	10
Firms within 1000 μ s	2.82	1.19	1	1	2	2	3	3	4	6	12

Notes: Please see the notes for Table 4.3 in the main text and the description in the text above this table.

Additional Versions of Percentage of 1st Successful and Failed Messages by Firm

Figure 4.2 in the main text reports the percentage of 1st successful and 1st failed messages in races, by firm, over all races in the FTSE 100. The following two appendix figures report the same figure for the FTSE 250 and full sample.

Figure B.1: **Percentage of 1st Successful and Failed Messages by Firm**



Notes: Please see the notes for Figure 4.2 in the main text and the description in the text above this figure.

Additional Details for Expected Number of Races by Chance Analysis

In Section 4.2 of the main text, in the subsection “Expected Number of Races by Chance,” we discussed the number of times per day we would see N messages on the same side of the order book within T microseconds, by chance, if orders arrive randomly according to a Poisson process. Poisson processes are memoryless meaning that the arrival of a message at one point in time does not make it any more or less likely for other messages to arrive in the interval of time thereafter. We concluded that clusters of messages within short time horizons would be very rare if messages arrive Poisson.

We also considered a sensitivity, to account for time-varying message arrival rates as in Engle and Russell (1998), in which we assume the entire trading day for a symbol-date is as busy as its busiest 30-minutes. Again, we concluded that clusters of messages within the kinds of ultra-short time horizons we observe in our data would be very rare.

This appendix table provides additional support for that discussion. In the first table we determine the Poisson arrival rate for each symbol-date based on the total number of potentially-race-relevant messages (i.e., marketable takes or cancels at the best bid or offer) for that symbol-date. In the second table we determine the Poisson arrival rate for each symbol-date based on the rate of potentially-race-relevant messages in that symbol-date’s busiest 30-minute increment. For each, we report the distribution over symbol-dates of the expected number of instances per day in which one would see N participants within T microseconds, and compare this to the actual observed number of races.

Table B.6: **Expected Number of Potential Race Events by Chance: Average Message Arrival Rate**

FTSE 100								
N	T	Mean	sd	Pct01	Pct25	Median	Pct75	Pct99
2	50	0.35	0.80	0.01	0.04	0.09	0.32	3.28
2	100	0.71	1.60	0.02	0.08	0.18	0.64	6.56
2	200	1.42	3.20	0.03	0.15	0.37	1.29	13.13
2	500	3.55	7.99	0.08	0.38	0.91	3.22	32.81
2	1000	7.09	15.96	0.15	0.77	1.83	6.44	65.57
3	1000	0.00	0.02	0.00	0.00	0.00	0.00	0.05
Actual Number of Races								
Baseline analysis		537.24	542.96	73	215	346	629	2,635
Sensitivity: 3+ within Info Horizon		228.98	206.88	28	100	161	278	1,002
FTSE 250								
N	T	Mean	sd	Pct01	Pct25	Median	Pct75	Pct99
2	50	0.00	0.01	0.00	0.00	0.00	0.00	0.04
2	100	0.01	0.02	0.00	0.00	0.00	0.01	0.09
2	200	0.02	0.04	0.00	0.00	0.01	0.02	0.17
2	500	0.04	0.10	0.00	0.00	0.02	0.04	0.43
2	1000	0.08	0.20	0.00	0.01	0.03	0.08	0.86
3	1000	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Actual Number of Races								
Baseline analysis		70.05	103.26	0	2	35	97	477
Sensitivity: 3+ within Info Horizon		30.68	49.17	0	0	12	43	223

Notes: Please see the notes for Table 4.5 in the main text. This table provides full distributional details for the columns marked “Average Rate” in Table 4.5 in the main text.

Table B.7: **Expected Number of Potential Race Events by Chance: Busiest 30 Minutes Message Arrival Rate**

FTSE 100								
N	T	Mean	sd	Pct01	Pct25	Median	Pct75	Pct99
2	50	1.33	3.22	0.03	0.12	0.31	1.19	13.75
2	100	2.65	6.43	0.05	0.25	0.62	2.38	27.49
2	200	5.31	12.86	0.10	0.50	1.25	4.76	54.97
2	500	13.26	32.12	0.26	1.24	3.12	11.90	137.31
2	1000	26.49	64.11	0.52	2.49	6.23	23.79	274.22
3	1000	0.03	0.14	0.00	0.00	0.00	0.01	0.39
Actual Number of Races								
Baseline analysis		537.24	542.96	73	215	346	629	2,635
Sensitivity: 3+ within Info Horizon		228.98	206.88	28	100	161	278	1,002
FTSE 250								
N	T	Mean	sd	Pct01	Pct25	Median	Pct75	Pct99
2	50	0.02	0.05	0.00	0.00	0.01	0.02	0.22
2	100	0.04	0.10	0.00	0.01	0.02	0.04	0.44
2	200	0.09	0.19	0.00	0.01	0.03	0.09	0.89
2	500	0.21	0.49	0.00	0.03	0.09	0.21	2.22
2	1000	0.43	0.97	0.00	0.06	0.17	0.43	4.43
3	1000	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Actual Number of Races								
Baseline analysis		70.05	103.26	0	2	35	97	477
Sensitivity: 3+ within Info Horizon		30.68	49.17	0	0	12	43	223

Notes: Please see the notes for Table 4.5 in the main text. This table provides full distributional details for the columns marked “Busiest 30 Mins” in Table 4.5 in the main text.

B.3 Additional Results for Section 4.3

Profits Per-Race Full Sample

Tables 4.6 and 4.7 in the main text reported data on the distribution of race profits for FTSE 100 and FTSE 250 symbols separately. The following tables report the same information for the combined full sample.

Table B.8: **Detail on Race Profits (Per-Share and Per-Race) Marked to Market at 10s (Appendix)**

Full Sample									
Description	Mean	sd	Pct01	Pct10	Pct25	Median	Pct75	Pct90	Pct99
Per-share profits (ticks)	0.55	3.92	-6.50	-1.50	-0.50	0.50	1.00	3.00	10.00
Per-share profits (GBX)	0.17	1.48	-2.00	-0.50	-0.05	0.01	0.25	1.00	3.50
Per-share profits (basis points)	1.66	8.71	-15.00	-4.26	-1.29	0.50	3.89	7.98	27.02
Per-race profits displayed depth (GBP)	1.85	16.27	-20.00	-2.76	-0.34	0.00	2.15	7.27	41.50
Per-race profits qty trade/cancel (GBP)	1.76	15.57	-18.13	-2.56	-0.32	0.00	2.02	6.78	38.44

Notes: Please see the notes for Table 4.6 in the main text. This table reports the same statistics for the full sample.

Table B.9: **Average Race Profits (Per-Share and Per-Race) for Different Mark to Market Horizons (Appendix)**

Full Sample									
Description	1ms	10ms	100ms	1s	10s	30s	60s	100s	
Mean per-share profits (ticks)	0.03	0.21	0.29	0.40	0.55	0.59	0.63	0.64	
Mean per-share profits (GBX)	0.03	0.08	0.10	0.13	0.17	0.18	0.18	0.18	
Mean per-share profits (basis points)	0.18	0.67	0.89	1.20	1.66	1.83	1.94	1.97	
Mean per-race profits displayed depth (GBP)	0.28	0.96	1.24	1.54	1.85	1.86	1.88	1.84	
Mean per-race profits qty trade/cancel (GBP)	0.31	0.94	1.18	1.45	1.76	1.76	1.77	1.74	

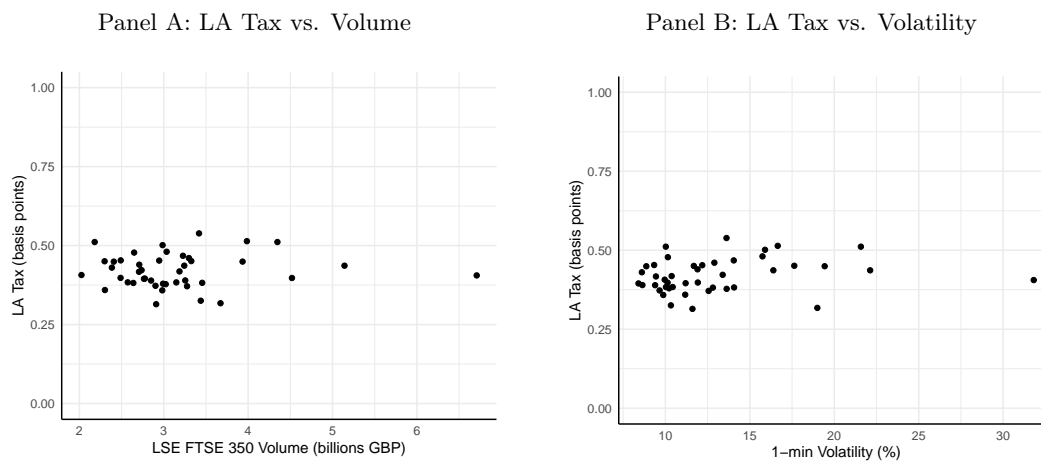
Notes: Please see the notes for Table 4.7 in the main text. This table reports the same statistics for the full sample.

B.4 Additional Results for Section 4.4

Latency Arbitrage Tax Correlation with Volume and Volatility

Figure 4.5 of the main text presents a scatterplot of latency arbitrage profits against volume and volatility. This figure is analogous but plots the latency arbitrage tax against volume and volatility.

Figure B.2: Latency Arbitrage Tax Correlation with Volume and Volatility



Notes: Panel A presents a scatterplot of the daily latency arbitrage tax, defined as daily race profits for the full sample divided by daily regular-hours trading volume, against regular-hours trading volume. Panel B presents a scatterplot of the daily latency arbitrage tax against daily realized 1-minute volatility for the FTSE 350 index. Please see the notes for Figure 4.5 in the main text which is closely related.

B.5 Additional Results for Section 4.5

Distribution of the Bid-Ask Spread by Symbol and Date

Table 4.10 in the main text presents a decomposition of the bid-ask spread into price impact in races, price impact not in races, loss avoidance, and the realized spread. For context on this analysis, this appendix table presents the distribution of the bid-ask spread across symbol (averaged over all dates) and dates (averaged over all symbols). Spreads are presented based on both the time-weighted displayed spread (Panel A) and the quantity-weighted traded spread (Panel B); this latter quantity-weighted spread corresponds to the term effective spread utilized in the literature and in the text of Section 4.5. For each analysis, we present results in both ticks (sub-panel A) and basis points (sub-panel B); this latter measurement corresponds to the spread decomposition reported in the text. All spreads are reported as the “half-spread”, i.e., half the distance between the bid and the offer, which corresponds to the difference between the tradable or traded price and the midpoint price. The half-spread is the standard measure in the literature.

Table B.10: **Spread by Date**

Panel A: Time-Weighted Average Half-Spread

Sub-Panel A: Ticks

Description	Mean	sd	Min	Pct10	Pct25	Median	Pct75	Pct90	Max
FTSE 100	0.97	0.06	0.86	0.92	0.93	0.96	1.00	1.04	1.20
FTSE 250	3.40	0.35	2.83	2.99	3.19	3.34	3.61	3.81	4.38
Full Sample	2.70	0.26	2.29	2.39	2.53	2.63	2.86	2.98	3.45

Sub-Panel B: Basis Points

Description	Mean	sd	Min	Pct10	Pct25	Median	Pct75	Pct90	Max
FTSE 100	3.77	0.20	3.42	3.54	3.66	3.76	3.82	3.97	4.39
FTSE 250	15.76	1.48	13.11	13.97	14.81	15.62	16.66	17.67	19.62
Full Sample	12.27	1.09	10.35	10.92	11.55	12.22	12.93	13.63	15.19

Panel B: Quantity-Weighted Average Half-Spread ("Effective Spread")

Sub-Panel A: Ticks

Description	Mean	sd	Min	Pct10	Pct25	Median	Pct75	Pct90	Max
FTSE 100	0.85	0.17	0.70	0.74	0.76	0.80	0.86	1.00	1.71
FTSE 250	1.44	0.13	1.15	1.31	1.37	1.44	1.47	1.53	1.82
Full Sample	0.93	0.15	0.77	0.83	0.85	0.88	0.95	1.06	1.66

Sub-Panel B: Basis Points

Description	Mean	sd	Min	Pct10	Pct25	Median	Pct75	Pct90	Max
FTSE 100	2.65	0.29	2.28	2.45	2.52	2.59	2.72	2.80	4.28
FTSE 250	6.76	0.58	5.72	6.24	6.44	6.66	6.95	7.19	8.97
Full Sample	3.17	0.27	2.74	2.92	3.06	3.12	3.22	3.38	4.52

Notes: Please see the description in the text above this table for a description of the spread variables. This table reports distributions of the spread at the date level, averaging over symbols.

Table B.11: **Spread by Symbol**

Panel A: Time-Weighted Average Half-Spread

Sub-Panel A: Ticks

Description	Mean	sd	Pct01	Pct10	Pct25	Median	Pct75	Pct90	Pct99
FTSE 100	0.97	0.32	0.56	0.64	0.83	0.92	1.02	1.32	2.14
FTSE 250	3.40	3.00	0.83	1.09	1.53	2.57	3.94	6.52	16.73
Full Sample	2.70	2.76	0.58	0.85	1.01	1.79	3.25	5.67	12.86

Sub-Panel B: Basis Points

Description	Mean	sd	Pct01	Pct10	Pct25	Median	Pct75	Pct90	Pct99
FTSE 100	3.77	1.56	1.09	1.70	2.56	3.77	4.85	5.49	7.59
FTSE 250	15.76	13.67	3.38	6.36	7.74	11.32	17.92	29.90	59.41
Full Sample	12.27	12.76	1.21	3.09	4.95	8.10	15.04	27.07	56.01

Panel B: Quantity-Weighted Average Half-Spread ("Effective Spread")

Sub-Panel A: Ticks

Description	Mean	sd	Pct01	Pct10	Pct25	Median	Pct75	Pct90	Pct99
FTSE 100	0.80	0.27	0.52	0.55	0.64	0.73	0.89	1.17	1.71
FTSE 250	2.09	1.42	0.60	0.84	1.13	1.75	2.58	3.80	6.62
Full Sample	1.72	1.34	0.54	0.66	0.81	1.32	2.17	3.21	6.38

Sub-Panel B: Basis Points

Description	Mean	sd	Pct01	Pct10	Pct25	Median	Pct75	Pct90	Pct99
FTSE 100	3.27	1.22	1.22	1.75	2.28	3.18	4.13	4.91	5.79
FTSE 250	11.61	9.53	2.66	4.90	5.99	8.22	13.67	22.96	47.35
Full Sample	9.18	8.90	1.29	2.59	4.21	6.26	10.38	18.47	40.07

Notes: Please see the description in the text above this table for a description of the spread variables. This table reports distributions of the spread at the symbol level, averaging over dates.

Spread Decomposition Full Sample by Date

Table 4.10 in the main text presents our spread decomposition for FTSE 100 symbols. The following table present the same spread decomposition for FTSE 250 symbols and for the full sample by date.

Table B.12: **Spread Decomposition (Appendix)**

Panel A: FTSE 250 by Symbol									
Description	Mean	sd	Pct01	Pct10	Pct25	Median	Pct75	Pct90	Pct99
Effective spread paid - overall (bps)	8.06	3.81	2.65	4.63	5.59	7.14	9.84	13.10	19.11
Effective spread paid - in races (bps)	6.74	3.03	2.42	4.32	4.97	6.08	7.63	9.96	15.62
Effective spread paid - not in races (bps)	8.22	3.87	2.72	4.70	5.72	7.31	9.94	13.34	19.55
Price impact - overall (bps)	8.09	3.54	2.64	4.96	5.71	7.10	9.40	12.95	19.91
Price impact - in races (bps)	12.22	6.19	4.04	7.17	8.82	10.72	13.75	18.12	33.42
Price impact - not in races (bps)	7.50	3.52	2.36	4.37	5.09	6.40	8.79	12.39	19.39
Loss avoidance (bps)	0.01	0.02	-0.02	0.00	0.00	0.01	0.01	0.02	0.07
Realized spread - overall (bps)	-0.04	1.14	-2.30	-1.02	-0.53	-0.14	0.34	0.96	2.67
Realized spread - in races (bps)	-5.48	3.68	-20.22	-9.36	-6.14	-4.43	-3.44	-2.73	-1.62
Realized spread - not in races (bps)	0.72	1.07	-0.97	-0.13	0.20	0.59	1.07	1.76	3.14
PI in races / PI total (%)	21.60	9.50	1.79	6.00	14.89	22.98	28.19	32.16	39.60
PI in races / Effective spread (%)	22.50	10.92	1.58	5.62	14.84	23.57	30.44	34.79	47.67

Panel B: Full Sample by Date									
Description	Mean	sd	Min	Pct10	Pct25	Median	Pct75	Pct90	Max
Effective spread paid - overall (bps)	3.17	0.27	2.74	2.92	3.06	3.12	3.22	3.38	4.52
Effective spread paid - in races (bps)	2.99	0.13	2.64	2.84	2.90	2.99	3.06	3.16	3.28
Effective spread paid - not in races (bps)	3.22	0.32	2.77	2.95	3.09	3.17	3.29	3.44	4.90
Price impact - overall (bps)	3.38	0.19	2.96	3.19	3.23	3.38	3.52	3.61	3.80
Price impact - in races (bps)	4.82	0.24	4.35	4.53	4.66	4.79	4.99	5.07	5.55
Price impact - not in races (bps)	2.99	0.19	2.57	2.79	2.86	2.95	3.13	3.29	3.38
Loss avoidance (bps)	0.01	0.00	-0.01	0.00	0.00	0.01	0.01	0.01	0.01
Realized spread - overall (bps)	-0.22	0.23	-0.62	-0.38	-0.31	-0.26	-0.15	-0.09	1.08
Realized spread - in races (bps)	-1.83	0.17	-2.43	-2.01	-1.92	-1.81	-1.74	-1.64	-1.51
Realized spread - not in races (bps)	0.23	0.26	-0.17	0.05	0.14	0.20	0.29	0.34	1.68
PI in races / PI total (%)	30.58	2.64	22.91	27.88	29.88	30.81	31.93	33.39	35.81
PI in races / Effective spread (%)	32.82	3.73	17.38	29.92	31.60	33.66	34.70	36.54	39.52

Notes: Please see the notes for Table 4.10 in the main text. Panel A reports the distribution for all symbols in the FTSE 250. We only include symbols that have at least 100 races summed over all dates; this drops about one-quarter of FTSE 250 symbols. Panel B reports the distribution by date for the full sample.

Spread Decomposition with Different Time Horizons

Table 4.10 in the main text reports results of our spread decomposition (Section 4.5 of the main text, Approach #1) using a 10 second mark-to-market time horizon for calculating price impact and loss avoidance. In this appendix we report the same decomposition but using 100 millisecond and 1 second time horizons instead. Notably, the realized spread appears to decline with the time horizon, from 100 millisecond to 1 second to 10 seconds, both in and out of races. While the overall sample realized spread is slightly negative at 10 seconds, it is slightly positive at 100 millisecond and 1 second. This pattern is consistent with price impact being smaller at shorter time horizons as discussed in Conrad and Wahal (2020).

Table B.13: Spread Decomposition - 100ms

Panel A: FTSE 100 by Symbol

Description	Mean	sd	Pct01	Pct10	Pct25	Median	Pct75	Pct90	Pct99
Effective spread paid - overall (bps)	3.27	1.22	1.22	1.75	2.28	3.18	4.13	4.91	5.79
Effective spread paid - in races (bps)	3.18	1.22	0.99	1.70	2.21	3.17	4.05	4.89	5.98
Effective spread paid - not in races (bps)	3.29	1.22	1.25	1.78	2.30	3.17	4.15	4.96	5.71
Price impact - overall (bps)	3.18	1.25	1.16	1.71	2.18	3.06	3.96	5.06	5.82
Price impact - in races (bps)	4.52	1.75	1.61	2.52	3.07	4.26	5.76	7.23	7.89
Price impact - not in races (bps)	2.75	1.03	1.03	1.47	1.92	2.72	3.36	4.25	4.94
Loss avoidance (bps)	0.00	0.01	-0.01	-0.00	0.00	0.00	0.00	0.01	0.02
Realized spread - overall (bps)	0.09	0.27	-0.43	-0.20	-0.03	0.06	0.18	0.37	1.06
Realized spread - in races (bps)	-1.33	0.62	-2.80	-2.32	-1.68	-1.11	-0.88	-0.71	-0.53
Realized spread - not in races (bps)	0.55	0.30	0.08	0.22	0.29	0.50	0.74	0.92	1.41
PI in races / PI total (%)	33.26	6.28	21.27	25.97	29.36	31.77	37.35	43.12	46.06
PI in races / Effective spread (%)	32.49	7.56	18.81	23.89	28.30	30.96	36.37	43.84	49.45

Panel B: FTSE 250 by Symbol

Description	Mean	sd	Pct01	Pct10	Pct25	Median	Pct75	Pct90	Pct99
Effective spread paid - overall (bps)	8.06	3.81	2.65	4.63	5.59	7.14	9.84	13.10	19.10
Effective spread paid - in races (bps)	6.74	3.03	2.42	4.32	4.97	6.08	7.63	9.96	15.62
Effective spread paid - not in races (bps)	8.22	3.87	2.72	4.70	5.72	7.31	9.94	13.34	19.55
Price impact - overall (bps)	5.99	2.47	2.24	3.58	4.34	5.44	7.09	9.23	14.30
Price impact - in races (bps)	9.38	4.87	3.50	5.39	6.51	8.23	11.07	13.93	26.88
Price impact - not in races (bps)	5.53	2.45	2.02	3.26	3.86	4.89	6.55	8.94	13.50
Loss avoidance (bps)	-0.00	0.02	-0.05	-0.02	-0.01	-0.00	0.00	0.01	0.06
Realized spread - overall (bps)	2.07	1.69	-0.04	0.45	1.17	1.82	2.57	3.51	6.97
Realized spread - in races (bps)	-2.64	2.75	-12.96	-5.92	-3.14	-1.97	-1.06	-0.47	0.99
Realized spread - not in races (bps)	2.69	1.70	0.42	1.22	1.74	2.44	3.18	4.28	7.07
PI in races / PI total (%)	21.82	9.31	2.14	7.49	15.08	23.34	28.22	32.29	39.41
PI in races / Effective spread (%)	17.14	8.59	1.30	4.59	10.97	17.30	22.54	27.63	37.15

Panel C: Full Sample by Date

Description	Mean	sd	Min	Pct10	Pct25	Median	Pct75	Pct90	Max
Effective spread paid - overall (bps)	3.17	0.27	2.74	2.92	3.06	3.12	3.22	3.38	4.52
Effective spread paid - in races (bps)	2.99	0.13	2.64	2.84	2.90	2.99	3.06	3.16	3.28
Effective spread paid - not in races (bps)	3.22	0.32	2.77	2.95	3.10	3.17	3.29	3.44	4.90
Price impact - overall (bps)	2.88	0.16	2.54	2.71	2.79	2.90	2.95	3.13	3.18
Price impact - in races (bps)	4.22	0.17	3.81	4.00	4.13	4.22	4.35	4.45	4.60
Price impact - not in races (bps)	2.52	0.15	2.19	2.33	2.43	2.52	2.58	2.72	2.84
Loss avoidance (bps)	0.00	0.00	-0.00	0.00	0.00	0.00	0.00	0.00	0.02
Realized spread - overall (bps)	0.29	0.23	0.11	0.17	0.20	0.24	0.30	0.39	1.66
Realized spread - in races (bps)	-1.24	0.08	-1.48	-1.33	-1.28	-1.23	-1.19	-1.13	-1.06
Realized spread - not in races (bps)	0.70	0.26	0.51	0.57	0.61	0.65	0.73	0.81	2.26
PI in races / PI total (%)	31.43	2.31	24.08	28.54	30.40	31.69	32.47	34.07	36.64
PI in races / Effective spread (%)	28.77	3.12	15.24	26.47	27.76	29.26	30.37	31.92	34.52

Notes: Please see the notes for Table 4.10 in the main text. This table is the same except that price impact and loss avoidance are calculated based on mark-to-market at 100 milliseconds instead of 10 seconds. Panel A reports the distributions for all symbols in the FTSE 100. Panel B reports the distribution for all symbols in the FTSE 250. We only include symbols that have at least 100 races summed over all dates; this drops about one-quarter of FTSE 250 symbols and does not drop any FTSE 100 symbols. Panel C reports the distribution of these statistics by date for the full sample.

Table B.14: Spread Decomposition - 1s

Panel A: FTSE 100 by Symbol

Description	Mean	sd	Pct01	Pct10	Pct25	Median	Pct75	Pct90	Pct99
Effective spread paid - overall (bps)	3.27	1.22	1.22	1.75	2.28	3.18	4.13	4.91	5.79
Effective spread paid - in races (bps)	3.18	1.22	0.99	1.70	2.21	3.17	4.05	4.89	5.98
Effective spread paid - not in races (bps)	3.29	1.22	1.25	1.78	2.30	3.17	4.15	4.96	5.71
Price impact - overall (bps)	3.39	1.29	1.27	1.85	2.34	3.34	4.15	5.20	6.30
Price impact - in races (bps)	4.81	1.78	1.83	2.78	3.33	4.63	6.04	7.44	8.33
Price impact - not in races (bps)	2.93	1.07	1.13	1.60	2.06	2.98	3.51	4.44	5.39
Loss avoidance (bps)	0.00	0.01	-0.00	-0.00	0.00	0.00	0.01	0.01	0.02
Realized spread - overall (bps)	-0.12	0.25	-0.56	-0.38	-0.25	-0.15	-0.00	0.14	0.76
Realized spread - in races (bps)	-1.63	0.62	-3.24	-2.54	-1.98	-1.48	-1.15	-0.91	-0.76
Realized spread - not in races (bps)	0.36	0.28	-0.09	0.06	0.16	0.32	0.55	0.72	1.13
PI in races / PI total (%)	33.29	6.26	20.88	25.73	29.49	32.11	37.49	42.69	46.16
PI in races / Effective spread (%)	34.74	7.42	19.79	26.20	30.94	34.06	39.08	44.93	49.85

Panel B: FTSE 250 by Symbol

Description	Mean	sd	Pct01	Pct10	Pct25	Median	Pct75	Pct90	Pct99
Effective spread paid - overall (bps)	8.06	3.81	2.65	4.63	5.59	7.14	9.84	13.10	19.11
Effective spread paid - in races (bps)	6.74	3.03	2.42	4.32	4.97	6.08	7.63	9.96	15.62
Effective spread paid - not in races (bps)	8.22	3.87	2.72	4.70	5.72	7.31	9.94	13.34	19.55
Price impact - overall (bps)	6.71	2.83	2.43	4.14	4.95	5.98	7.79	10.34	17.10
Price impact - in races (bps)	10.44	5.46	3.75	6.14	7.33	9.10	12.28	15.39	29.90
Price impact - not in races (bps)	6.20	2.82	2.18	3.63	4.41	5.41	7.23	9.85	16.38
Loss avoidance (bps)	-0.00	0.01	-0.04	-0.01	-0.00	-0.00	0.00	0.01	0.07
Realized spread - overall (bps)	1.35	1.44	-0.46	0.06	0.57	1.11	1.73	2.66	5.68
Realized spread - in races (bps)	-3.70	3.14	-16.39	-6.99	-4.13	-2.65	-1.99	-1.44	-0.69
Realized spread - not in races (bps)	2.02	1.44	0.22	0.81	1.25	1.80	2.43	3.38	5.89
PI in races / PI total (%)	21.79	9.41	2.10	6.72	15.03	23.58	28.40	32.31	39.77
PI in races / Effective spread (%)	19.03	9.41	1.61	5.19	12.08	19.61	25.39	30.01	41.32

Panel C: Full Sample by Date

Description	Mean	sd	Min	Pct10	Pct25	Median	Pct75	Pct90	Max
Effective spread paid - overall (bps)	3.17	0.27	2.74	2.92	3.06	3.12	3.22	3.38	4.52
Effective spread paid - in races (bps)	2.99	0.13	2.64	2.84	2.90	2.99	3.06	3.16	3.28
Effective spread paid - not in races (bps)	3.22	0.32	2.77	2.95	3.10	3.17	3.29	3.44	4.90
Price impact - overall (bps)	3.10	0.17	2.72	2.90	3.00	3.11	3.21	3.36	3.44
Price impact - in races (bps)	4.51	0.20	4.08	4.26	4.39	4.51	4.66	4.75	4.98
Price impact - not in races (bps)	2.71	0.17	2.35	2.54	2.61	2.71	2.78	2.99	3.06
Loss avoidance (bps)	0.00	0.00	-0.01	0.00	0.00	0.00	0.00	0.01	0.01
Realized spread - overall (bps)	0.07	0.22	-0.11	-0.06	-0.02	0.02	0.12	0.19	1.31
Realized spread - in races (bps)	-1.52	0.11	-1.86	-1.65	-1.58	-1.52	-1.45	-1.40	-1.32
Realized spread - not in races (bps)	0.50	0.24	0.29	0.36	0.41	0.46	0.55	0.62	1.89
PI in races / PI total (%)	31.24	2.41	23.10	28.32	30.29	31.69	32.37	33.99	36.59
PI in races / Effective spread (%)	30.71	3.37	16.41	28.06	29.47	31.27	32.89	34.03	36.64

Notes: Please see the notes for Table 4.10 in the main text. This table is the same except that price impact and loss avoidance are calculated based on mark-to-market at 1 second instead of 10 seconds. Panel A reports the distributions for all symbols in the FTSE 100. Panel B reports the distribution for all symbols in the FTSE 250. We only include symbols that have at least 100 races summed over all dates; this drops about one-quarter of FTSE 250 symbols and does not drop any FTSE 100 symbols. Panel C reports the distribution of these statistics by date for the full sample.

C Supporting Details for Section 5.1: Sensitivity Analysis

This appendix presents detailed sensitivity analyses for the main results presented in the body of the paper, as discussed in detail in Section 5.1. Section C.1 explores sensitivity to the race horizon, i.e., to the definition of what counts as “at the same time.” Section C.2 explores sensitivity to the number of race participants, e.g., requiring 3+ participants at the same time rather than 2+. Section C.3 explores sensitivity to requiring cancel attempts in the race, i.e., to not counting races that contain only aggressive orders, and also explores stricter requirements on the number of aggressive orders. Section C.4 explores varying the definition of what counts as a success and a fail. Together, then, Sections C.1-C.4 explore sensitivity to the four components of our race definition: multiple participants, at the same time, at least some of whom are aggressive, and at least some of whom succeed and some of whom fail.

C.1 Sensitivity to Race Horizon

As a reminder, our baseline method requires that messages satisfying the baseline race requirements (i.e., 2+ messages from distinct users, 1+ aggressing, 1+ success, and 1+ fail) arrive within the “information horizon” of the first message of the race. The information horizon, which is the window of time such that we can be essentially certain that inbound messages in the race are not responding to earlier messages’ outbound reports (see Section 3.4 of the main text) has a median of 186 microseconds in our data, a 10th percentile of 137 and a 90th percentile of 332. The 500 microsecond truncation binds 4% of the time.

Table C.1 presents sensitivity analysis for changes to the race horizon. The first column of the table re-presents our main results from Section 4 in the main text for this baseline specification, to facilitate comparison. The next set of columns presents these same results using fixed race horizons of varying lengths, from 50 microseconds to 3 milliseconds. That is, instead of using the information horizon method, under which the race window will vary with the observed lag in information processing by the LSE’s matching engine, we just fix a time window, and consider a wide range of such windows. The 50 microsecond window roughly corresponds to the minimum observed information horizon (which is 43 microseconds), the 200 microsecond window roughly corresponds to the median observed information horizon, and 500 microseconds corresponds to the upper bound on the information horizon we determined in consultation with FCA supervisory experts. The horizons beyond that are included to capture races among firms of varying technological sophistication that could still be considered racing one another. For instance, the threshold could be wide enough to include a firm that is not utilizing the fastest connections to exchanges in the United States or elsewhere, but is using the next-fastest.⁴ We consulted with HFT industry contacts and FCA supervisors to agree on an appropriate horizon. Following these discussions, we determined 3 milliseconds would capture most of these additional potential races, though for races originating from signals far from London (e.g., Chicago) differences in speed between cutting-edge

⁴Other sources of speed differential include using code and hardware that is not optimized for speed, not being co-located, and not using microwave connections where possible to do so.

HFTs and relatively sophisticated firms could easily exceed that number. The last set of columns runs a sensitivity analysis specifically on the choice of the truncation parameter for the information horizon method.

Focus first on the number of races per day per symbol in the FTSE 100, the first row of the table. In the baseline there are 537 races per symbol per day. In the 50 microsecond column, this number is reduced to 297. As the race horizon increases, so does the number of races detected. The growth is especially steep up to 500 microseconds, reaching 793 races per symbol per day, and then tapers off, with 870 races at a horizon of 1 millisecond and 946 races at a horizon of 3 milliseconds.⁵ Varying the truncation parameter for the information horizon method does not yield much additional insight beyond what is already learned from the baseline and the fixed horizon columns. Using a 100 microsecond truncation parameter yields results that are very similar to the 100 microsecond fixed race horizon, which makes sense since this truncation parameter will bind most of the time. Using a 1 millisecond truncation parameter yields results that are similar to the baseline with the 500 microsecond truncation parameter, which again makes sense because neither truncation parameter will bind very much.

Turn next to the measures of per-race profits. Interestingly, per-race profits, whether measured per-share (ticks, pence (GBX), basis points) or in GBP per-race (either displayed depth or quantity actually traded/canceled), are relatively similar across these different specifications. This tells us that the additional races being picked up by the longer race horizons are, on average, of similar profitability to the races being picked up at shorter race horizons. This will not be the case for some of the subsequent sensitivities.

As a result, the latency arbitrage tax measures are all increasing with the race horizon. At a 50 microsecond race horizon, the FTSE 350 latency arbitrage tax, using the all-volume measure, is 0.20 basis points, versus 0.42 basis points in our baseline specification. At the 3 millisecond race horizon, the latency arbitrage tax is 0.81 basis points, or 4 times higher, roughly proportional to the increase in the number of races. The effect on the second measure of the latency arbitrage tax, based on non-race trading volume, is even larger, because as the numerator (race profits) is increasing, the denominator (non-race volume) is also shrinking. This figure increases from 0.22 basis points at 50 microseconds, to 0.53 basis points in our baseline specification, all the way up to 1.55 basis points at 3 milliseconds. For FTSE 250 stocks, the latency arbitrage tax is as high as 2.49 basis points at 3 milliseconds.

Last we discuss the implied reduction in the cost of liquidity. In our baseline, eliminating latency arbitrage would reduce the cost of liquidity by 20.0% for the average FTSE 100 symbol and by 16.7% for the market overall. Using a 50 microsecond race horizon lowers these figures to 8.0% and 7.0%, respectively. Using a 3 millisecond race horizon increases these figures all the way to 59.2% and 48.8%, respectively. Again, this large change relative to the baseline is driven by both the increase in the numerator (race profits) and decrease in the denominator (non-race effective spread paid).

⁵These numbers are slightly higher than the numbers reported in Table 5.1 in the main text because they use the baseline definition of fails rather than the stricter definition of fails that is appropriate for longer time horizons. See further discussion in Appendix C.4.

Table C.1: Sensitivity Analysis: Different Race Horizons

Measure	Baseline	Fixed Race Horizon of Duration T					Info Horizon, Max T			
		50 μ s	100 μ s	200 μ s	500 μ s	1ms	2ms	3ms	100 μ s	1ms
Frequency and Duration of Races										
Races per day										
FTSE 100 - per symbol	537.24	296.66	388.58	521.53	793.01	869.73	921.08	946.48	387.96	542.99
FTSE 250 - per symbol	70.05	41.37	52.78	69.22	112.99	127.04	134.06	138.37	52.71	70.28
Mean race duration (μ s)	78.65	16.12	30.80	72.18	194.20	304.96	450.87	572.12	30.61	84.85
% of races with wrong winner	4.30	8.18	6.41	4.21	1.98	1.67	1.43	1.32	6.42	4.24
% of volume in races										
FTSE 100	22.15	9.99	13.65	19.71	37.43	43.53	47.11	48.61	13.64	22.65
FTSE 250	16.90	8.36	11.20	15.99	33.34	38.37	41.23	42.63	11.20	17.07
Full Sample	21.46	9.77	13.32	19.21	36.88	42.84	46.33	47.82	13.32	21.92
Mean number of messages within 500 μ s	3.46	3.51	3.51	3.51	3.39	3.01	2.83	2.76	3.51	3.44
Per-Race Profits										
Per-share profits										
ticks	0.55	0.54	0.53	0.51	0.53	0.55	0.56	0.57	0.53	0.56
GBX	0.17	0.16	0.16	0.16	0.16	0.16	0.17	0.16	0.16	0.17
basis points	1.66	1.68	1.63	1.57	1.61	1.64	1.65	1.65	1.63	1.67
Per-race profits GBP										
displayed depth	1.85	1.58	1.59	1.60	1.84	1.94	1.97	1.97	1.60	1.90
qty trade/cancel	1.76	1.38	1.44	1.51	1.84	1.95	2.00	2.00	1.44	1.81
Aggregate Profits and LA Tax										
Daily Profits										
FTSE 100 - per symbol	1,047	490	647	872	1,520	1,769	1,909	1,965	647	1,089
FTSE 250 - per symbol	108	57	73	96	184	211	226	231	73	110
Full Sample - aggregate	132,378	63,573	83,233	111,722	198,700	230,586	248,291	255,408	83,181	137,173
Latency Arbitrage Tax, All Volume (bps)										
FTSE 100	0.38	0.18	0.24	0.32	0.56	0.65	0.70	0.72	0.24	0.40
FTSE 250	0.66	0.35	0.45	0.59	1.13	1.30	1.39	1.42	0.45	0.68
Full Sample	0.42	0.20	0.26	0.35	0.63	0.73	0.78	0.81	0.26	0.43
Latency Arbitrage Tax, Non-Race Volume (bps)										
FTSE 100	0.49	0.20	0.27	0.40	0.89	1.15	1.32	1.40	0.27	0.52
FTSE 250	0.80	0.38	0.50	0.71	1.70	2.12	2.37	2.49	0.50	0.82
Full Sample	0.53	0.22	0.30	0.44	1.00	1.28	1.47	1.55	0.30	0.56
Spread Decomposition										
Price impact in races / All price impact %	30.58	12.84	17.89	25.69	49.79	58.71	64.34	66.82	17.88	31.76
Price impact in races / Effective spread %	32.82	13.77	19.19	27.57	53.42	62.99	69.03	71.69	19.19	34.08
Loss avoidance / Effective spread %	0.19	0.07	0.13	0.26	0.53	0.94	1.31	1.48	0.13	0.20
Implied Reduction in Cost of Liquidity										
% Reduction in liquidity cost										
FTSE 100 - by symbol	19.95	7.98	10.97	15.91	35.73	46.95	55.24	59.20	10.97	21.00
FTSE 250 - by symbol	11.93	6.17	7.96	10.79	24.36	28.46	31.70	32.90	7.95	12.17
Full Sample - by date	16.73	6.96	9.49	13.62	30.38	39.20	45.62	48.75	9.49	17.49

Notes: For descriptions of the sensitivity scenarios please see the text of Section C.1. Descriptions of each of the items in this table can be found in the following table notes in Section 4. Races per day: Table 4.1. Mean race duration and % of races with wrong winner: Figure 4.1. % of Volume in Races: Table 4.2. Mean number of messages: Table 4.3. Per-race profits: Table 4.6. Aggregate profits: Table 4.8. Latency Arbitrage Tax: Table 4.9. Spread decomposition: Table 4.10. Implied Reduction in Cost of Liquidity: Table 4.12.

C.2 Sensitivity to Number of Race Participants

Our baseline method requires that there are at least 2 race participants within the information horizon. Table C.2 presents sensitivity analysis for requiring 3+ participants; Table C.3 presents the same table for 5+ participants. In both cases, the other race criteria are held the same, specifically we require 1+ aggressors, 1+ successes, and 1+ fails. Given the large effect that the race’s time horizon had on the number of races and race profits, we include this sensitivity for multiple race horizons, including the baseline information horizon method and fixed race horizons from 50 microseconds to 3 milliseconds.

Focus first on the 3+ race participants within information horizon column; this column is exactly the same as the baseline but replacing 2+ race participants with 3+. Requiring 3+ race participants reduces the number of races by about 60%; for example, for the FTSE 100 the number of races per symbol per day declines from 537 to 229. However, these races are significantly more profitable, on a per-share basis and particularly on a GBP per-race basis. The net effect is that total race profits are reduced by about 30%. This roughly 30% reduction can be seen in the aggregate race profits measures, the latency arbitrage tax measures, and the liquidity cost reduction measures.

Increasing the race horizon increases the number of races detected, just as in the baseline case with 2+ participants. At a 50 microsecond race horizon there are 87 3+ participant races per day for the average FTSE 100 symbol, up to 482 races per symbol per day at a 500 microsecond race horizon, and up to 686 races at a 3 millisecond race horizon. With this increase in the number of races detected comes a commensurate increase in the various race profits measures and harm-to-liquidity measures.

We note that the 3+ race participants within 500 microseconds sensitivity is on most measures relatively similar to the baseline case of 2+ race participants within the information horizon. The number of races is a bit smaller but they are more profitable on average, with the net effect that the overall profits measures and liquidity-harm measures are about 20-30% higher than in the baseline. The 3+ race participants within 1 millisecond sensitivity yields a latency arbitrage tax (all-volume) of 0.65, versus 0.42 in baseline, and yields an implied harm to the cost of liquidity of 30.7%, versus 16.7% in baseline. In this sense, our baseline specification is meaningfully more conservative than the requirement of 3+ participants within 1 millisecond.

Now turn to the sensitivity for 5+ participants (Table C.3). There are very few (38) races per FTSE 100 symbol per day within the information horizon, versus 537 in the baseline and 229 with 3+. That said, these few races are quite profitable: they are about twice as profitable per share and more than three times as profitable in GBP per race as in the baseline. Increasing the race horizon to 500 microseconds yields 122 races per FTSE 100 symbol per day, and to 1 millisecond yields 202 races per day, again with races that are significantly more profitable per race than in the baseline. As a consequence, the sensitivity for 5+ participants within 500 microseconds yields overall profits that are about 60% of the baseline, and the sensitivity for 5+ participants within 1 millisecond yields overall profits and harm to liquidity that are just about the same as in the baseline.

We also include for completeness a sensitivity requiring 2+ unique firms as opposed to our baseline requirement of 2+ unique participants (Table C.4). As mentioned in the main text, some

firms use different UserIDs for different trading desks. Typically, this will be the case when the trading desks are operated sufficiently separately that if they happen to trade with each other the firm would not be in violation of wash-trade requirements. This economic separation is the reason why our baseline uses UserIDs as the measurement of the number of participants. This sensitivity reduces the number of races and various profits measures by about 10%.

Table C.2: Sensitivity Analysis: 3+ Race Participants

Measure	Baseline	3+ Race Participants Within								
		InfoHor	50 μ s	100 μ s	200 μ s	500 μ s	1ms	2ms	3ms	
Frequency and Duration of Races										
Races per day										
FTSE 100 - per symbol	537.24	228.98	86.73	134.38	236.14	482.47	585.98	655.57	685.98	
FTSE 250 - per symbol	70.05	30.68	13.40	19.92	32.98	67.54	82.20	90.01	93.47	
Mean race duration (μ s)	78.65	77.56	14.26	28.54	75.54	194.56	305.95	449.76	553.83	
% of races with wrong winner	4.30	5.08	10.66	7.88	4.59	2.01	1.71	1.44	1.33	
% of volume in races										
FTSE 100	22.15	12.75	3.84	6.32	11.57	27.97	35.53	39.99	41.83	
FTSE 250	16.90	9.33	3.38	5.24	9.35	23.68	29.59	32.95	34.39	
Full Sample	21.46	12.30	3.78	6.17	11.28	27.40	34.74	39.06	40.85	
Mean number of messages within 500 μ s	3.46	4.68	4.83	4.82	4.62	4.21	3.58	3.28	3.17	
Per-Race Profits										
Per-share profits										
ticks	0.55	0.71	0.73	0.71	0.64	0.61	0.63	0.64	0.65	
GBX	0.17	0.23	0.23	0.22	0.20	0.19	0.19	0.19	0.19	
basis points	1.66	2.24	2.36	2.29	2.03	1.90	1.91	1.92	1.92	
Per-race profits GBP	1.85	2.98	2.55	2.60	2.43	2.52	2.57	2.58	2.58	
displayed depth	1.76	2.87	2.29	2.40	2.33	2.55	2.62	2.65	2.64	
qty trade/cancel										
Aggregate Profits and LA Tax										
Daily Profits										
FTSE 100 - per symbol	1,047	736	238	375	612	1,273	1,583	1,770	1,848	
FTSE 250 - per symbol	108	70	27	41	65	147	181	200	208	
Full Sample - aggregate	132,378	91,506	30,701	47,980	77,738	164,760	204,272	228,064	237,757	
Latency Arbitrage Tax, All Volume (bps)										
FTSE 100	0.38	0.27	0.09	0.14	0.23	0.47	0.58	0.65	0.68	
FTSE 250	0.66	0.43	0.17	0.25	0.40	0.90	1.11	1.23	1.28	
Full Sample	0.42	0.29	0.10	0.15	0.25	0.52	0.65	0.72	0.75	
Latency Arbitrage Tax, Non-Race Volume (bps)										
FTSE 100	0.49	0.35	0.10	0.16	0.28	0.75	1.03	1.23	1.32	
FTSE 250	0.80	0.51	0.18	0.28	0.48	1.36	1.81	2.11	2.24	
Full Sample	0.53	0.37	0.11	0.18	0.31	0.83	1.14	1.35	1.45	
Spread Decomposition										
Price impact in races / All price impact %	30.58	19.13	5.64	9.34	16.39	38.37	48.96	55.92	58.99	
Price impact in races / Effective spread %	32.82	20.54	6.05	10.03	17.61	41.17	52.54	60.01	63.31	
Loss avoidance / Effective spread %	0.19	0.18	0.07	0.13	0.27	0.63	1.09	1.50	1.65	
Implied Reduction in Cost of Liquidity										
% Reduction in liquidity cost										
FTSE 100 - by symbol	19.95	12.46	3.69	5.94	10.23	26.26	36.90	45.30	49.49	
FTSE 250 - by symbol	11.93	7.67	3.03	4.50	7.09	19.05	24.47	28.10	28.85	
Full Sample - by date	16.73	10.43	3.19	5.13	8.79	22.18	30.65	37.13	40.29	

Notes: For descriptions of the sensitivity scenarios please see the text of Section C.2. Descriptions of each of the items in this table can be found in the following table notes in Section 4. Races per day: Table 4.1. Mean race duration and % of races with wrong winner: Figure 4.1. % of Volume in Races: Table 4.2. Mean number of messages: Table 4.3. Per-race profits: Table 4.6. Aggregate profits: Table 4.8. Latency Arbitrage Tax: Table 4.9. Spread decomposition: Table 4.10. Implied Reduction in Cost of Liquidity: Table 4.12.

Table C.3: Sensitivity Analysis: 5+ Race Participants

Measure	Baseline	5+ Race Participants Within								
		InfoHor	50 μ s	100 μ s	200 μ s	500 μ s	1ms	2ms	3ms	
Frequency and Duration of Races										
Races per day										
FTSE 100 - per symbol	537.24	37.83	5.96	13.58	35.27	121.76	202.00	268.66	297.78	
FTSE 250 - per symbol	70.05	4.91	0.88	2.03	5.21	16.36	26.67	33.77	36.62	
Mean race duration (μ s)	78.65	73.23	11.14	23.94	61.66	170.05	304.84	469.80	582.24	
% of races with wrong winner	4.30	5.62	14.93	9.48	4.98	2.32	1.84	1.45	1.30	
% of volume in races										
FTSE 100	22.15	3.39	0.38	0.99	2.70	10.36	17.94	23.31	25.51	
FTSE 250	16.90	2.23	0.33	0.78	2.12	7.65	12.81	16.43	17.87	
Full Sample	21.46	3.24	0.37	0.97	2.62	10.01	17.27	22.41	24.52	
Mean number of messages within 500 μ s	3.46	7.04	7.37	7.37	7.06	6.23	4.79	4.11	3.90	
Per-Race Profits										
Per-share profits										
ticks	0.55	1.01	1.02	0.98	0.92	0.84	0.83	0.86	0.87	
GBX	0.17	0.34	0.29	0.31	0.30	0.28	0.27	0.27	0.27	
basis points	1.66	3.39	3.25	3.30	3.13	2.79	2.69	2.66	2.64	
Per-race profits GBP	1.85	6.30	4.52	5.14	5.01	4.89	4.82	4.63	4.58	
displayed depth	1.76	6.29	4.28	4.91	4.96	5.06	5.03	4.84	4.80	
qty trade/cancel										
Aggregate Profits and LA Tax										
Daily Profits										
FTSE 100 - per symbol	1,047	262	29	77	195	637	1,037	1,310	1,433	
FTSE 250 - per symbol	108	21	3	7	19	63	102	129	139	
Full Sample - aggregate	132,378	31,663	3,699	9,609	24,265	79,717	129,773	163,927	178,855	
Latency Arbitrage Tax, All Volume (bps)										
FTSE 100	0.38	0.10	0.01	0.03	0.07	0.24	0.38	0.48	0.53	
FTSE 250	0.66	0.13	0.02	0.04	0.12	0.38	0.62	0.79	0.85	
Full Sample	0.42	0.10	0.01	0.03	0.08	0.25	0.41	0.52	0.57	
Latency Arbitrage Tax, Non-Race Volume (bps)										
FTSE 100	0.49	0.12	0.01	0.03	0.09	0.38	0.68	0.91	1.03	
FTSE 250	0.80	0.16	0.02	0.05	0.14	0.58	1.02	1.36	1.50	
Full Sample	0.53	0.13	0.01	0.04	0.10	0.40	0.73	0.97	1.09	
Spread Decomposition										
Price impact in races / All price impact %	30.58	6.01	0.63	1.72	4.61	16.26	27.53	35.79	39.63	
Price impact in races / Effective spread %	32.82	6.46	0.68	1.85	4.96	17.46	29.56	38.43	42.55	
Loss avoidance / Effective spread %	0.19	0.06	0.01	0.03	0.10	0.42	0.92	1.38	1.61	
Implied Reduction in Cost of Liquidity										
% Reduction in liquidity cost										
FTSE 100 - by symbol	19.95	3.88	0.50	1.15	2.91	10.46	18.79	25.71	29.26	
FTSE 250 - by symbol	11.93	2.76	0.77	1.18	2.40	7.29	12.09	15.75	17.14	
Full Sample - by date	16.73	3.31	0.38	0.99	2.55	8.94	15.82	21.39	24.19	

Notes: Please see the notes and surrounding text for Table C.2. This table is identical except it conditions on 5+ participants in a race instead of 3+ participants.

Table C.4: Sensitivity Analysis: 2+ Participating Firms

Measure	Baseline	2+ Participating Firms Within								
		InfoHor	50 μ s	100 μ s	200 μ s	500 μ s	1ms	2ms	3ms	
Frequency and Duration of Races										
Races per day										
FTSE 100 - per symbol	537.24	479.32	247.25	332.99	462.39	736.14	818.92	871.31	891.49	
FTSE 250 - per symbol	70.05	60.44	32.74	43.39	59.97	102.41	116.32	122.08	124.89	
Mean race duration (μ s)	78.65	81.59	16.08	31.24	74.03	196.91	306.40	447.04	552.85	
% of races with wrong winner	4.30	4.67	9.46	7.22	4.57	2.05	1.73	1.48	1.38	
% of volume in races										
FTSE 100	22.15	20.08	8.15	11.51	17.62	35.79	42.20	45.86	47.26	
FTSE 250	16.90	15.19	6.89	9.54	14.37	31.57	36.83	39.52	40.63	
Full Sample	21.46	19.44	7.98	11.25	17.19	35.23	41.49	45.02	46.39	
Mean number of messages within 500 μ s	3.46	3.52	3.52	3.54	3.58	3.47	3.08	2.90	2.83	
Per-Race Profits										
Per-share profits										
ticks	0.55	0.54	0.51	0.50	0.50	0.53	0.56	0.57	0.58	
GBX	0.17	0.17	0.16	0.16	0.16	0.16	0.17	0.17	0.17	
basis points	1.66	1.65	1.62	1.58	1.56	1.62	1.67	1.69	1.70	
Per-race profits GBP										
displayed depth	1.85	1.93	1.58	1.60	1.65	1.94	2.04	2.08	2.09	
qty trade/cancel	1.76	1.83	1.40	1.45	1.56	1.94	2.05	2.10	2.12	
Aggregate Profits and LA Tax										
Daily Profits										
FTSE 100 - per symbol	1,047	971	404	553	793	1,482	1,744	1,889	1,945	
FTSE 250 - per symbol	108	98	46	62	87	176	205	221	226	
Full Sample - aggregate	132,378	122,218	52,221	70,992	101,416	192,912	226,603	245,049	252,001	
Latency Arbitrage Tax, All Volume (bps)										
FTSE 100	0.38	0.36	0.15	0.20	0.29	0.54	0.64	0.69	0.71	
FTSE 250	0.66	0.60	0.29	0.38	0.53	1.08	1.26	1.35	1.39	
Full Sample	0.42	0.39	0.17	0.23	0.32	0.61	0.72	0.77	0.80	
Latency Arbitrage Tax, Non-Race Volume (bps)										
FTSE 100	0.49	0.46	0.17	0.24	0.36	0.87	1.13	1.31	1.39	
FTSE 250	0.80	0.72	0.31	0.43	0.64	1.63	2.06	2.32	2.43	
Full Sample	0.53	0.49	0.18	0.26	0.40	0.97	1.26	1.45	1.53	
Spread Decomposition										
Price impact in races / All price impact %	30.58	28.12	10.66	15.27	23.26	48.00	57.26	62.91	65.21	
Price impact in races / Effective spread %	32.82	30.18	11.43	16.38	24.97	51.50	61.44	67.50	69.97	
Loss avoidance / Effective spread %	0.19	0.19	0.07	0.13	0.26	0.53	0.94	1.32	1.48	
Implied Reduction in Cost of Liquidity										
% Reduction in liquidity cost										
FTSE 100 - by symbol	19.95	18.15	6.55	9.30	14.28	34.09	45.35	53.50	57.24	
FTSE 250 - by symbol	11.93	10.49	4.97	6.68	9.52	22.75	28.20	30.26	31.44	
Full Sample - by date	16.73	15.15	5.66	7.99	12.18	28.90	37.77	44.04	46.86	

Notes: Please see the description in the text above this table for a description. The table is identical to Table C.1 except it conditions on 2+ unique firms in a race whereas the baseline conditions on 2+ unique participants.

C.3 Sensitivity to Requiring Cancels or Multiple Takes

Our baseline method requires that of the 2+ messages in a race, at least 1 is aggressive. Thus, a race could have 1+ aggressive messages and 1+ cancel messages, or it could have 2+ aggressive messages and 0 cancel messages. Table C.5 presents sensitivity analysis for these requirements. In the first set of columns after the baseline, we require 1+ cancel message and 1+ aggressive message, i.e., exclude races with 0 cancels (and hence 2+ aggressive messages). In the second set of columns, we require 2+ aggressive messages, i.e., exclude races with exactly 1 aggressive message (and hence 1+ cancel messages).

Focus first on the 1+ cancel within information horizon column. Requiring a cancel attempt within the race horizon window reduces the number of races significantly, from 537 to 173 per day for the average symbol in the FTSE 100. These races are also less profitable on average. This reduction in profitability is driven by races with exactly 1 aggressive message. If we require 2+ aggressive messages alongside a cancel, profits per race are higher than in the baseline, especially in GBP per race where profits are nearly double.

Looking across the different race horizons does not change this picture much. The number of races goes up with the race horizon, as before, but the number of races and overall profitability are meaningfully smaller than without the 1+ cancel requirement, at all horizons. This pattern is consistent with our findings in Section 4 of the main text that most message activity in races is take attempts and most races are won by takers.

If we require at least 1 cancel within the information horizon, in addition to our other baseline race requirements, the harm to liquidity and the latency-arbitrage tax are each about 30% of baseline. That said, if we consider races with 1+ cancel within 3 milliseconds the results are closer to baseline, at about 85% of the harm to liquidity and level of latency-arbitrage tax.

Now focus on the columns that require at least 2 aggressive messages; that is, a race must have 2+ takes, along with 1+ success and 1+ fail, within the race horizon. Relative to the baseline, this excludes races with exactly 1 take and with 1+ cancels, which as we just discussed are relatively unprofitable. The number of races with 2+ takes within the information horizon is 424 for FTSE 100 symbols, versus 537 under the baseline scenario, a reduction of about 20%. These races are more profitable on average than the baseline races, so the net effect on profits and the harm-to-liquidity measures is smaller, roughly 10-15%. This magnitude of reduction relative to the baseline requirements persists across the other time horizons.

These overall patterns, as discussed in detail in the main text, are consistent with equilibria of the BCS model in which many of the fastest traders primarily engage in sniping as opposed to liquidity provision, and significant liquidity is provided by market participants not at the cutting edge of speed.

Table C.5: Sensitivity Analysis: Requiring Cancels or Multiple Takes

Measure	Baseline			1 + Cancel Within			2 + Takes Within			
	InfoHor	50 μ s	500 μ s	1ms	3ms	InfoHor	50 μ s	500 μ s	1ms	3ms
Frequency and Duration of Races										
Races per day										
FTSE 100 - per symbol	537.24	172.70	242.59	303.68	380.88	423.86	241.67	695.44	774.40	851.41
FTSE 250 - per symbol	70.05	14.40	23.42	31.10	40.89	60.91	36.30	103.85	117.33	127.31
Mean race duration (μ s)	78.65	92.77	206.89	373.48	768.59	74.52	15.42	194.72	300.04	547.27
% of races with wrong winner	4.30	3.15	2.01	1.50	0.99	4.63	8.34	1.89	1.62	1.30
% of volume in races										
FTSE 100	22.15	8.49	12.71	17.30	22.75	17.40	8.39	33.90	40.55	46.15
FTSE 250	16.90	3.31	6.17	9.21	13.02	15.20	7.65	32.02	37.04	41.19
Full Sample	21.46	7.82	11.87	16.26	21.49	17.11	8.28	33.64	40.08	45.49
Mean number of messages within 500 μ s	3.46	3.36	3.65	3.09	2.73	3.66	3.65	3.50	3.11	2.85
Per-Race Profits										
Per-share profits										
ticks	0.55	0.37	0.37	0.39	0.40	0.62	0.62	0.57	0.59	0.62
GBX	0.17	0.11	0.07	0.12	0.12	0.19	0.19	0.17	0.18	0.18
basis points	1.66	0.99	0.70	1.03	1.14	1.92	1.93	1.75	1.79	1.82
Per-race profits GBP										
displayed depth	1.85	1.92	1.18	2.14	2.24	2.03	1.74	1.96	2.08	2.19
qty trade/cancel	1.76	1.82	0.95	2.07	2.19	1.92	1.54	1.97	2.11	2.24
Aggregate Profits and LA Tax										
Daily Profits										
FTSE 100 - per symbol	1,047	361	505	705	917	907	441	1,418	1,690	1,968
FTSE 250 - per symbol	108	15	28	44	64	104	55	181	209	233
Full Sample - aggregate	132,378	40,205	57,933	81,993	108,273	117,054	57,996	187,719	222,151	256,194
Latency Arbitrage Tax, All Volume (bps)										
FTSE 100	0.38	0.13	0.19	0.26	0.34	0.33	0.16	0.52	0.62	0.72
FTSE 250	0.66	0.10	0.18	0.27	0.39	0.63	0.34	1.11	1.28	1.43
Full Sample	0.42	0.13	0.19	0.26	0.35	0.37	0.18	0.59	0.70	0.81
Latency Arbitrage Tax, Non-Race Volume (bps)										
FTSE 100	0.49	0.17	0.30	0.46	0.66	0.43	0.18	0.83	1.10	1.40
FTSE 250	0.80	0.11	0.27	0.45	0.69	0.76	0.37	1.67	2.10	2.52
Full Sample	0.53	0.16	0.30	0.46	0.66	0.47	0.20	0.94	1.23	1.56
Spread Decomposition										
Price impact in races / All price impact %	30.58	11.86	16.95	23.55	31.72	24.14	11.02	44.66	53.98	63.58
Price impact in races / Effective spread %	32.82	12.73	18.20	25.28	34.05	25.91	11.83	47.92	57.92	68.22
Loss avoidance / Effective spread %	0.19	0.19	0.53	0.94	1.48	0.16	0.06	0.59	1.09	1.76
Implied Reduction in Cost of Liquidity										
% Reduction in liquidity cost										
FTSE 100 - by symbol	19.95	5.41	8.17	12.44	17.83	16.24	7.17	31.12	41.37	54.89
FTSE 250 - by symbol	11.93	1.57	2.85	4.43	6.60	11.32	5.94	24.11	27.57	32.13
Full Sample - by date	16.73	4.49	6.80	10.21	14.63	13.80	6.23	26.82	35.18	45.82

Notes: For descriptions of the sensitivity scenarios please see the text of Section C.3. Descriptions of each of the items in this table can be found in the following table notes in Section 4. Races per day: Table 4.1. Mean race duration and % of races with wrong winner: Figure 4.1. % of Volume in Races: Table 4.2. Mean number of messages: Table 4.3. Per-race profits: Table 4.6. Aggregate profits: Table 4.8. Latency Arbitrage Tax: Table 4.9. Spread decomposition: Table 4.10. Implied Reduction in Cost of Liquidity: Table 4.12.

C.4 Sensitivity to Varying the Definitions of Success and Fail

Our baseline method defined success and fail as follows. A take attempt succeeds if it executes at least in part, and otherwise fails. A cancel attempt succeeds if at least some of the order’s quantity is successfully canceled, and otherwise fails. As discussed in Section 3.3, while the definition of success might sound quite loose — e.g., if there are 10,000 shares in the book, an attempt to take 10,000 shares that “succeeds” in taking just 100 shares is counted as a success — it has some real bite in conjunction with the requirement that a race has a fail, because someone else likely got or canceled the other 9,900 shares, for there then to be yet another participant who then fails to get anything or cancel anything. The exception is if there is a successful take attempt for a small amount (e.g., the order is for just 100 shares) followed by a cancel attempt for a small amount (e.g., 100 shares) where, by coincidence, the cancel fails because it was that user’s 100 shares that just got taken. Thus, to deal with this possibility, our first sensitivity imposes that 100% of the depth at the race level is cleared, either through takes or cancels. As can be seen this reduces the number of races by about 5% (from 537 to 514), and reduces our measures of aggregate profits, latency arbitrage tax, and harm to liquidity by about 5% as well.

For our definition of fail, the concern we mentioned in Section 3.3 is that we count limit orders that post to the book as a fail. A worry, especially at longer race horizons, is that we are picking up as “latency arbitrage races” cases where the “fail” is in fact simply a participant posting new liquidity at a new price, using a plain vanilla limit order, at a price that happened to be the price of the last successful trade. As a sensitivity, therefore, we only allow failed IOCs and failed cancels to count as fails.⁶ That is, we do not allow ordinary limit orders to count as fails, even though some participants may in fact use them in latency arbitrage races, because of the fee advantage described earlier in the main text.

In the baseline, the strict fail criterion only reduces the number of races detected by about 8% (from 537 to 494), and race profits by about 5%. At longer horizons, as expected, the strict fails criterion reduces the number of races detected, and overall race profits, by larger amounts—for instance, at 3ms, the reduction in the number of races is about 15% (from 946 to 800) and the reduction in total profits is about 10% (from 255,000 per day to 232,000 per day). This makes sense because at longer horizons we should be more concerned about mistaking limit orders that post to the book as failed race attempts. For this reason, when we consider what the sensitivity analyses suggest about upper bounds on race profits in Table 5.1 in the main text, when we use longer race horizons we always do so in conjunction with the strict fail requirement.

⁶Note as well that this sensitivity has the interpretation of only allowing as fails the “error messages”—failed IOCs and failed cancel attempts—that are unique to our message data relative to ordinary limit-order book data.

Table C.6: Sensitivity Analysis: Definitions of Success and Fail

Measure	Success:		Strict Fail			
	Baseline	All Cleared	50 μ s	500 μ s		
Frequency and Duration of Races						
Races per day					1ms	3ms
FTSE 100 - per symbol	537.24	514.13	266.32	719.71	768.02	799.91
FTSE 250 - per symbol	70.05	68.35	38.08	105.09	115.94	123.01
Mean race duration (microseconds)	78.65	76.57	15.87	195.83	294.52	509.37
% of races with wrong winner	4.30	4.27	8.79	2.07	1.79	1.47
% of volume in races						
FTSE 100	22.15	20.69	8.83	35.33	40.54	44.27
FTSE 250	16.90	16.46	7.87	32.16	36.79	40.17
Full Sample	21.46	20.14	8.70	34.89	40.03	43.72
Mean number of messages within 500 μ s	3.46	3.47	3.54	3.48	3.15	2.94
Per-Race Profits						
Per-share profits						
ticks	0.55	0.55	0.52	0.52	0.54	0.55
GBX	0.17	0.17	0.16	0.16	0.16	0.16
basis points	1.66	1.69	1.64	1.59	1.64	1.64
Per-race profits GBP						
displayed depth	1.85	1.86	1.57	1.91	2.04	2.09
qty trade/cancel	1.76	1.78	1.38	1.92	2.07	2.16
Aggregate Profits and LA Tax						
Daily Profits						
FTSE 100 - per symbol	1,047	1,003	437	1,437	1,650	1,783
FTSE 250 - per symbol	108	108	53	174	200	213
Full Sample - aggregate	132,378	127,928	57,048	187,989	215,794	232,457
Latency Arbitrage Tax, All Volume (bps)						
FTSE 100	0.38	0.37	0.16	0.53	0.61	0.65
FTSE 250	0.66	0.66	0.32	1.07	1.23	1.31
Full Sample	0.42	0.40	0.18	0.60	0.68	0.74
Latency Arbitrage Tax, Non-Race Volume (bps)						
FTSE 100	0.49	0.47	0.18	0.82	1.02	1.18
FTSE 250	0.80	0.80	0.35	1.59	1.95	2.20
Full Sample	0.53	0.52	0.20	0.92	1.14	1.31
Spread Decomposition						
Price impact in races / All price impact %	30.58	28.85	11.70	47.38	55.27	61.61
Price impact in races / Effective spread %	32.82	30.97	12.56	50.84	59.31	66.11
Loss avoidance / Effective spread %	0.19	0.18	0.06	0.62	1.07	1.51
Implied Reduction in Cost of Liquidity						
% Reduction in liquidity cost						
FTSE 100 - by symbol	19.95	18.63	7.24	33.19	42.45	50.59
FTSE 250 - by symbol	11.93	11.82	5.64	23.12	27.56	29.90
Full Sample - by date	16.73	15.86	6.25	28.12	35.37	41.64

Notes: For descriptions of the sensitivity scenarios please see the text of Section C.4. Descriptions of each of the items in this table can be found in the following table notes in Section 4: Races per day: Table 4.1. Mean race duration and % of races with wrong winner: Figure 4.1. % of Volume in Races: Table 4.2. Mean number of messages: Table 4.3. Per-race profits: Table 4.6. Aggregate profits: Table 4.8. Latency Arbitrage Tax: Table 4.9. Spread decomposition: Table 4.10. Implied Reduction in Cost of Liquidity: Table 4.12.

D Supporting Details for Section 5.2: Additional Robustness Checks

D.1 Races with Negative Profits Ex-Post

The following table provides additional data to support the discussion in Section 5.2 of the main text on races with negative profits. Each column represents a race scenario: the first column is our baseline race specification; the next set of columns is for 2+ participants over a wide range of fixed time horizons; the next set of columns is for 3+ participants, over a wide range of time horizons; and the last set of columns is for 5+ participants over a wide range of time horizons.

For each specification, we report data on the proportion of races with strictly positive, zero, and strictly negative profits at different time horizons. For example, at 1 second, our baseline method has 47.2% of races with strictly positive profits, 31.2% of races with zero profits, and 21.6% of races with strictly negative profits. We then report the proportion of races that have strictly negative profits continuously throughout the specified time interval. For example, in our baseline method, 10.5% of races have strictly negative profits continuously for 1 second and 7.5% of races have strictly negative profits continuously for 10 seconds. We then report data for price impact. For example, in our baseline, at 1 second, 81.5% of races have strictly positive price impact, 15.4% have zero price impact, and 3.1% have strictly negative price impact.

We have two main takeaways from these data. First, in our main specifications (baseline and others we emphasize in the text) there is a meaningful proportion of races with negative profits. Most of the time, these are cases where price impact is weakly positive but not enough to cover the aggressor's half-spread. Second, even in our most strenuous sensitivity analysis, some races have strictly negative profits.

Overall, trading in races seems to lie comfortably between the two extremes of 100/0 pure arbitrage and 51/49 tiny statistical edges.

Table D.1: Sign of Race Profits and Price Impact

Panel A: 2+ Participants

		2+, IH	2+, 50 μ s	2+, 100 μ s	2+, 500 μ s	2+, 1ms	2+, 3ms
% Races with Profits							
1ms	>0	32.95	36.21	34.94	31.42	30.98	27.76
	=0	44.94	44.70	44.65	46.07	44.52	41.48
	<0	22.11	19.09	20.41	22.51	24.50	30.76
10ms	>0	41.23	42.71	41.78	40.26	41.25	42.41
	=0	39.45	39.68	39.62	39.67	38.58	37.20
	<0	19.32	17.61	18.60	20.07	20.17	20.39
100ms	>0	43.94	44.71	44.04	43.03	43.87	44.81
	=0	36.20	36.78	36.58	36.28	35.31	34.15
	<0	19.86	18.51	19.38	20.70	20.81	21.04
1s	>0	47.20	47.68	47.15	46.31	46.89	47.44
	=0	31.15	31.85	31.60	31.23	30.50	29.62
	<0	21.64	20.46	21.25	22.47	22.61	22.94
10s	>0	50.02	50.38	49.99	49.44	49.68	49.78
	=0	20.95	21.46	21.27	20.97	20.61	20.08
	<0	29.02	28.15	28.74	29.59	29.71	30.14
Always <0 for 10ms		15.52	13.97	14.89	15.98	15.87	16.53
Always <0 for 100ms		12.90	11.74	12.50	13.43	13.28	13.74
Always <0 for 1s		10.53	9.63	10.23	11.06	10.93	11.34
Always <0 for 10s		7.52	6.85	7.29	7.94	7.84	8.19
% Races with PI							
10ms	>0	86.08	85.00	84.40	88.06	88.21	87.74
	=0	13.24	14.35	14.94	11.29	11.17	11.67
	<0	0.67	0.65	0.67	0.65	0.62	0.59
100ms	>0	84.26	83.18	82.58	86.03	86.15	85.68
	=0	14.29	15.40	15.98	12.55	12.49	13.04
	<0	1.45	1.42	1.43	1.42	1.36	1.28
1s	>0	81.50	80.53	79.91	83.16	83.25	82.74
	=0	15.42	16.48	17.07	13.77	13.76	14.38
	<0	3.07	2.99	3.02	3.07	2.99	2.88
10s	>0	72.31	71.46	70.90	73.81	73.91	73.38
	=0	17.14	18.24	18.80	15.45	15.49	16.16
	<0	10.55	10.30	10.30	10.74	10.60	10.46

Notes: This table reports the proportion of races with strictly positive, zero, and strictly negative profits and price impact at different time horizons. The race scenario is indicated in the column. For computational tractability, we define race profits to be always <0 for T time if profits are strictly negative at each round-number interval up to and including T: for example, “Always <0 for 1s” is defined as true if profits are strictly negative at 1ms, 10ms, 100ms and 1s.

Panel B: 3+ Participants

		3+, IH	3+, 50 μ s	3+, 100 μ s	3+, 500 μ s	3+, 1ms	3+, 3ms
% Races with Profits							
1ms	>0	38.77	43.21	41.49	34.27	32.79	28.85
	=0	42.04	41.25	41.23	45.47	43.64	39.17
	<0	19.20	15.54	17.28	20.25	23.57	31.98
10ms	>0	48.52	50.63	49.32	44.11	44.60	45.94
	=0	35.94	35.67	35.72	38.41	37.46	35.75
	<0	15.54	13.70	14.96	17.47	17.94	18.31
100ms	>0	51.03	52.45	51.37	46.76	47.14	48.22
	=0	33.04	33.16	33.12	35.23	34.36	32.93
	<0	15.94	14.39	15.52	18.01	18.49	18.85
1s	>0	53.71	54.82	53.93	49.74	49.93	50.56
	=0	28.76	29.07	28.95	30.56	29.87	28.77
	<0	17.53	16.12	17.12	19.70	20.21	20.67
10s	>0	54.63	55.45	54.85	51.84	51.79	51.90
	=0	20.31	20.71	20.57	21.06	20.62	19.91
	<0	25.06	23.84	24.58	27.09	27.59	28.20
Always <0 for 10ms		12.20	10.59	11.74	13.66	13.91	14.77
Always <0 for 100ms		9.89	8.71	9.65	11.34	11.46	12.03
Always <0 for 1s		7.99	7.11	7.83	9.26	9.35	9.82
Always <0 for 10s		5.66	4.99	5.52	6.61	6.67	7.02
% Races with PI							
10ms	>0	91.33	91.20	90.30	90.94	90.71	90.06
	=0	8.12	8.29	9.16	8.50	8.74	9.40
	<0	0.54	0.52	0.53	0.56	0.55	0.54
100ms	>0	89.71	89.58	88.70	89.09	88.80	88.16
	=0	9.07	9.24	10.11	9.66	9.96	10.65
	<0	1.21	1.19	1.19	1.26	1.24	1.19
1s	>0	87.27	87.19	86.31	86.40	86.09	85.41
	=0	10.20	10.32	11.21	10.88	11.21	11.95
	<0	2.53	2.49	2.48	2.71	2.71	2.64
10s	>0	78.23	78.28	77.49	77.08	76.72	75.97
	=0	12.50	12.63	13.53	13.01	13.31	14.07
	<0	9.27	9.10	8.98	9.91	9.97	9.96

Panel C: 5+ Participants

		5+, IH	5+, 50 μ s	5+, 100 μ s	5+, 500 μ s	5+, 1ms	5+, 3ms
% Races with Profits							
1ms	>0	49.48	56.42	53.97	44.31	40.08	32.65
	=0	33.97	31.71	31.84	39.08	37.69	31.99
	<0	16.55	11.87	14.18	16.61	22.24	35.37
10ms	>0	61.70	64.47	62.66	55.42	54.69	55.61
	=0	27.32	26.17	26.46	31.67	31.30	29.67
	<0	10.98	9.36	10.89	12.91	14.01	14.72
100ms	>0	63.40	65.01	63.75	57.43	56.71	57.29
	=0	25.35	25.00	25.00	29.21	28.86	27.57
	<0	11.25	9.99	11.26	13.36	14.42	15.14
1s	>0	65.12	65.96	64.95	59.55	58.67	58.76
	=0	22.47	22.56	22.61	25.68	25.49	24.54
	<0	12.41	11.48	12.45	14.77	15.84	16.70
10s	>0	62.95	62.94	62.55	58.86	57.96	57.45
	=0	17.58	17.60	17.75	19.08	18.81	18.11
	<0	19.48	19.46	19.70	22.06	23.22	24.44
Always <0 for 10ms		8.34	6.99	8.28	9.69	10.49	11.89
Always <0 for 100ms		6.52	5.58	6.61	7.74	8.29	9.23
Always <0 for 1s		5.15	4.48	5.27	6.19	6.61	7.36
Always <0 for 10s		3.60	3.14	3.70	4.36	4.65	5.19
% Races with PI							
10ms	>0	95.13	95.22	94.58	93.96	93.18	92.59
	=0	4.54	4.41	5.06	5.63	6.39	7.00
	<0	0.33	0.37	0.36	0.42	0.43	0.41
100ms	>0	93.87	93.89	93.28	92.47	91.62	91.01
	=0	5.33	5.16	5.86	6.56	7.38	8.04
	<0	0.80	0.95	0.86	0.97	0.99	0.95
1s	>0	92.08	91.94	91.42	90.34	89.43	88.73
	=0	6.23	5.98	6.75	7.67	8.49	9.21
	<0	1.69	2.08	1.83	2.00	2.08	2.06
10s	>0	84.07	83.20	83.24	81.79	80.72	79.70
	=0	9.09	8.69	9.57	10.40	11.19	11.93
	<0	6.84	8.11	7.19	7.81	8.09	8.37

D.2 Races Triggered by Order Book Activity

The following table provides additional data to support the discussion in Section 5.2 of the main text on races triggered by order book activity. Each column represents a window of time before the race starts. For each time window we report what percentage of races had a change in the race side best bid or best offer — e.g., for the time window 100 microseconds, if the race was at the offer we check whether the best offer changed in the 100 microseconds leading up to the race time. We then compute whether the change was price improving (a lower offer or a higher bid) or price worsening (the reverse), and provide statistics for each category.

Our takeaway is that a meaningful proportion of races have changes in the race-side best bid or best offer just leading up to the race: 14% in the 100 microseconds before the race, and 21% in the 500 microseconds before the race. When there are changes, they tend to be price improving and tend to have fewer cancelations, both of which seem consistent with the theories of Li, Wang and Ye (2020) and Foucault, Kozhan and Tham (2016).

We also find that changes in the race-side best bid or offer are comparatively more likely than changes in the non-race side best bid or offer: in the 100 microseconds before the race, a change in the race-side price occurs 14% of the time, versus 4% for the non-race-side price. This also seems to support the story in Li, Wang and Ye (2020) and Foucault, Kozhan and Tham (2016) being a feature of the data.

Still, the large majority of races have stable prices leading up to the race. This suggests that most races are triggered by some public signal external to the symbol's own order book, as in the model of Budish, Cramton and Shim (2015).

Table D.2: Races Triggered by Order Book Activity

Measure	Check for Pre-Race Order Book Price Changes Within				
	10 μ s	50 μ s	100 μ s	500 μ s	1ms
All Races - FTSE 100					
% Stable Pre-Race Price	98.65	91.12	86.46	78.69	76.01
% Change in Pre-Race Price	1.35	8.88	13.54	21.31	23.99
Of Races with Changes in Pre-Race Price					
% Price Improves	96.47	96.28	89.06	78.10	71.16
% Price Worsens	3.53	3.72	10.94	21.90	28.84
Number of Cancels within 500 μ s					
Stable Pre-Race Price	0.40	0.41	0.40	0.38	0.37
Change in Pre-Race Price	0.49	0.32	0.36	0.46	0.48
Price Improves	0.46	0.28	0.24	0.25	0.27
Price Worsens	1.45	1.36	1.24	1.12	0.96
% Liquidity Provided by Non-Top 6 Firms					
Stable Pre-Race Price	57.58	57.20	57.06	56.78	56.85
Change in Pre-Race Price	83.76	70.40	66.73	64.38	62.77
Price Improves	88.47	72.49	70.63	69.45	67.35
Price Worsens	40.43	39.06	43.43	49.18	52.91
Race Duration (μ s)					
Stable Pre-Race Price	79.17	83.30	85.01	86.03	86.09
Change in Pre-Race Price	45.63	34.49	41.71	53.41	57.09
Price Improves	44.52	32.79	38.12	43.99	45.80
Price Worsens	77.15	75.94	68.39	83.50	82.07
Effective Spread (bps)					
Stable Pre-Race Price	3.19	3.23	3.23	3.22	3.18
Change in Pre-Race Price	3.21	2.80	2.96	3.10	3.23
Price Improves	3.10	2.70	2.67	2.65	2.75
Price Worsens	6.36	6.03	5.80	4.78	4.43
Per-Share Profits (ticks)					
Stable Pre-Race Price	0.55	0.54	0.53	0.54	0.54
Change in Pre-Race Price	0.59	0.64	0.63	0.59	0.58
Price Improves	0.55	0.62	0.62	0.56	0.53
Price Worsens	1.49	1.07	0.70	0.69	0.68

Notes: This table groups races detected by our baseline method (see Section 3 of the main text for detailed description) into the following categories according to the pre-race order book activity and report statistics for each category: races with stable race-side BBO price pre-race and with changes in race-side BBO price pre-race. The latter is further divided into races with strictly improving and weakly worsening pre-race price. Descriptions of each of the items in this table can be found in the following table notes in Section 4. Number of cancels: Table 4.3. Proportion of liquidity provided by non-top 6 firms: Figure 4.3. Race duration: Figure 4.1. Effective spread: 4.10. Per-share profits: 4.6.

E Additional Results for Section 6: Magnitudes

E.1 Additional Extrapolation Models

Table 6.1 in the main text presents regressions of daily latency arbitrage profits on volume and 1-minute realized volatility. These regressions were used for the purpose of out-of-sample extrapolation in Section 6. The following appendix table presents analogous regressions using additional volatility variables, as was discussed in the main text. Columns (1)-(4) are analogous to Columns (3)-(6) in Table 6.1, but using 5-minute realized volatility instead of 1-minute realized volatility. Columns (5)-(8) are analogous to the same columns in Table 6.1, but using midpoint distance traveled (Budish, Cramton and Shim, 2015) as the volatility measure. As discussed in the main text, the fit is worse with 5-minute realized volatility than with 1-minute realized volatility, and is slightly better with midpoint distance traveled. We nevertheless utilize 1-minute realized volatility in the main text since it is more easily interpreted, and its measurement does not depend on the number of significant digits of the trading index (or the tick size if using a futures contract price for the index) in the way that distance traveled does.

Table E.1: Extrapolation Models (Appendix)

	<i>Dependent variable:</i>							
	Latency Arbitrage Profits (GBP)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Volume (10,000 GBP)			0.4237*** (0.0583)	0.4123*** (0.0320)			0.2561*** (0.0790)	0.2833*** (0.0578)
Volatility (5 min) * Average Volume	0.0147*** (0.0020)	0.0276*** (0.0013)	0.0004 (0.0024)	0.0006 (0.0022)				
Volatility (Midpoint Distance Travelled) * Average Volume					0.0072*** (0.0006)	0.0090*** (0.0002)	0.0032*** (0.0013)	0.0030*** (0.0012)
Constant	68,085*** (9,796)		-2,768 (11,717)		28,891*** (8,771)		5,464 (10,709)	
Observations	43	43	43	43	43	43	43	43
R ²	0.561	0.134	0.811	0.811	0.791	0.742	0.835	0.834

*p<0.1; **p<0.05; ***p<0.01

Notes: Please see the description in the text above this table and the notes for Table 6.1 in the main text. 5-minutes volatility is computed as realized 5-minute volatility for the FTSE 350 index in percentage points, using TRTH data. The distance travelled for each day is calculated as the sum of the absolute value of the change in midpoint on each update of the FTSE350. The FTSE350 is disseminated 4 times a second, or every 250 milliseconds.

E.2 Global Magnitudes for 2020

This table is analogous to Table 6.3 in the main text but using 2020 data instead of 2018 data. Please note that 2020 was an unusually high-volume and high-volatility year due to the Covid-19 pandemic.

Table E.2: **Annual Latency Arbitrage Profits in Global Equity Markets in 2020 (USD Millions)**

Exchange Group	(1) Volume- Volatility	(2) Volume- Only	(3) Low Scenario	(4) High Scenario
Nasdaq - U.S.	1,385	1,272	607	2,222
NYSE Group	1,201	1,103	526	1,926
BATS Global Markets - U.S.	1,171	1,075	513	1,878
Shenzhen Stock Exchange	747	764	364	1,334
Shanghai Stock Exchange	533	545	260	953
Japan Exchange Group	305	289	138	504
Korea Exchange	277	239	114	417
Hong Kong Exchanges and Clearing	142	141	67	246
Euronext	116	109	52	191
London Stock Exchange Group**	112	103	49	180
Deutsche Börse Group	104	98	47	171
TMX Group	88	84	40	147
National Stock Exchange of India	87	77	37	134
BATS Chi-X Europe	75	70	34	123
B3 - Brasil Bolsa Balcao	82	61	29	107
Global Total (WFE Data Universe)	6,957	6,529	3,114	11,404

**London Stock Exchange Group includes London Stock Exchange as well as Borsa Italiana

Note: Please see the notes for Table 6.3 in the main text. This table is the same except that we use the volume and volatility data in 2020. Trading volume is from the World Federation of Exchanges (2021). Volatility is computed using TRTH data for the following indices. NYSE, BATS and Nasdaq: S&P 500. Shenzhen and Shanghai: Shanghai composite. Japan: Nikkei225. Korea: KOSPI. Hong Kong: Hang Seng. Euronext, BATS Chi-X, Deutsche Börse: EuroStoxx600. LSE Group: FTSE 350. Canada TMX Group: TSX Composite. India: NIFTY 50. Brazil: BOVESPA.

F Theory Appendix

This theory appendix covers three topics. First, discussion of equilibrium in the case where the firm providing liquidity is slow. Second, the analysis behind the bid-ask spread decomposition (4.3). Third, the algebra in support of equation (4.6) and its empirical counterpart (4.7), which express the proportional reduction of the cost of liquidity if latency arbitrage were eliminated.

F.1 Equilibrium with Slow Liquidity Providers

In the equilibria of the continuous limit order book market studied in Budish, Cramton and Shim (2015), fast trading firms both engage in stale-quote sniping and provide all of the market's liquidity. There is a fringe of slow trading firms but they play no role in these equilibria (see especially Section VI.D and Proposition 3). The slow firms only play a role in equilibrium in Budish, Cramton and Shim (2015) under the frequent batch auctions market design.

In the BCS equilibria of the continuous market, fast trading firms are indifferent between liquidity provision and stale-quote sniping at the equilibrium bid-ask spread s^{CLOB} , characterized by

$$\lambda_{invest} \frac{s^{CLOB}}{2} = \lambda_{public} L\left(\frac{s^{CLOB}}{2}\right), \quad (\text{F.1})$$

where λ_{invest} denotes the arrival rate of investors (i.e., liquidity traders), λ_{public} denotes the arrival rate of new public information, and $L\left(\frac{s^{CLOB}}{2}\right) \equiv \Pr(J \geq \frac{s^{CLOB}}{2})\mathbb{E}(J - \frac{s^{CLOB}}{2} | J \geq \frac{s^{CLOB}}{2})$ denotes the expected loss to a liquidity provider if there is a jump larger than their half-spread and they get sniped (J is the random variable describing the absolute value of jump sizes). In the event of a jump larger than the half-spread, stale-quote snipers are successful $\frac{1}{N}$ of the time, where N is the number of fast trading firms, and hence earn expected profits of $\frac{1}{N}\lambda_{public}L\left(\frac{s^{CLOB}}{2}\right)$. A fast trading firm that provides liquidity earns revenues of $\lambda_{invest}\frac{s^{CLOB}}{2}$ from providing liquidity to investors, but, if there is a public jump, they get sniped with probability $\frac{N-1}{N}$, hence incurring costs of $\frac{N-1}{N}\lambda_{public}L\left(\frac{s^{CLOB}}{2}\right)$. At the equilibrium spread, the revenue benefits of liquidity provision less these sniping costs net to the same $\frac{1}{N}\lambda_{public}L\left(\frac{s^{CLOB}}{2}\right)$ earned by snipers. This net profit can be interpreted as the fast liquidity provider earning the opportunity cost of not sniping.

Under slightly different modeling formalities, introduced in Budish, Lee and Shim (2019), there also exist equilibria in which slow trading firms provide liquidity, at exactly the same bid-ask spread $\frac{s^{CLOB}}{2}$ characterized by (F.1), and the N fast trading firms all engage in stale-quote sniping. The economic intuition for why this can also be an equilibrium is as follows. First, at this bid-ask spread, slow trading firms earn zero profits from liquidity provision, so slow trading firms are indifferent between liquidity provision here, and doing nothing as before. Second, with all N fast trading firms now engaged in sniping, and the bid-ask spread the same as before, the fast trading firms all earn the same profits of $\frac{1}{N}\lambda_{public}L\left(\frac{s^{CLOB}}{2}\right)$ as before. And, as before, at this bid-ask spread the fast trading firms are indifferent between providing liquidity or being one of $N - 1$ snipers, so they do not strictly prefer to change from sniping to liquidity provision. Thus, the same bid-ask spread that leaves fast trading firms indifferent between liquidity provision and stale-quote sniping, and hence

can support equilibrium with fast trading firms engaged in liquidity provision, leaves slow trading firms indifferent between liquidity provision and not (i.e., with zero profits), and hence can support equilibrium with slow trading firms engaged in liquidity provision.

Formally, the configuration of play in which a slow trading firm provides liquidity at the spread characterized by (F.1) (or its slight generalization to include adverse selection as well, presented as equation (4.2) in the main text) is an Order Book Equilibrium as defined in Budish, Lee and Shim (2019). The argument that this play constitutes an Order Book Equilibrium is as follows:

- If the slow TF deviates by widening their spread to $s' > s^{CLOB}$: another TF (whether slow or fast) can profitably undercut the deviation by providing liquidity at a better spread. Order Book Equilibrium requires that any deviation be robust to another TF providing better liquidity in response, so this potential deviation does not violate Order Book Equilibrium.
- If the slow TF deviates by narrowing their spread to $s' < s^{CLOB}$: they earn strictly negative profits as opposed to zero profits, so this is not a profitable deviation.
- If a fast TF undercuts the slow TF's spread to $s' < s^{CLOB}$: this is a profitable unilateral deviation for a fast TF for s' close enough to s^{CLOB} , because the fast TF gets to both earn positive expected profits from liquidity provision, of just less than $\frac{1}{N}\lambda_{public}L(\frac{s^{CLOB}}{2})$, and potentially snipe the slow TF (the “have your cake and eat it too” deviation). However, the deviation is not robust to the slow TF canceling in response. Order Book Equilibrium requires that deviations are robust to other firms' responses with either cancels or price improvements (“no robust deviations”).⁷
- If any other slow TF undercuts to $s' < s^{CLOB}$: this is not a profitable unilateral deviation for slow TFs, because s^{CLOB} is the bid-ask spread at which slow TFs earn zero expected profits from liquidity provision. (The reason why providing liquidity at s' close enough to s^{CLOB} is profitable for a fast TF but not a slow TF is that fast TFs get sniped with probability $\frac{N-1}{N}$, whereas slow TFs get sniped with probability 1.)

Thus there exist order book equilibria in which fast TFs provide all liquidity as well as order book equilibria in which slow TFs provide all liquidity. It follows that there also exist order book equilibria in which, proportion $\rho_{fast} \in (0, 1)$ of the time, a fast TF provides liquidity at s^{CLOB} , while the remaining $1 - \rho_{fast}$ of the time a slow TF provides liquidity at s^{CLOB} . Either way, the spread is the same, the profits of all fast TFs are the same ($\frac{1}{N}\lambda_{public}L(\frac{s^{CLOB}}{2})$), and the profits of all slow TFs are zero.

⁷This case is the key technical difference between the modeling approach in Budish, Lee and Shim (2019) versus that in BCS. In the continuous-time game form considered in BCS a fast TF undercutting a slow TF in this way is a profitable deviation for the fast trading firm, because, in the small amount of time before a slow trading firm is able to respond to this deviation, the deviating fast trading firm both earns potential revenues from liquidity provision and earns potential profits from sniping the slow trading firm. In contrast, the Order Book Equilibrium concept introduced in Budish, Lee and Shim (2019) requires that the order book is at a resting point, where, if any one trading firm can profitably deviate from this resting point the deviation is no longer profitable after other trading firms respond with either price improvements or cancelations.

F.2 Support for Bid-Ask Spread Decomposition (4.3)

Equation (4.3) in the main text provides a novel bid-ask spread decomposition that includes Price Impact both in and out of races, as well as a Loss Avoidance term for the case where a liquidity provider successfully cancels in a race. In this section we provide formal support for this decomposition.

Begin with the bid-ask spread characterization presented in the main text as (4.2),

$$\lambda_{invest} \frac{s^{CLOB}}{2} = (\lambda_{public} + \lambda_{private}) \cdot L\left(\frac{s^{CLOB}}{2}\right),$$

where λ_{public} and $\lambda_{private}$ denote the arrival rate of public and private information, respectively, and $L(\frac{s^{CLOB}}{2})$ denotes the expected loss to a liquidity provider conditional on getting sniped or adversely selected. For simplicity, we assume that the jump size J is identically distributed for public and private information, and that all jumps are of size of at least the equilibrium half-spread $\frac{s^{CLOB}}{2}$, so all jumps generate attempts to trade. These assumptions can be relaxed but at considerable notational burden.⁸ With these assumptions, we have $L(\frac{s^{CLOB}}{2}) = E(J) - \frac{s^{CLOB}}{2}$.⁹

As discussed in the previous subsection, there exist equilibria in which only fast TFs provide liquidity, only slow TFs provide liquidity, and in which both fast and slow TFs provide liquidity. The former case was emphasized in BCS but the latter case appears to fit the data better. Let $\rho_{fast} \in [0, 1]$ denote the proportion of liquidity provided by fast TFs in equilibrium with the remaining $1 - \rho_{fast}$ provided by slow TFs. We can now formally define the terms utilized in equation (4.3).

- *EffectiveSpread* is equal to $[\lambda_{invest} + \lambda_{public}(1 - \frac{\rho_{fast}}{N}) + \lambda_{private}] \cdot \frac{s^{CLOB}}{2}$. Trade occurs whenever an investor arrives (at rate λ_{invest}), whenever an informed trader arrives ($\lambda_{private}$), and whenever there is public news (λ_{public}) and the race is won by a sniper: which occurs with probability $\frac{N-1}{N}$ if the TF providing liquidity is fast, where N is the number of fast traders, and probability 1 if the TF providing liquidity is slow, hence total probability of $\rho_{fast} \frac{N-1}{N} + (1 - \rho_{fast}) = 1 - \frac{\rho_{fast}}{N}$.
- *PriceImpact_{Race}* is equal to $\lambda_{public}(1 - \frac{\rho_{fast}}{N}) \cdot E(J)$: the $\lambda_{public}(1 - \frac{\rho_{fast}}{N})$ probability that a sniper wins a race, times the size of the jump $E(J)$, which will be the change in the midpoint. Using $L(\frac{s^{CLOB}}{2}) = E(J) - \frac{s^{CLOB}}{2}$ this can be rewritten as $\lambda_{public}(1 - \frac{\rho_{fast}}{N})E(J) = \lambda_{public}(1 - \frac{\rho_{fast}}{N})(\frac{s^{CLOB}}{2} + L(\frac{s^{CLOB}}{2}))$.
- *PriceImpact_{NonRace}*, by similar logic, is equal to $\lambda_{private}E(J)$: the $\lambda_{private}$ probability that

⁸Formally, if $J_{private}$ and J_{public} are, respectively, the jump distributions for private and public information, with cumulative distribution functions $F_{private}(x)$ and $F_{public}(x)$, respectively, then the conditional distributions of interest are $J_{private}^*$ and J_{public}^* with cdf's $F_{private}^*(x) = \frac{F_{private}(x) - F_{private}^-(\frac{s^{CLOB}}{2})}{1 - F_{private}^-(\frac{s^{CLOB}}{2})}$ and $F_{public}^*(x) = \frac{F_{public}(x) - F_{public}^-(\frac{s^{CLOB}}{2})}{1 - F_{public}^-(\frac{s^{CLOB}}{2})}$, respectively, for $x \geq \frac{s^{CLOB}}{2}$ and $F_{private}^*(x) = F_{public}^*(x) = 0$ for $x < \frac{s^{CLOB}}{2}$.

⁹In the generalization described in the previous footnote the appropriate formulas to use are $L_{private}(\frac{s^{CLOB}}{2}) \equiv E(J_{private}^*) - \frac{s^{CLOB}}{2}$ and $L_{public}(\frac{s^{CLOB}}{2}) \equiv E(J_{public}^*) - \frac{s^{CLOB}}{2}$. In the mathematics that follows it is then convenient to define $\lambda_{public}^* = \lambda_{public}(1 - F_{public}^-(\frac{s^{CLOB}}{2}))$ and $\lambda_{private}^* = \lambda_{private}(1 - F_{private}^-(\frac{s^{CLOB}}{2}))$ as the arrival rates of jumps that are larger than the equilibrium spread.

there is an informed trader times the size of the jump $E(J)$, which will be the change in the midpoint. This can be rewritten as $\lambda_{private}E(J) = \lambda_{private}(\frac{s^{CLOB}}{2} + L(\frac{s^{CLOB}}{2}))$.

- *LossAvoidance* is equal to $\lambda_{public} \frac{\rho_{fast}}{N} L(\frac{s^{CLOB}}{2})$: the $\lambda_{public} \frac{\rho_{fast}}{N}$ probability that a fast liquidity provider wins a race with a cancel, times the size of the avoided loss $L(\frac{s^{CLOB}}{2})$.

Now take the equilibrium bid-ask spread as characterized in equation (4.2),

$$\lambda_{invest} \frac{s^{CLOB}}{2} = (\lambda_{public} + \lambda_{private}) \cdot L(\frac{s^{CLOB}}{2}),$$

and add $(\lambda_{public}(1 - \frac{\rho_{fast}}{N}) + \lambda_{private}) \cdot \frac{s^{CLOB}}{2}$ to both sides of the equation. This yields

$$\begin{aligned} & \left(\lambda_{invest} + \lambda_{public}(1 - \frac{\rho_{fast}}{N}) + \lambda_{private} \right) \cdot \frac{s^{CLOB}}{2} \\ &= \left(\lambda_{public}(1 - \frac{\rho_{fast}}{N}) + \lambda_{private} \right) \cdot \left(\frac{s^{CLOB}}{2} + L(\frac{s^{CLOB}}{2}) \right) + \lambda_{public} \frac{\rho_{fast}}{N} L(\frac{s^{CLOB}}{2}). \end{aligned}$$

If we substitute in terms as defined above, this in turn yields

$$EffectiveSpread = PriceImpact_{Race} + PriceImpact_{NonRace} + LossAvoidance.$$

We follow the spread decomposition literature and include *RealizedSpread* as the residual in this equation for the purpose of bringing it to data, yielding equation (5.3) in the text:

$$EffectiveSpread = PriceImpact_{Race} + PriceImpact_{NonRace} + LossAvoidance + RealizedSpread.$$

F.3 Support for the Proportional Reduction in Cost of Liquidity Equations (4.6)-(4.7)

We start with equation (4.4) in the main text, which defines this proportional reduction theoretically:

$$\frac{\frac{s^{CLOB}}{2} - \frac{s^{FBA}}{2}}{\frac{s^{CLOB}}{2}}$$

where s^{CLOB} denotes the equilibrium bid-ask spread in the continuous limit order book market, and s^{FBA} denotes the equilibrium bid-ask spread in the frequent batch auctions market, which eliminates sniping. Next, multiply both the numerator and denominator by $(\lambda_{invest} + \lambda_{private})$:

$$\frac{(\lambda_{invest} + \lambda_{private})\left(\frac{s^{CLOB}}{2} - \frac{s^{FBA}}{2}\right)}{(\lambda_{invest} + \lambda_{private})\frac{s^{CLOB}}{2}}$$

Next, use the bid-ask spread characterization (4.2) in the main text to solve out for $\lambda_{invest} \frac{s^{CLOB}}{2}$

in the numerator:

$$\frac{(\lambda_{public} + \lambda_{private}) \cdot L\left(\frac{s^{CLOB}}{2}\right) + \lambda_{private} \frac{s^{CLOB}}{2} - (\lambda_{invest} + \lambda_{private}) \left(\frac{s^{FBA}}{2}\right)}{(\lambda_{invest} + \lambda_{private}) \frac{s^{CLOB}}{2}}$$

Analogously, use equation (5.1) of Budish, Lee and Shim (2019) to solve out for $\lambda_{invest} \frac{s^{FBA}}{2}$ in the numerator:

$$\frac{(\lambda_{public} + \lambda_{private}) \cdot L\left(\frac{s^{CLOB}}{2}\right) + \lambda_{private} \frac{s^{CLOB}}{2} - \lambda_{private} L\left(\frac{s^{FBA}}{2}\right) - \lambda_{private} \left(\frac{s^{FBA}}{2}\right)}{(\lambda_{invest} + \lambda_{private}) \frac{s^{CLOB}}{2}}$$

Next, regroup terms to place $\lambda_{public} \cdot L\left(\frac{s^{CLOB}}{2}\right)$ on the left of the numerator, and then utilize $L\left(\frac{s}{2}\right) = E(J) - \frac{s}{2}$ for $\lambda_{private} L\left(\frac{s^{CLOB}}{2}\right)$ and $\lambda_{private} L\left(\frac{s^{FBA}}{2}\right)$:

$$\frac{\lambda_{public} \cdot L\left(\frac{s^{CLOB}}{2}\right) + \lambda_{private} \left(E(J) - \frac{s^{CLOB}}{2}\right) + \lambda_{private} \frac{s^{CLOB}}{2} - \lambda_{private} \left(E(J) - \frac{s^{FBA}}{2}\right) - \lambda_{private} \left(\frac{s^{FBA}}{2}\right)}{(\lambda_{invest} + \lambda_{private}) \frac{s^{CLOB}}{2}}$$

Observe that most of the terms in the numerator cancel. Specifically, we have $\lambda_{private} \left(E(J) - \frac{s^{CLOB}}{2}\right) + \lambda_{private} \frac{s^{CLOB}}{2} - \lambda_{private} \left(E(J) - \frac{s^{FBA}}{2}\right) - \lambda_{private} \left(\frac{s^{FBA}}{2}\right) = 0$. This leaves us with:

$$\frac{\lambda_{public} \cdot L\left(\frac{s^{CLOB}}{2}\right)}{(\lambda_{invest} + \lambda_{private}) \frac{s^{CLOB}}{2}}$$

as claimed in the text as equation (4.6). Equation (4.6)'s empirical implementation, equation (4.7), then follows immediately as described in the main text.