

Experiments in High-Frequency Trading: Comparing Two Market Institutions

Online Appendices

Eric M. Aldrich*

Kristian López Vargas†

Department of Economics

Department of Economics

University of California, Santa Cruz

University of California, Santa Cruz

February 28, 2019

A Calibration

[Aldrich and Friedman \(2017\)](#) obtain proprietary data from the IEX exchange for the month of December, 2016. IEX classifies each participant as either an "agency" or "proprietary" trader, the former being the class of traders with a fundamental interest to buy or sell assets (i.e. to maintain an inventory for portfolio reasons). [Aldrich and Friedman \(2017\)](#) report that IEX agency transactions comprised 10,498,518 shares of the S&P 500 exchange traded fund (ticker SPY) during the 21 trading days or $21 \times 6.5 \times 60 = 8190$ trading minutes during December, 2016. Since the median trade size is the minimum block of 100 shares, this amounts to $10,498,518/100 \approx 105,000$ total trades during the month, or $105,000/8190 = 12.82$ investor arrivals per minute, or roughly 1 investor arrival every 4.68 seconds. Although IEX represents only a small fraction of equities market share, we believe that most fundamental traders (investors) will utilize IEX in conjunction with other equities exchanges, and hence that their arrival rates would be suggestive of aggregate investor arrival intensities. The result is that in raw time (we discuss time rescaling below), the Poisson intensity parameter for investor arrivals is $\tilde{\lambda}_I = 1/4.68$.

*Email: ealdrich@ucsc.edu.

†Email: kristian@ucsc.edu.

To calibrate $\tilde{\lambda}_V$, the intensity of the Poisson process governing jumps in the fundamental value, we utilize SPY quotation data at Nasdaq, which, given its liquidity and overall market share, is a good surrogate for the SPY national best bid and offer (NBBO). Our sample covers the period 16 June – 11 September, 2014. There are 26,216,524 quotations in the 62-day period, which comprises 1,450,800,000 milliseconds during trading hours, or approximately 1 quote every 55 milliseconds. Defining a jump as any midpoint price change of magnitude at least \$0.01 over the period of four quotations, or 220 milliseconds, resulted in a median of approximately 3978 jumps per day, or one jump every 5.88 seconds (assuming 23,400 seconds during the 6.5 hour equities market trading day). Hence, $\tilde{\lambda}_J = 1/5.88$, prior to time rescaling.

Following Aldrich et al. (2016), we assume the trade-time distribution of asset price changes, $\Delta V(t)$, is Gaussian with mean zero. Using the SPY data above, we find that $Std(\Delta V(t)) = \$0.007$, or slightly less than the minimum spread of \$0.01. However, given the preponderance of liquidity at SPY best bid and offer and the fact that the minimum spread is set by the SEC, the unconstrained equilibrium spread is widely considered to be less than \$0.01. This suggests that the standard deviation of value changes should be of similar magnitude to the equilibrium minimum spread. As the magnitude of the scale parameter is otherwise arbitrary, we set $\Delta V(t) \sim \mathcal{N}(0, \sigma = 0.5)$, resulting in a scale parameter that is of the same approximate magnitude as the CDA equilibrium spreads. Further, the choice region for maker spreads in the experiment was set to encompass a region of 4σ around the fundamental value.

We set $\lambda_V = 1/4$ in our baseline calibration, which leads us to interpret 1 second of lab time as $\lambda_V/\tilde{\lambda}_V = 5.88/4 = 1.47$ seconds of raw financial market time. Consequently, we interpret a single four-minute experimental period as approximately $4 \times 1.47 = 5.88$ minutes of financial market time, and the full eight-period session as approximately $8 \times 4 \times 1.47 \approx 47$ minutes of market time. Further, rescaling the investor intensity parameter to experimental time results in $\lambda_I = 1.47/4.68 \approx 1/3$, which is the value of our baseline calibration.

To calibrate the cost of fast communication technology, c_s , we use the pricing schedule of McKay Brothers LLC, a premier microwave transmission service. To transmit a single symbol on the long-haul route between the CME data center in Aurora, IL to an equities data center in New Jersey costs \$10,600 per month, or about \$0.02 per second (assuming 22 trading days per month, and 6.5 trading hours per day). Scaling to experimental time, this results in about \$0.015 per second. In our baseline calibration we set $c_s = \$0.01$ per second.

B Off-Equilibrium Sniping

In the BCS equilibrium under CDA, sniping transactions occur between one of the $N - 1$ snipers and the single maker. In practice, if there is more than one maker (off-equilibrium) at the time of a jump in the fundamental value, the maker whose orders are repriced first may snipe the maker who is repriced after. These maker-to-maker sniping transactions are not explicitly dealt with in [Budish et al. \(2015\)](#) and raise the question as to whether they could alter the equilibrium. Empirically, we find the prevalence of these transactions is substantial. However, augmenting the model with this strategy does not extend the set of equilibria, which remains unique up to the aggregate composition of making and sniping strategies.

We illustrate this type event with an example. Consider a market operating under the CDA. Suppose at $t = 0$ the value of the asset is $V = 100$ and only two traders are present in the market as makers. Let us refer to these traders as M1 and M2, and assume that they both have a spread of $s = 2$ (bids at 99 and offers at 101), and that M1 has purchased speed services (operating at a latency of 100 ms) and M2 has not (operating at a latency of 500 ms). If the value jumps to $V = 105$ at $t = 700$ ms, M1's algorithm will submit messages to the exchange, updating her bid and offer to 104 and 106, respectively. Those messages will be received at the messaging server at $t = 800$ ms and immediately passed to the order book, where M1's new bid of 104 will cross with M2's stale offer of 101. Under typical exchange rules, the order is filled at 101, resulting in a profit (loss) of $105 - 101 = 4$ to M1 (M2).¹

In the equilibrium under FBA, on the other hand, sniping transactions occur between investors and any of the N makers when the value jump occurs too late for makers to update quotes before batch end. In practice, if some makers purchase speed technology (off-equilibrium), maker-to-maker sniping can occur. However, since these events can only happen when jumps occur close to the end of the batch, their prevalence is much less common than in the CDA. Additionally, if one or more traders decide to act as snipers (off-equilibrium), both slow and fast makers will be sniped by investors, fast makers and fast snipers, in that order. As shown below, this is the pattern we find in the data.

Table 1 reports frequency counts of snipes under each of the CDA and FBA configurations. In each case, the number of snipes is decomposed by the roles of the traders participating in the sniping transaction. While the parties of each transaction are

¹Similar events occur less frequently when makers (fast or slow) are filled by an investor and the asset value changes multiple times during the interval in which new orders are being routed to the exchange. We do not focus on these, as they account for very little total volume.

clearly defined in the CDA, the same is not true of the FBA: there is no specific attribution of which traders transact with each other in a call-type auction. To make such attributions, we paired sniped traders (those with negative profits) with non-sniped traders in the following priority: (1) fast snipers, (2) fast makers, and (3) investors. Since slow makers and slow snipers never have an opportunity to snipe in the FBA, we make no such attributions. Although our particular attribution order may be somewhat arbitrary, it corresponds to the correct attribution in batches with a single transaction and the aggregate counts are comparable to those of the CDA.

| | Investor | | Fast Maker | | Slow Maker | | Fast Sniper | | Slow Sniper | | Total | |
|---------------------|----------|-----|------------|-----|------------|-----|-------------|-----|-------------|-----|-------|------|
| | CDA | FBA | CDA | FBA | CDA | FBA | CDA | FBA | CDA | FBA | CDA | FBA |
| (a) Configuration 1 | | | | | | | | | | | | |
| Fast Maker | 32 | 14 | 641 | 10 | 7 | NA | 432 | 0 | 8 | NA | 1,120 | 24 |
| Slow Maker | 22 | 193 | 870 | 112 | 28 | NA | 413 | 55 | 51 | NA | 1,384 | 360 |
| (b) Configuration 2 | | | | | | | | | | | | |
| Fast Maker | 38 | 49 | 397 | 61 | 25 | NA | 1,885 | 0 | 21 | NA | 2,366 | 110 |
| Slow Maker | 29 | 271 | 536 | 639 | 39 | NA | 1,473 | 181 | 83 | NA | 2,160 | 1092 |
| (c) Configuration 3 | | | | | | | | | | | | |
| Fast Maker | 135 | 77 | 388 | 31 | 18 | NA | 2,102 | 2 | 23 | NA | 2,666 | 110 |
| Slow Maker | 69 | 448 | 416 | 479 | 22 | NA | 1,048 | 158 | 153 | NA | 1,708 | 1085 |

Table 1: Frequency counts of snipes under each of the CDA and FBA configurations, decomposed by the role of the sniper and the trader being sniped.

The values in Table 1 demonstrate that sniping is much more prevalent under the CDA, with the bulk of CDA snipes done by fast snipers or fast makers. Interestingly, consistent with the off-equilibrium insight at the beginning of this section, fast makers account for a larger share of sniping in CDA configuration 1, relative to CDA configurations 2 and 3. This latter result is due to the higher population of snipers under the latter two configurations (see Figure 3 and Table 2). Under FBA, the number of times that fast makers are sniped is about 20 times less than under CDA, and slow makers are sniped roughly 2 to 4 times less. Further, we make the bulk of FBA attributions to fast makers and investors rather than fast snipers.

As successful sniping of any form occurs infrequently under FBA, the consequences of unintentional maker snipes are more important under CDA. That is, makers have access to similar profit opportunities as snipers, with the additional benefit (cost) of investor transactions (being sniped). We detail the trade-offs to this expanded strategy space below, and prove that the model equilibrium is unchanged. Despite the identical equilibrium, it is possible that the perceived benefit of maker-to-maker snipes may

partly explain why our experimental observations are attenuated relative to predicted BCS equilibrium values.

B.1 Equilibrium with Maker-Snipers in the CDA

We now show that extending the BCS model to explicitly account for maker-to-maker sniping leaves the CDA equilibrium intact.

Proposition B.1. *Consider an augmented BCS model that allows repriced maker orders to transact with stale limit orders posted by other traders. The equilibrium of the augmented model is identical to that of [Budish et al. \(2015\)](#).*

Proof We sketch the proof for the case of the CDA with endogenous entry. Suppose an equilibrium exists with N trading firms acting as market makers, quoting the same spread, s , and all purchasing fast communication technology. The profit to each firm would be

$$\lambda_I \cdot \frac{s}{2N} - \lambda_V \cdot \Pr(J > s) \cdot \mathbb{E} \left[J - \frac{s}{2} | J > s \right] \cdot \Pr(\text{Sniped}) \\ + \lambda_V \cdot \Pr(J > s) \cdot \mathbb{E} \left[J - \frac{s}{2} | J > s \right] \cdot \Pr(\text{Sniping}) = c_{speed} \quad (\text{A.1})$$

$$\Rightarrow \lambda_I \cdot \frac{s}{2N} = c_{speed}. \quad (\text{A.2})$$

The LHS of Equation (A.1) has positive and negative terms related to sniping which cancel, resulting in Equation (A.2). We examine the terms on the LHS of Equation (A.1) and compare them with the typical profit condition for a maker in the BCS equilibrium (Equation (1) of the paper):

1. Investor profits under the all-maker equilibrium are shared among all N firms: $\lambda_I \cdot \frac{s}{2N}$. In the BCS equilibrium, a single trading firm acts as market maker and captures all of the profit alone: $\lambda_I \cdot \frac{s}{2}$.
2. Makers can snipe other makers, but it requires a larger jump to do so: the jump must be at least s in magnitude in order for the bid (offer) of a sniping maker to cross with the stale offer (bid) of another maker. In the BCS equilibrium, a jump of magnitude $s/2$ suffices for a sniper to transact with a stale maker quote. In both cases, the sniping profit is $J - \frac{s}{2}$. Thus, the second term on the LHS of Equation (A.1) breaks down in the following way: (i) λ_V is the jump intensity, (ii) $\Pr(J > s)$ is the probability that a jump is big enough for the maker to be sniped, (iii) $\mathbb{E} \left[J - \frac{s}{2} | J > s \right]$ is the expected sniping profit (conditional on sufficiently large

jump), and (iv) $\Pr(\text{Sniped})$ is the probability that a maker is sniped. The same terms in Equation (1) of the paper are: (i) λ_V , (ii) $\Pr\left(J > \frac{s}{2}\right)$, which is larger than $\Pr(J > s)$, (iii) $\mathbb{E}\left[J - \frac{s}{2} | J > \frac{s}{2}\right]$, which is less than $\mathbb{E}\left[J - \frac{s}{2} | J > s\right]$, and (iv) $\frac{N-1}{N}$, which is an upper bound for $\Pr(\text{Sniped})$ (see below).

Letting $\Pr(\text{Sniping})$ denote the probability of a maker successfully sniping another maker, it can be readily shown that the $\Pr(\text{Sniping}) = \Pr(\text{Sniped})$ (e.g. the High School Prom Theorem – see Davis, 2009). Furthermore, these probabilities lie within the open interval $\left(\frac{\lfloor N/2 \rfloor}{N}, \frac{N-1}{N}\right)$. Thus, the losses attributed to sniping in Equation (A.1) are exactly offset by the gains, resulting in a cancellation of the second and third terms on the LHS of Equation (A.1).

We now explore the profitability of two possible deviations: (1) a single maker quoting a more narrow spread and (2) a single maker switching roles to act as a pure sniper. Maintaining the notation s^* for the BCS equilibrium spread, we let \hat{s} denote the all-maker equilibrium spread. For the first deviation (narrower spread) to be profitable, the following would need to hold

$$\begin{aligned} \lambda_I \cdot \frac{\hat{s}}{2N} &< \lambda_I \cdot \frac{\hat{s} - \varepsilon}{2} \\ &- \lambda_V \cdot \Pr\left(J > \frac{\hat{s} + (\hat{s} - \varepsilon)}{2}\right) \\ &\quad \times \mathbb{E}\left[J - \frac{\hat{s} - \varepsilon}{2} | J > \frac{\hat{s} + (\hat{s} - \varepsilon)}{2}\right] \cdot \frac{N-1}{N} \quad (\text{A.3}) \\ &+ \lambda_V \cdot \Pr\left(J > \frac{\hat{s} + (\hat{s} - \varepsilon)}{2}\right) \\ &\quad \times \mathbb{E}\left[J - \frac{\hat{s} - \varepsilon}{2} | J > \frac{\hat{s} + (\hat{s} - \varepsilon)}{2}\right] \cdot \Pr(\text{Snipe}) \end{aligned} \quad (\text{A.4})$$

$$\begin{aligned} \Rightarrow N(\hat{s} - \varepsilon) - \hat{s} &> 2 \cdot \frac{\lambda_V}{\lambda_I} \cdot \Pr\left(J > \frac{\hat{s} + (\hat{s} - \varepsilon)}{2}\right) \\ &\quad \times \mathbb{E}\left[J - \frac{\hat{s} - \varepsilon}{2} | J > \frac{\hat{s} + (\hat{s} - \varepsilon)}{2}\right] \cdot (N-1). \end{aligned} \quad (\text{A.5})$$

Intuitively, the deviating maker would capture all investor profits at the expense of increased probability, $\frac{N-1}{N} > \Pr(\text{Sniped})$, of getting sniped. In Equation (A.5), we have eliminated a positive constant which does not affect the inequality. Taking the limit as $\varepsilon \rightarrow 0$, Equation (A.5) can be expressed as

$$\hat{s} > 2 \cdot \frac{\lambda_V}{\lambda_I} \cdot \Pr(J > \hat{s}) \cdot \mathbb{E}\left[J - \frac{\hat{s}}{2} | J > \hat{s}\right]. \quad (\text{A.6})$$

Alternatively, if a single maker deviates to act as a sniper, the resulting profit would be

$$\lambda_V \cdot \Pr\left(J > \frac{\hat{s}}{2}\right) \cdot \mathbb{E}\left[J - \frac{\hat{s}}{2} \mid J > \frac{\hat{s}}{2}\right] \cdot \frac{N-1}{N} = c_{speed} \quad (\text{A.7})$$

Thus, a sniper deviation is profitable if

$$\lambda_V \cdot \Pr\left(J > \frac{\hat{s}}{2}\right) \cdot \mathbb{E}\left[J - \frac{\hat{s}}{2} \mid J > \frac{\hat{s}}{2}\right] \cdot \frac{N-1}{N} > \lambda_I \cdot \frac{\hat{s}}{2N} \quad (\text{A.8})$$

$$\Rightarrow \hat{s} < 2 \cdot \frac{\lambda_V}{\lambda_I} \cdot \Pr\left(J > \frac{\hat{s}}{2}\right) \cdot \mathbb{E}\left[J - \frac{\hat{s}}{2} \mid J > \frac{\hat{s}}{2}\right] \cdot (N-1). \quad (\text{A.9})$$

To determine if a profitable deviation exists, we compare Equations (A.6) and (A.9). If at least one of the conditions is always satisfied, the all-maker equilibrium cannot be supported. This is the case if the RHS of Equation (A.9) is always greater than the RHS of Equation (A.6):

$$\begin{aligned} \Pr(J > \hat{s}) \cdot \mathbb{E}\left[J - \frac{\hat{s}}{2} \mid J > \hat{s}\right] &< \Pr\left(J > \frac{\hat{s}}{2}\right) \cdot \mathbb{E}\left[J - \frac{\hat{s}}{2} \mid J > \frac{\hat{s}}{2}\right] \cdot (N-1) \\ \Rightarrow \Pr(J > \hat{s}) \frac{\int_{\hat{s}}(z - \frac{\hat{s}}{2})2\phi(z)dz}{\Pr(J > \hat{s})} &< \Pr\left(J > \frac{\hat{s}}{2}\right) \frac{\int_{\frac{\hat{s}}{2}}(z - \frac{\hat{s}}{2})2\phi(z)dz}{\Pr(J > \frac{\hat{s}}{2})} \cdot (N-1) \\ \Rightarrow \int_{\hat{s}}(z - \frac{\hat{s}}{2})2\phi(z)dz &< \int_{\frac{\hat{s}}{2}}(z - \frac{\hat{s}}{2})2\phi(z)dz \cdot (N-1) \\ \Rightarrow \int_{\hat{s}}(z - \frac{\hat{s}}{2})2\phi(z)dz &< \int_{\frac{\hat{s}}{2}}(z - \frac{\hat{s}}{2})(N-1)2\phi(z)dz \quad (\text{A.10}) \\ \Rightarrow \int_{\hat{s}}(z - \frac{\hat{s}}{2})2\phi(z)dz &< \int_{\frac{\hat{s}}{2}}^{\hat{s}}(z - \frac{\hat{s}}{2})(N-1)2\phi(z)dz \\ &\quad + \int_{\hat{s}}(z - \frac{\hat{s}}{2})(N-1)2\phi(z)dz \\ \Rightarrow 0 &< \int_{\frac{\hat{s}}{2}}^{\hat{s}}(z - \frac{\hat{s}}{2})(N-1)2\phi(z)dz \\ &\quad + \int_{\hat{s}}(z - \frac{\hat{s}}{2})(N-2)2\phi(z)dz. \end{aligned}$$

Since the RHS of the last equation of System (A.10) is always positive for $N \geq 3$, we conclude that the all-maker equilibrium cannot be supported. Further deviations can be derived inductively, resulting in the equilibrium as stated in Budish et al. (2015).

C Collusive Play

In this section, we address the possibility of collusive behavior in our experimental markets.² In short, our general finding is that human subjects in both market formats exhibited behavior more congruent with competitive, rather than collusive, play. This is consistent with the theoretical insight and empirical findings of existing literature, suggesting that cooperation becomes more difficult when a game has three or more players.

In the FBA, which is a repeated game in discrete time, collusive play would imply all traders choosing to be a maker, at positive (arguably maximal) spread, and not buying speed services. Although there is no previous result on collusion that applies directly to this game, (Budish et al., 2015 focused only on the pure-strategy, instantaneous Nash Equilibrium), theory and empirical work on Bertrand games and social dilemmas (e.g., prisoner's dilemma) relate to our setting and therefore provide relevant insights.

In Bertrand games, as well as prisoner's dilemma (PD) games, where the Nash Equilibrium is inefficient from the perspective of players, cooperation is not possible if the stage game is repeated a finite number of times. Sub-game perfection implies non-cooperative behavior in the last period and inductively unravels cooperation in all earlier periods.

However, for the PD, as documented by Selten and Stoecker (1986), Friedman and Oprea (2012), Embrey et al. (2018) and others, unraveling is limited. Indeed when the game has many periods, cooperation is often observed in two-person PD (Friedman and Oprea, 2012; Embrey et al., 2018). Departures from non cooperative behavior have been attributed to different underlying features of decision making. As noted in Embrey et al. (2018), cooperation in a PD stage game is predicted by models where: (1) there is uncertainty about the type and payoffs of other players (Kreps et al., 1982), (2) players generate *epsilon equilibria* (Radner, 1986; Friedman and Oprea, 2012) by optimizing up to an epsilon distance from the maximal payoff, (3) learning is present (Mengel, 2014), and (4) traders exhibit limited forward reasoning (Mantovani, 2016). Some forms of social preferences could rationalize cooperative behavior as well.

PD games with $N > 2$ tend to display low levels of cooperation (e.g. Barcelo and Capraro, 2015). For N -player, finitely-repeated Bertrand games, empirical and theoretical insights are more sparse. In one shot encounters there is evidence of collusive behavior with two players, but collusion is infrequently observed when there are three or more players (Dufwenberg and Gneezy, 2000; Abbink and Brandts, 2005, 2008; Orzen,

²We thank the editor as well as an anonymous referee for suggesting and encouraging us to include this discussion.

2008; Potters and Suetens, 2013). For the case of Cournot games, Friedman et al. (2014) shows that with imitation strategies and long time horizons (e.g., over 1000), collusive play is commonly achieved in duopolies and, less commonly, in triopolies. However, Oechssler et al. (2016) show that with four firms, none of the oligopolies becomes permanently collusive. For the case of Bertrand oligopoly, although there is evidence of cooperation in repeated two-player games (see Argenton and Müller, 2012), to the best of our knowledge there are no relevant studies with multiple players, fixed matching, and repetition over many periods. Therefore, the relative contribution of time horizon (in repeated games) and number of participants on cooperative play is an open question for this class of games. Our experiments indicate that even with hundreds of periods, the number of players would be the prevailing force for shaping (the lack of) cooperation.

The CDA, on the other hand, is a finite-horizon, continuous-time form of an undercutting game. The body of research on undercutting and cooperation games is small, but has relevant insights. Following the work of Radner (1986) and Simon and Stinchcombe (1989), Friedman and Oprea (2012) formalize a model for a continuous-time prisoners dilemma (in addition to a fast-paced, discrete-time PD). They use the notion of epsilon-equilibrium to derive an equilibrium with near-perfect cooperation, where unraveling forces are limited if (1) players are willing to give up a tiny part of their profits and (2) they can react almost immediately. The logic is that the ability to react quickly reduces the risk of loss related to defection and makes a cooperative equilibrium easier to sustain. Their empirical findings are congruent with the existence of a high-cooperation equilibrium in a continuous-time PD. Park (2014) formalizes a similar model where cooperation is sustained on heterogeneous reaction times that are private information.

Broad theoretical findings for cooperation in continuous-time PD games are, however, ambiguous. Although the literature above predicts full cooperation, a simple extension of a one-shot game predicts that cooperation would not exist, whereas an extended Folk Theorem gives broad predictions with differing degrees of cooperation. Empirically, the experiments of Horstmann et al. (2015) find reduced cooperation in continuous time, in direct contrast to the results of Friedman and Oprea (2012). Their work, however, finds that an increase in the number of players substantially reduces cooperation, and is therefore consistent with the findings of the PD, Cournot and Bertrand literature cited above.

Together, these insights are consistent with our experimental findings: very little cooperation, most likely due to the number of players in each market (six).

C.1 Predicted Collusive Profits in FBA and CDA

Collusive behavior in the FBA would result in all traders choosing to be slow makers, quoting the maximal spread, $\bar{S} = 2$. Under these circumstances, expected profits for each player are equal to:

$$\pi_c^{FBA} = \frac{\bar{S}}{N} \cdot T \cdot E[N_I^e], \quad (\text{A.11})$$

where N is the number of traders in the market, T is the number of batches (80 per 4-minute trading period), and $E[N_I^e]$ is the expected absolute difference in the number of buy and sell investors that arrive within a single batch (recall that investors cross with each other before human participants in the FBA). More precisely, N_I^e is the absolute value of a Skellam random variable and has expectation

$$E[N_I^e] = \sum_{k=0}^{\infty} k P(|N_{sell} - N_{buy}| = k) \quad (\text{A.12})$$

$$= \sum_{k=1}^{\infty} 2k \sum_{l=k}^{\infty} P(N_{sell} = l) P(N_{buy} = l - k) \quad (\text{A.13})$$

$$= \sum_{k=1}^{\infty} 2k \sum_{l=k}^{\infty} \frac{(\lambda_I \tau)^{2l-k}}{l!(l-k)!} e^{-2\lambda_I \tau}, \quad (\text{A.14})$$

where N_{sell} and N_{buy} are the number of selling and buying investors, respectively, that arrive in a batch and τ is the batch length (3 seconds in our setting). In the configurations we study, $\pi_c^{FBA} = (27.93, 20.32, 35.19)$ ECUs for configurations 1–3, respectively. These values are substantially larger than the observed per period/player profits in our experiments, reported in Table 2 of the paper: $(0.435, 0.372, 1.52)$ ECUs for configurations 1–3, respectively.

Similar to the FBA, collusive behavior under the CDA would result in all traders choosing to be slow makers, quoting the maximal spread, $\bar{S} = 2$. In this case, expected profits for each player are equal to:

$$\pi_c^{CDA} = \frac{\bar{S}}{N} \cdot T \cdot \lambda_I, \quad (\text{A.15})$$

where T is the number of seconds (240 per trading period) and N is the number of traders in a market. In the configurations we consider, $\pi_c^{CDA} = (26.67, 16, 40)$ ECUs for configurations 1–3, respectively, and are likewise substantially larger than the observed values reported in Table 2 of the paper: $(0.0869, 0.603, 4.31)$ ECUs for configurations

1–3, respectively.

D Equilibrium Market Statistics

Proposition D.1. *In the equilibrium of the BCS model under CDA, the standard deviation of changes in transactions prices is*

$$\text{Std}(P_t - P_{t-1}) = \sqrt{\text{Var}(P_t - P_{t-1})} = s\sqrt{\left(\frac{A}{2} - \frac{A^2}{4}\right)}, \quad (\text{A.16})$$

where $A = \lambda_I + \lambda_V \text{Pr}(J > \frac{s}{2})$.

Proof Price changes are always zero or s . Suppose that the last transaction, P_{t-1} , occurred at the best bid. If the next transaction is attributed to an investor arrival, it occurs at the same price ($P_t - P_{t-1} = 0$) with probability $\lambda_I/2$ and on the best offer ($P_t - P_{t-1} = s$), with probability $\lambda_I/2$. If the next transaction is attributed to a positive value jump, a sniping transaction will occur on the (stale) best offer ($P_t - P_{t-1} = s$) with probability $\frac{1}{2}\lambda_V \text{Pr}(J > \frac{s}{2})$, whereas if it is attributed to a negative value jump, it occurs on the (stale) best bid ($P_t - P_{t-1} = 0$) with the same probability. The case of P_{t-1} occurring on the best offer is symmetric. Collecting these results, we find

$$\begin{aligned} E[(P_t - P_{t-1})^2] &= \lambda_I \left(\frac{1}{2}0^2 + \frac{1}{2}s^2\right) + \lambda_V \text{Pr}\left(J > \frac{s}{2}\right) \left(\frac{1}{2}0^2 + \frac{1}{2}s^2\right) \\ &= \frac{s^2}{2} \left(\lambda_I + \lambda_V \text{Pr}\left(J > \frac{s}{2}\right)\right) \end{aligned} \quad (\text{A.17})$$

and

$$E[P_t - P_{t-1}]^2 = \frac{s^2}{4} \left(\lambda_I + \lambda_V \text{Pr}\left(J > \frac{s}{2}\right)\right)^2. \quad (\text{A.18})$$

Thus,

$$\begin{aligned} \text{Std}(P_t - P_{t-1}) &= \sqrt{\text{Var}(P_t - P_{t-1})} \\ &= \sqrt{E[(P_t - P_{t-1})^2] - E[P_t - P_{t-1}]^2} \\ &= s\sqrt{\left(\frac{A}{2} - \frac{A^2}{4}\right)}. \end{aligned} \quad (\text{A.19})$$

Proposition D.2. *In the equilibrium of the BCS model under FBA, the standard de-*

viation of changes in transactions prices is

$$\text{Std}(P_t - P_{t-1}) = \sigma \sqrt{(1 - e^{-\tau\lambda_V})(1 - e^{-\tau\lambda_I})}. \quad (\text{A.20})$$

Proof. In the FBA equilibrium, all traders act as makers and quote zero spread. As a result, there are no sniping transactions and, in the absence of jumps, $P_t = P_{t-1} = V_t$ at the time of auctions (when the auction results in a transaction). The only instances that result in $P_t \neq V_t$ are batches with at least one change in the fundamental value and at least one investor arrival. This implies that price changes are perfectly correlated with value changes and that the standard deviation of $P_t - P_{t-1}$ should be identical to the standard deviation of $V_t - V_{t-1}$, but scaled by the probability that a value jump occurs in the batch and at least one investor arrives:

$$\text{Std}(P_t - P_{t-1}) = \sigma \sqrt{(1 - e^{-\tau\lambda_V})(1 - e^{-\tau\lambda_I})}. \quad (\text{A.21})$$

Proposition D.3. *In the equilibrium of the BCS model under CDA, the root mean squared deviation of transactions prices from the fundamental value is*

$$\text{RMSD}(P_t - V_t) = \sqrt{\lambda_I \left(\frac{s}{2}\right)^2 + \lambda_V \int_{\frac{s}{2}} \left(J - \frac{s}{2}\right)^2 dF}. \quad (\text{A.22})$$

Proof. In equilibrium, investor transactions, occurring with probability λ_I , always take place at the bid or offer, such that $P_t - V_t = s/2$, where s is the equilibrium spread of the single market maker. At the time of jumps that are sufficiently large to result in sniping transactions, occurring with probability $\lambda_V Pr\left(J > \frac{s}{2}\right)$ (where $J = |\Delta V_t|$), the difference between the new fundamental V'_t and the transaction price P_t is $J - \frac{s}{2}$. Collecting these results,

$$\begin{aligned} \text{RMSD}(P_t - V_t) &= \sqrt{E[(P_t - V_t)^2]} \\ &= \sqrt{\lambda_I \left(\frac{s}{2}\right)^2 + \lambda_V Pr\left(J > \frac{s}{2}\right) E\left[\left(J - \frac{s}{2}\right)^2 \mid J > \frac{s}{2}\right]} \\ &= \sqrt{\lambda_I \left(\frac{s}{2}\right)^2 + \lambda_V \int_{\frac{s}{2}} \left(J - \frac{s}{2}\right)^2 dF}. \end{aligned} \quad (\text{A.23})$$

Proposition D.4. *In the equilibrium of the BCS model under FBA, the root mean*

squared deviation of transactions prices from the fundamental value is

$$RMSD(P_t - V_t) = \sigma \sqrt{(1 - e^{-\delta_{slow}\lambda_V})(1 - e^{-\tau\lambda_I})}. \quad (\text{A.24})$$

Proof. Since all traders are makers and quote zero spreads in the FBA equilibrium, $P_t \neq V_t$ only in instances that the fundamental value changes in the period $(t - \delta_{slow}, t]$ prior to the auction at time t , and when at least one investor arrives (otherwise there are no transactions). The probability of the former event is $1 - e^{-\delta_{slow}\lambda_V}$ and the probability of the latter is $1 - e^{-\tau\lambda_I}$. Thus, price changes are perfectly correlated with value changes and the standard deviation of $P_t - P_{t-1}$ should be identical to the standard deviation of $V_t - V_{t-1}$, but scaled by the aforementioned probabilities, which is the stated result.

E Instructions

Below we provide a static version of the experiment instructions for both market formats. Instructions with the animated figures (GIFS) are also available online. To access the animated instructions, go to: <https://leeps.ucsc.edu/media/alv2019/>

Experiment Instructions (FBA)

Welcome, and thank you for participating!

From now until the end of the experiment, please turn off your phone and do not communicate with other participants. If you have any questions, please raise your hand; an experimenter will come and answer your question. Please pay careful attention to the instructions as real money is at stake.

During the experiment you will earn *Experimental Currency Units* (ECUs) and at the end of the experiment all your earnings will be converted to US dollars at the rate of one dollar for every two ECUs. You are guaranteed a show up fee of \$7.00 but can earn considerably more.

Basic Ideas

In this experiment you will participate in a simple **automated financial market**. Using information displayed on your screen as in Figure 1, you will **set and adjust trading algorithms (bots)** so as to earn as many ECUs as you can. As explained below, your earnings will depend on the settings you choose and on the choices of **the other five participants** in your group. There will be **eight trading periods**, each lasting **four minutes**, after which the experiment will end and you will be paid.

(removed)
Figure 1

Stock value. There is just one stock traded in your financial market. The stock has a *value* (henceforth “*V*”) that will jump up or down by a random amount at random times. The average time between jumps is about 1 second, but shorter or much longer times are possible. The **Choice Box** on your screen will always center on the current value of *V*, depicted by a horizontal gray line (see Figure 1). Every time there is a jump in *V*, there will be a yellow flash around the center line, and all other prices (described below) will appear to shift until they adjust to the new value of *V*.

Buy Orders, Sell Orders, and Spreads. All participants control algorithms (bots) that submit *buy orders* and *sell orders*. Submitting a *buy order for 99*, for example, means that you are willing to pay 99 ECUs or less to buy a share of the stock. Similarly, if you submit a *sell order for 101*, you are willing to receive 101 ECUs or more to sell a share of the stock. The price difference between your buy and sell orders is referred to as your *spread*. In the 99/101 example, your spread is 2.

Investor arrivals. At random times an automated investor arrives in the market to buy or to sell a share of the stock. Investors transact immediately at the best price they find in the market upon arrival.

In today's experiment, the average time between investor arrivals is about 2 seconds, but again shorter or much longer times are possible.

Auctions. Trading occurs in auctions that are conducted exactly once every three seconds. Although participants may submit and cancel orders at any time, the only orders that matter are those that are active at the moment an auction is conducted. At that moment, a computer algorithm finds a **clearing price** P and fills:

- all buy orders with prices above P ,
- all sell orders with prices below P ,
- as many as possible buy and sell orders placed exactly at P .

Filled buy orders pay the clearing price P , and filled sell orders get paid the clearing price P . Thus, every filled order gets at least as good a price as the order specifies. The clearing price P is chosen so that the number of filled sell orders is equal to the number of filled buy orders. For example, if the clearing price is $P=101$ and you have a sell order at 98, then your order will be filled at 101 (i.e. you get paid 101 for your share).

The **clearing price** of the last auction will be depicted by a short red tick in the choice box (See Figure 1). Unfilled orders are automatically carried over into the next auction, but of course you can cancel or reprice your orders if you wish.

Your Choices and Earnings

Participants choose among three possible roles: *maker*, *sniper* and *out*. You can change your role at any time of the trading period, using the buttons on the right side of the *choice box* (see GIF 1a).

Maker. As a maker, you post both a buy order below V and a sell order above V using the computer program (henceforth, "your bot"). These orders are symmetrically located around V and are represented by the horizontal blue marks above and below the center line representing V . Your spread is the difference between your sell and buy marks (depicted by the light green or light blue shaded region between those marks). To adjust your orders (and spread), simply click at the desired point in the choice box (see GIF 1b). Your bot will cancel your old buy and sell orders and send new orders to the market with the spread that you just chose. Since it takes time for the new orders to travel to the market, your new orders are initially represented as horizontal lines at the right side of the choice box which slide toward the center of the box, at which time they replace the old orders (see GIF 1b).

You can make a profit at the moment of an auction when more "buy" investors arrived for that action than "sell" investors or vice versa. (When the same number of buy and sell investor arrive, they transact amongst themselves.) If your buy order is above the clearing price P , you buy at P and your bot automatically sells the unit at V . You earn a profit of $V - P$ (see GIF 1c). However, if your buy order is below the clearing price, it just gets carried over into the next auction. Similarly, if you have a sell order below the clearing price P , you will sell at that price and earn $P - V$, because your bot buys a replacement unit at the value V . These profits are represented by flashing green lines in the choice box, and are also displayed as vertical green lines in the profit box (see GIF 1c).

You can also lose money (earn a negative profit) while in the maker role. When V jumps, your bot sends updated orders to the market with prices that are centered on the new value of V . Your new orders take half a second to get to the market (or one tenth of a second if you have your speed switch on) and there are three seconds between each auction. In only one out of ten jumps, if you have speed off, (and, in one out of 100 jumps, if you have speed on) your new orders will not arrive to the market on time for the auction. If your orders do not make it to the market, the auction will use your previous orders and there is a chance you will earn a negative profit.

For example, suppose that $V=100$ and your current spread as a maker is 2 (that is you have a buy order at 99 and sell order at 101). If V jumped to 105 just one tenth of a second before the auction time, your bot will send messages to the market updating your buy and sell orders to prices 104 and 106, respectively. But because this jump by chance happened very close to the auction time, you are still committed to sell for 101 (cheap!). When this happens, we say that your orders are “stale”. If another participant’s bot places a buy order and her order does make it to the market before the auction (this could happen if she has speed on), she will possibly buy cheap from you and make a profit. And because you sold too cheap compared to the value V you will lose money. When this happens we say you have been “sniped”; and your loss is represented by the vertical distance between V and the clearing price. Losses from snipes are represented by flashing red lines between V and the corresponding clearing price in the choice box.

| | | |
|-----------------------|-----------------------------|----------------------------|
| (removed) | (removed) | (removed) |
| GIF 1a Role Choice | GIF 1b Change Own Spread | GIF 1c Investor Arrival |

Sniper. If you choose to be a sniper, your bot attempts to transact with stale orders. Whenever there is an jump in V , your bot will send an order that tries to buy cheap from a stale sell order (for an up jump in V) or tries to sell at a high price to a stale buy order (for a down jump). As with maker orders, it takes time for sniper orders to reach the market. You will be able to snipe stale orders only in the cases where jumps happen a fraction of a second before the auction, makers’ bots are not able to get their updates to the market on time, and your own order does arrive before the auction. When you transact, your choice box will flash green to indicate your profit (see GIF 2c).

Out. If you choose to be out, you do not participate in the market. In this role you never earn profits nor take losses.

Choosing Speed

As a *maker* or a *sniper*, it takes time to send your orders to the market. At **normal speed** (the default option) it takes half of a second to update orders. Alternatively, you can choose **fast speed** to reduce the update time to one tenth of a second, but it costs 0.022 ECUs per second for the faster service.

You can activate or deactivate the speed option at any time. For example, if you activate speed now and deactivate it 30 seconds later, you are charged 0.66 ECUs, as you will see in the Profit Box.

| | | |
|--------------------|------------------------------------|-----------------------------|
| (removed) | (removed) | (removed) |
| 2a Value Change | 2b Value Change, Getting Sniped | 2c Value Change, Sniping |

Trading Period Payoffs. You begin each *trading period* with an endowment of 20 ECUs. Gains and losses from the trading period will be added to and subtracted from your endowment to form your profits for the period.

Interim Screens. In between trading periods, your screen will display a summary of role and strategy choices as well as profits made in the last period. On the left side of the screen, you will see the average profits per minute of a player with certain strategy. For example, there will be a bar indicating the profit per minute of choosing the role of a Maker with the lowest spread and with speed ON, another bar indicating the profit per minute of choosing the role of a Maker with the lowest spread and with speed OFF, and so on. On the right side of the screen, you will see a pie chart that summarizes the frequency of the main role and speed choices in the last period. For example, if the Maker-Fast strategy makes up 20% of the pie, it means that, on average (during the whole duration of the trading period), 20% of players chose the role of a Maker and had speed switched to ON.

Session Earnings. Your final earnings for the session will be the sum of your show up fee (\$7.00) and the profits of one randomly chosen trading period, converted into US dollars.

FAQs.

Q1. Do I face the same five participants every period, or do they change?

A1. Yes, it is always the same five, chosen randomly at the beginning of the session.

Q2. If some other market maker has a narrower spread than mine, can I make money?

A2. Unlikely. When investors arrive, the makers with narrower spreads get to trade first. Other makers never get to trade profitably. (They will trade unprofitably if they get sniped, which can happen occasionally.)

Q3. Is this experiment a test of how fast I can react?

A3. No. The competition is intended to focus on which role to pick, whether to invest in speed and (if a market maker) what spread to choose. Since there is no way to predict the exact time of investor arrivals or jumps in V , **what matters is your average choice**, not how fast you react.

Experiment Instructions (CDA)

Welcome, and thank you for participating!

From now until the end of the experiment, please do not communicate with other participants. If you have any questions, please raise your hand; an experimenter will come and answer your question. Please pay careful attention to the instructions as real money is at stake.

During the experiment you will earn *Experimental Currency Units* (ECUs) and at the end of the experiment all your earnings will be converted to US dollars at the rate of one dollar for every two ECUs. You are guaranteed a show up fee of \$7.00 but can earn considerably more.

Basic Ideas

In this experiment you will participate in a simple **automated financial market**. Using information displayed on your screen as in Figure 1, you will **set and adjust trading algorithms (bots)** so as to earn as many ECUs as you can. As explained below, your earnings will depend on the settings you choose and on the choices of **the other five participants** in your group. There will be **eight trading periods**, each lasting **four minutes**, after which the experiment will end and you will be paid.

(removed)
Figure 1

Stock value. There is just one stock traded in your financial market. The stock has a *value* (henceforth “*V*”) that will jump up or down by a random amount at random times. The average time between jumps is about 1 second, but shorter or much longer times are possible. The Choice Box on your screen will always center on the current value of *V*, depicted by a horizontal gray line (see Figure 1). Every time there is a jump in *V*, there will be a yellow flash around the center line, and all other prices (described below) will appear to shift until they adjust to the new value of *V*.

Buy Orders, Sell Orders, and Spreads. All participants control algorithms (bots) that submit *buy orders* and *sell orders*. Submitting a *buy order for 99*, for example, means that you are willing to pay 99 ECUs or less to buy a share of the stock. Similarly, if you submit a *sell order for 101*, you are willing to receive 101 ECUs or more to sell a share of the stock. The price difference between your buy and sell orders is referred to as your *spread*. In the 99/101 example, your spread is 2.

Investor arrivals. At random times an automated investor arrives in the market to buy a share of the stock at the best available (lowest) selling price or to sell at the best available (highest) buying price. In today’s experiment, the average time between investor arrivals is about 2 seconds, but again shorter or much longer times are possible.

Your Choices and Earnings

Participants choose among three possible roles: *maker*, *sniper* and *out*. You can change your role at any time of the trading period, using the buttons on the right side of the *choice box* (see GIF 1a).

Maker. As a maker, you post both a buy order below V and a sell order above V using the computer program (henceforth, “your bot”) shown in GIF 1b. These orders are symmetrically located around V and are represented by the horizontal blue marks above and below the center line representing V . Your spread is the green or blue shaded region between those marks. To adjust your orders (and spread), simply click at the desired point in the choice box (see GIF 1b). Your bot will cancel your old buy and sell orders and send new orders to the market with the spread that you just chose. Since it takes time for the new orders to travel to the market (as explained below), your new orders are initially represented as horizontal lines at the right side of the choice box which slide toward the center of the box, at which time they replace the old orders.

You make a profit when an investor randomly arrives to buy if you are the maker with the lowest sell order (best price). Your bot automatically sells at your posted price and buys a replacement unit at the asset value V , so your profit is the vertical distance between the sell order and V . Similarly, you also make a profit when an investor randomly arrives to sell and you are the maker with the highest buy order (best price). Your bot will automatically buy at your posted price and sell it at the value V , so your profit is the difference between the value V and your buy order (see GIF 1c). Profits from transactions are represented by flashing green lines between V and the corresponding buy or sell order in the choice box, and are also displayed as vertical green lines in the profit box.

You can also lose money (earn a negative profit) while in the maker role. When V jumps up or down, your bot cancels your buy and sell orders and sends new orders to the market with prices that are centered on the new value of V . However, it takes time for your bot to send the orders and get them processed, and during that time you may be “sniped” by another participant. For example, suppose again that $V = 100$ and your current spread is 2. If V jumps to 105, your bot will send messages to the market cancelling your 99/101 orders and placing new buy and sell orders at 104 and 106. (In the choice box, you will see the blue marks representing your orders, and the shading between them, jump down relative to V . The blue marks and shaded box will then jump back to their original positions when your bot’s new orders are received, a fraction of a second later.) Before the new orders are processed, you are still committed to sell for 101 (cheap!) despite the fact the value of the asset is 105 -- we say that your orders are “stale”. You will be sniped if anyone places a buy order before your orders update. In this example, sniping would result in a loss of $105 - 101 = 4$ -- you sold too cheap. (see GIF 2b). Whenever you are sniped, the loss is the vertical distance between V and your old sell order. Similarly, if V jumps down, your stale buy order is vulnerable to being sniped. Losses from snipes are represented by flashing red lines between V and the corresponding buy or sell order in the choice box, and are likewise displayed as vertical red lines in the profit diagram.

| | | |
|-----------------------|-----------------------------|----------------------------|
| (removed) | (removed) | (removed) |
| GIF 1a Role Choice | GIF 1b Change Own Spread | GIF 1c Investor Arrival |

Sniper. If you choose to be a sniper, you attempt to transact with stale orders of makers. Whenever there is a jump event (yellow flash), your bot will send an order that tries to buy at a mispriced sell order (for an up jump in V) or sell to a mispriced buy order (for a down jump). As with maker orders, it takes time for sniper orders to reach the market. If your order is processed by the market before the stale orders are updated, you will transact and your choice box will flash green to indicate your profit (see GIF 2c).

| | | |
|------------------------|--|---------------------------------|
| (removed) | (removed) | (removed) |
| GIF 2a Value Change | GIF 2b Value Change, Getting Sniped | GIF 2c Value Change, Sniping |

Out. If you choose to be out, you do not participate in the market. In this role you never earn profits nor take losses.

Choosing Speed. As a *maker* or a *sniper*, it takes time to send your orders to the market. At **normal speed** (the default option) it takes half of a second to update orders. Alternatively, you can choose **fast speed** to reduce the update time to one tenth of a second, but it costs 0.02 ECUs per second for the faster service. You can activate or deactivate the speed option at any time. For example, if you activate speed now and deactivate it 30 seconds later, you are charged 0.60 ECUs, as you will see in the Profit Box. Choosing speed reduces the chance that a maker will get sniped and increases the chance that a sniper will be successful.

Tie Breaking. Whenever more than one sniper or maker is trying to update orders at the same time, the exchange will break ties randomly and execute the orders one at a time.

Trading Period Payoffs. You begin each *trading period* with an endowment of 20 ECUs. Gains and losses from the trading period will be added to and subtracted from your endowment to form your profits for the period.

Interim Screens. In between trading periods, your screen will display a summary of role and strategy choices as well as profits made in the last period (See Figure 3). On the left side of the screen, you will see the average profits per minute of a player with certain strategy. For example, there will be a bar indicating the profit per minute of choosing the role of a Maker with the lowest spread and with speed ON, another bar indicating the profit per minute of choosing the role of a Maker with the lowest spread and with speed OFF, and so on. On the right side of the screen, you will see a pie chart that summarizes the frequency of the main role and speed choices in the last period. For example, if the

Maker-Fast strategy makes up 20% of the pie, it means that, on average (during the whole duration of the trading period), 20% of players chose the role of a Maker and had speed switched to ON.

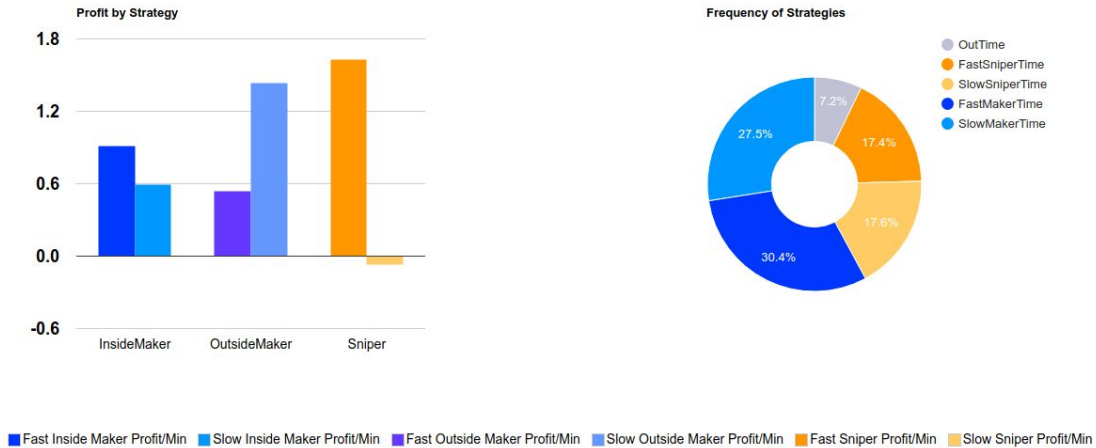


Figure 3: Interim Screens

Session Earnings. Your final earnings for the session will be the sum of your show up fee (\$7.00) and the profits of one randomly chosen trading period, converted into US dollars.

FAQs.

Q1. Do I face the same five participants every period, or do they change?

A1. Yes, it is always the same five, chosen randomly at the beginning of the session.

Q2. If some other market maker has a narrower spread than mine, can I make money?

A2. No. When an investor arrives, only the maker with narrowest spread gets to trade. Other makers never get to trade profitably. (They will trade unprofitably if they get sniped, which can happen occasionally.)

Q3. Is this experiment a test of how fast I can react?

A3. No. The competition is intended to focus on which role to pick, whether to invest in speed and (if a market maker) what spread to choose. Since there is no way to predict the exact time of investor arrivals or jumps in V , **what matters is your average choice**, not how fast you react.

References

- Abbink, K. and Brandts, J. (2005), “Price Competition Under Cost Uncertainty: A Laboratory Analysis,” *Economic Inquiry*, 43, 636–648.
- (2008), “Pricing in Bertrand competition with increasing marginal costs,” *Games and Economic Behavior*, 63, 1–31.
- Aldrich, E. M. and Friedman, D. (2017), “Order Protection through Delayed Messaging,” *Working Paper*, 1–43.
- Aldrich, E. M., Grundfest, J. A., and Laughlin, G. (2016), “The Flash Crash: A New Deconstruction,” *Working Paper*.
- Argenton, C. and Müller, W. (2012), “Collusion in experimental Bertrand duopolies with convex costs: The role of cost asymmetry,” *International Journal of Industrial Organization*, 30, 508–517.
- Barcelo, H. and Capraro, V. (2015), “Group size effect on cooperation in one-shot social dilemmas,” *Scientific Reports*, 5, 1–8.
- Budish, E., Cramton, P., and Shim, J. (2015), “The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response,” *The Quarterly Journal of Economics*, 130, 1547–1621.
- Davis, M. (2009), “Sex and the mathematician: The High School Prom Theorem,” *Games and Economic Behavior*, 66, 600.
- Dufwenberg, M. and Gneezy, U. (2000), “Price competition and market concentration: an experimental study,” *International Journal of Industrial Organization*, 18, 7–22.
- Embrey, M., Frechette, G. R., and Yuksel, S. (2018), “Cooperation in the Finitely Repeated Prisoner’s Dilemma,” *Quarterly Journal of Economics*, 133, 509–551.
- Friedman, B. D. and Oprea, R. (2012), “A Continuous Dilemma,” 102, 337–363.
- Friedman, D., Huck, S., Oprea, R., and Weidenholzer, S. (2014), “From imitation to collusion: Long-run learning in a low-information environment,” *Journal of Economic Theory*, 155, 185–205.
- Horstmann, N., Kramer, J., and Schnurr, D. (2015), “Oligopoly Competition in Continuous Time,” *Working Paper*.

- Kreps, D. M., Milgrom, P., Roberts, J., and Wilson, R. (1982), “Rational Cooperation in the Finitely Repeated Prisoners’ Dilemma,” *Journal of Economic Theory*, 27, 245–252.
- Mantovani, M. (2016), “Limited foresight in sequential games: an experiment,” *Working Paper*.
- Mengel, F. (2014), “Computer Games and Prosocial Behaviour,” *PLoS ONE*, 9, 1–5.
- Oechssler, J., Roomets, A., and Roth, S. (2016), “From imitation to collusion: a replication,” *Journal of the Economic Science Association*, 2, 13–21.
- Orzen, H. (2008), “Counterintuitive number effects in experimental oligopolies,” *Experimental Economics*, 11, 390–401.
- Park, I.-u. (2014), “Cooperation in Continuous Dilemma and Uncertain Reaction Lag,” *Working Paper*, 1–29.
- Potters, J. and Suetens, S. (2013), “Oligopoly Experiments in the Current Millennium,” *Journal of Economic Surveys*, 27, 439–460.
- Radner, R. (1986), “Can Bounded Rationality Resolve the Prisoners’ Dilemma?” in *Contributions to Mathematical Economics*, eds. Mas-Colell, A. and Hildenbrand, W., Amsterdam: North-Holland, chap. 20, pp. 387–399.
- Selten, R. and Stoecker, R. (1986), “End Behavior in Sequences of Finite Prisoner’s Dilemma Supergames: A Learning Theory Approach,” *Journal of Economic Behavior and Organization*, 7, 47–70.
- Simon, L. and Stinchcombe, M. B. (1989), “Extensive Form Games in Continuous Time: Pure Strategies,” *Econometrica*, 57, 1171–1214.