

Welfare Effects of Market Making in Continuous Double Auctions

Elaine Wah

*IEX Group, Inc., 4 World Trade Center
150 Greenwich St, 44th Floor, New York, NY 10007*

ELAINE.WAH@IEXTRADING.COM

Mason Wright

MASONDW@UMICH.EDU

Michael P. Wellman

*Computer Science & Engineering, University of Michigan
2260 Hayward St, Ann Arbor, MI 48109*

WELLMAN@UMICH.EDU

Abstract

We investigate the effects of market making on market performance, focusing on allocative efficiency as well as gains from trade accrued by background traders. We employ empirical simulation-based methods to evaluate heuristic strategies for market makers as well as background investors in a variety of complex trading environments. Our market model incorporates private and common valuation elements, with dynamic fundamental value and asymmetric information. In this context, we compare the surplus achieved by background traders in strategic equilibrium, with and without a market maker. Our findings indicate that the presence of the market maker strongly tends to increase total welfare across various environments. Market-maker profit may or may not exceed the welfare gain, thus the effect on background-investor surplus is ambiguous. We find that market making tends to benefit investors in relatively thin markets, and situations where background traders are impatient, due to limited trading opportunities. The presence of additional market makers increases these benefits, as competition drives the market makers to provide liquidity at lower price spreads. A thorough sensitivity analysis indicates that these results are robust to reasonable changes in model parameters.

1. Introduction

A *market maker* (MM) facilitates trade in a two-sided auction market by simultaneously maintaining offers to buy and sell. An ever-present MM supplies *liquidity* to the market. Liquidity refers to the availability of immediate trading opportunities at prices that reasonably reflect current market conditions. In compensation for liquidity provision, MMs profit from the *spread*, the difference between their buy and sell offers. MM activity is generally understood to stabilize prices and facilitate discovery of accurate prices in the market (Schwartz & Peng, 2013).

The exact role of market makers varies across market institutions. In a *pure dealer market*, multiple MMs competitively quote prices, and incoming market orders from investors trade at the best available MM price (Huang & Stoll, 1996). In a *pure limit-order market*, both investors and MMs submit orders with price limits, and whenever an incoming order matches an existing order, they trade at the incumbent order's limit price. This market mechanism is also called a *continuous double auction* (CDA), the name we use here. In a *specialist market*, there is a single MM designated to act as dealer, with an affirmative obligation to maintain fair and orderly markets (Saar, 2010). With the transition to electronic markets, pure limit-order markets are predominating (Frey & Grammig, 2006; Glosten, 1994), so this is the market mechanism we employ in our study.

Providing liquidity can generate profits from investors, but also runs the risk of *adverse selection*: when traders with newer or otherwise better information take advantage of the MM's standing offers. Much of the market making literature focuses on this tradeoff and its implications for MM strategies (Glosten & Milgrom, 1985; Kyle, 1985); other prior research has investigated the effects of MM on liquidity (e.g., as measured by price spreads) (Das & Magdon-Ismael, 2008; Eldor, Hauser, Pilo, & Shurki, 2006) and price discovery (Leach & Madhavan, 1992). Although liquidity and price discovery are generally expected to be positive factors for market performance and therefore welfare, there has been a notable dearth of prior research modeling this directly. Of the existing work addressing welfare, the focus has been on the need for an affirmative MM obligation due to adverse selection (Bessembinder, Hao, & Lemmon, 2011; Bessembinder, Hao, & Zheng, 2015), the cost structure of market participation in supplying liquidity (Huang & Wang, 2010), and trading mechanisms to incentivize market making (Brusco & Jackson, 1999).

The impact of MMs may depend on the number of market makers present. Prior works have considered the effect of MMs competing for the orders of other traders. These studies have examined market maker competition in models where background-trader orders are split across separate markets, each with an MM (Bernhardt & Hughson, 1997); where background traders can deal with each MM separately (Dennert, 1993); or, as in this study, where MMs compete in the same market through a common limit order book (Biais, Martimort, & Rochet, 2000; Glosten & Milgrom, 1985).

We investigate the effects of MM on market performance, focusing on allocative efficiency as well as gains from trade accrued by background investors. In our model, a single security is traded via the CDA mechanism in a market environment comprising multiple background traders, and in some cases one or more market makers. The fundamental value of the security evolves according to a mean-reverting stochastic process. An investor's value for units of the security is given by this fundamental plus an agent-specific private value that decreases in marginal value with the number of units held. The background traders enter and reenter according to a stochastic arrival process, each time to offer to buy or sell a single unit of the security. The stochastic arrival process can be interpreted to represent agents' perception of information or trading opportunities at random intervals. It forces agents to weigh the benefit of trading immediately at a known price against the uncertain possibility of a future arrival. MMs in our model have no private value, and thus aim to profit through maintaining buy and sell offers with a positive price spread.

To compare outcomes both with and without market making, we search for strategy configurations where traders best-respond to the environment and other-agent behavior. As analytic game-theoretic solution of this rich dynamic model appears intractable, we employ empirical simulation-based methods to derive equilibria over a restricted strategy space. For background traders, we consider parameterized strategies based on Zero Intelligence agents (Gode & Sunder, 1993). For the MM, we consider heuristic strategies loosely based on that defined by Chakraborty and Kearns (2011). From extensive simulation over thousands of strategy profiles, we estimate game models for various instances of the target scenario.

Analysis of the empirical games provides strong support for overall welfare benefits of market making in most settings. We derive empirical equilibria with and without market making in 30 environments, finding that the mix of background-trader strategies in equilibrium varies depending on the presence and strategy choice of the MM(s). In all of our environments, a single market maker is profitable in equilibrium, and in all but four equilibrium comparisons, the presence of MM increases overall welfare (background-trader surplus combined with MM profit). These findings are

robust to changes in many environment parameters, including the number of traders in the market, the distribution of private values, the duration of simulations, and the speed of background traders.

Whether market making benefits background traders (i.e., increases welfare net of MM profits) is more ambiguous, however. A single market maker made the investors better off in many environments tested, and was more likely to do so in relatively thin markets. For impatient investors with relatively infrequent trading opportunities, there is some evidence the MM is more beneficial in settings with a short trading horizon.

In contrast, we find that when multiple MMs compete, background traders always earn higher surplus than when there is a monopolist market maker or none at all. With two or four MMs present, every environment tested showed higher social welfare and greater background-trader surplus than the corresponding environment with zero or one MM (all other conditions being equal). Market makers' equilibrium strategies had narrower spreads in settings with multiple MMs, compared to settings with a single MM. These narrower spreads may partly account for the greater surplus achieved by background traders when multiple MMs are present.

In the next section we explain by way of example the potential role of market makers in alleviating allocative inefficiencies. We describe relevant work in Section 3. Section 4 discusses the market environment including background-trader and MM strategies. In Section 5 we present our empirical game-theoretic analysis, and we conclude in Section 6.

2. Motivating Example

We illustrate the problem of allocative inefficiency in CDAs, and the influence of market makers, with the following simple example. Suppose a market with four background traders: two buyers and two sellers. The buyers have values b_1 and b_2 , and seller values are s_1 and s_2 , with $b_1 > s_1 > b_2 > s_2$. Let us further assume for this illustration that the traders submit orders at their valuations.

Suppose that the orders arrive at the market in the order shown in Figure 1. Then buyer 1 trades with seller 1, and buyer 2 with seller 2, achieving a total surplus of $(b_1 - s_1) + (b_2 - s_2)$. The socially optimal allocation, in contrast, would have buyer 1 trading with seller 2, for a total surplus of $b_1 - s_2$. The difference between the optimal and achieved surplus is $\Delta = s_1 - b_2 > 0$. We can attribute this loss to the vagaries of the sequencing of limit orders, combined with the greedy matching implemented by the CDA mechanism. We choose to depict in the figure a sequence that leads to a suboptimal allocation; however, this is not the only one. In fact, only one-third of the possible orderings of these bids (8 out of 24) would result in the optimal allocation, with the remaining two-thirds under-performing by Δ .

Now suppose there is a market maker who continually maintains buy and sell offers in the auction, with difference δ between them. As long as the MM's offer to buy is within the interval (s_2, s_1) , and its offer to sell falls within (b_2, b_1) , then for this sequence of order arrivals, buyer 1 and seller 2 will trade with the MM, and the allocation will be efficient. If the MM quotes lie within the narrower interval of *competitive equilibrium* prices¹ $[b_2, s_1]$, then the efficient allocation is achieved for *any* sequence. In such cases, the MM accrues a profit of δ , with the remaining surplus divided among background traders.

The MM promotes efficiency in this example by providing liquidity to the market. In the absence of MM, when buyer 1 arrives, it has nobody to trade with. Seller 1 fills the vacuum and makes a

1. A competitive equilibrium price balances supply and demand with price-taking bidders. Here the balance is with respect to cumulative orders over the time horizon.




Figure 1: A sequence of CDA orders leading to a suboptimal allocation.

profitable trade with this buyer, but at a price far removed from that which would match supply and demand aggregated over time. An MM with quotes approximating this long-run price, in contrast, allows arriving bidders to trade near prevailing prices. Equally important, it prevents bidders who should not trade based on their valuations from doing so.²

Even assuming that the MM improves overall efficiency, does it make the background traders better off? In the specific scenario of Figure 1, the background traders benefit (in aggregate) if $\delta < \Delta$. If instead we consider the same set of four bids, but submitted in random order, then the background traders are clearly worse off in the third of instances where they would have achieved the efficient allocation without the MM's help. With random sequencing, the background traders benefit in expectation if and only if $\delta < \frac{2}{3}\Delta$.

More generally, we see that the question of whether MM presence is welfare-improving for background traders depends on specific details of the market setting. For background traders, the MM contribution may be sensitive to the distribution of valuations and bids, as well as their pattern of arrival over time. It also depends pivotally on the MM strategy—how well it tracks the prevailing market price and how large a spread the MM maintains between its buy and sell offers. In realistic environments, valuations include a combination of common and private elements and may evolve over time. Based on time and role, agents may have differential information about the common-value component. Thus for time-varying environments, we cannot assume the MM knows the underlying market equilibrium; it must instead act adaptively based on observations and statistical assumptions.

Moreover, individual traders may reenter the market to revise bids or reverse transactions, or to trade multiple units of the good. If such reentry were costless, market making would not be necessary to achieve allocative efficiency, as the traders could exchange among themselves to quiescence (Huang & Wang, 2010). As long as the traders do not indefinitely hold out for strict profits, the market would converge to an efficient allocation. In other words, *liquidity has economic value only to the extent that patience and market participation have costs or limits*.

With such complications, it seems unlikely we will be able to establish general analytical conditions for the benefits of MM. We therefore adopt a simulation approach, employing empirical game-theoretic techniques to search for strategically stable background-trader and MM strategies.

2. A modest amount of bid shading can also prevent inefficient trades, and indeed equilibrium shading strategies often lead to more efficient outcomes than truthful bidding in CDAs (Zhan & Friedman, 2007).

Our model includes all of the elements listed above, within an extensible framework that could incorporate (in future work) additional relevant features of financial markets.

3. Related Work

Literature on market making lies predominantly within the field of *market microstructure*, which examines the process by which prices, information, and transactions are formed by detailed interactions of traders in a market mechanism (Biais, Glosten, & Spatt, 2005; Madhavan, 2000; O'Hara, 1995). Early work focused on dealer markets, in which a monopolistic MM (the dealer) controls trading by acting as the middleman. Garman (1976) presents an explicit formulation of the market maker's optimization problem. O'Hara and Oldfield (1986) and Amihud and Mendelson (1980) concentrate on the impact of dealer inventory on spreads, while the seminal model of Glosten and Milgrom (1985) frames spreads as arising from adverse selection. Others focus on the consequences of informed trading on MM (Kyle, 1985; Chowdhry & Nanda, 1991; Das, 2008), as well as the role of market makers as liquidity providers (Grossman & Miller, 1988; Seppi, 1997).

The standard model of Glosten and Milgrom (1985) assumes perfect competition, constraining the MMs to set prices to achieve zero expected profit. Another line of research investigates models where MMs have some market power and may obtain positive profits. Such works typically derive the prices that would be set by rational market makers, either as monopolists (Das, 2008) or in oligopolistic competition with each other (Bernhardt & Hughson, 1997; Biais et al., 2000; Denert, 1993). In particular, the model of Biais et al. (2000) suggests that multiple market makers in imperfect competition will earn positive expected profit, but approach zero profit in the limit as the number of MMs goes to infinity.

Much of the relevant theoretical literature, however, relies on simplifying assumptions about MM behavior and trader interactions (Biais et al., 2005). Empirical studies have provided insight into the effects of market makers in real-world markets (Conrad, Wahal, & Xiang, 2015; Eldor et al., 2006; Frey & Grammig, 2006; Hasbrouck & Sofianos, 1993; Manaster & Mann, 1996; Menkveld, 2013; Sandás, 2001). Historical data alone, however, cannot elucidate the strategic choices faced by market participants. Agent-based modeling (ABM) and simulation of financial markets has proven conducive to exploring these questions (LeBaron, 2006); however, only a handful of ABM finance papers focus on market making (Chan & Shelton, 2001; Darley, Outkin, Plate, & Gao, 2000; Das, 2008).

Outside of microstructure, researchers have developed MM strategies for a variety of settings, including prediction markets (Hanson, 2007; Chen & Pennock, 2007; Abernethy, Chen, & Wortman Vaughan, 2011), dealer-mediated markets (Das, 2005; Jumadinova & Dasgupta, 2010), CDAs (Feng, Yu, & Stone, 2004), and environments where prices are generated exogenously (Abernethy & Kale, 2013). In this last category, Chakraborty and Kearns (2011) demonstrate the profitability of market making, given a mean-reverting price series. They propose a simple MM algorithm to submit a ladder of prices; the market makers we investigate can be viewed as variations on this strategy.

None of these studies, however, addresses questions about allocative efficiency in the market. Existing studies are largely concerned with how adverse selection affects allocative efficiency. For example, Bessembinder et al. (2011) demonstrate that restricting spread widths improves allocative efficiency and encourages more traders to become informed. Their results suggest that MMs enhance efficiency primarily when information asymmetries are significant. Brusco and Jackson

(1999) illustrate the inefficiencies of competitive markets in a two-period model in which the market maker position is designated via an auction. They also design a system of trading rules to reach an efficient allocation by identifying and incentivizing MM agents. Huang and Wang (2010) propose a model in which provision of liquidity is endogenous, finding that mandating participation tends to improve welfare, but that the welfare effects of lowering costs for liquidity provision per se are ambiguous. In a similar vein, Bessembinder et al. (2015) present a model in which a firm can sell an asset to an investor in an IPO, with the option of paying a designated market maker (DMM) in exchange for liquidity provision in a secondary market. When the secondary market is illiquid due to asymmetric information and uncertainty regarding the asset's fundamental value, social welfare can be improved if the firm enters into a DMM contract. Menkveld and Wang (2013) provide empirical evidence for DMM benefits in otherwise illiquid markets for small-cap stocks.

4. Market Environment

To investigate the effect of market making on allocative efficiency, we construct a simple model of a single security traded in a continuous double auction market. Prices are fine-grained but discrete, taking values at integer multiples of the *tick size*, p_{ts} . Time is likewise fine-grained and discrete, with finite horizon T . The horizon places pressure on agents to accomplish trades quickly, thus modulating T allows us to study the impact of varying immediacy preference. Agents arrive at designated times, and submit limit orders to the market. The CDA maintains price quotes reflecting the best outstanding orders. BID_t corresponds to the highest buy offer at time t , and ASK_t to the lowest offer to sell. Other bids in the order book are not visible to traders. The market environment is populated by multiple background traders, representing investors, and (optionally) one or more market makers. At any time, the background investors are restricted to placing a single order to buy or sell one unit, whereas the MM may maintain orders to buy and sell any number of units at various prices.

The source code of the financial market simulator employed for this study is publicly available online (Strategic Reasoning Group, 2016).

4.1 Valuation Model

Each background trader has an individual valuation for the security composed of private and common components. We denote by r_t the common *fundamental value* for the security at time t . The fundamental time series is generated by a mean-reverting stochastic process:

$$r_t = \max \{0, \kappa \bar{r} + (1 - \kappa) r_{t-1} + u_t\}.$$

Parameter $\kappa \in [0, 1]$ specifies the degree to which the fundamental reverts back to the mean \bar{r} , and $u_t \sim \mathcal{N}(0, \sigma_s^2)$ is a random shock at time t .

The private component for agent i is a vector Θ_i representing differences in private benefits of trading given the trader's net position, similar to the model of Goettler, Parlour, and Rajan (2009).³ The vector is of size $2q_{\max}$, where $q_{\max} > 0$ is the maximum number of units the agent can be long

3. Goettler et al. assume an infinite horizon and employ exponential discounting to represent the cost of trading delays. We adopt a finite horizon, which provides an alternative force for trading urgency. The stochastic agent arrivals in combination with a finite trading horizon capture, in effect, the role of discounting in an infinite horizon model.

or short at any time, with

$$\Theta_i = \left(\theta_i^{-q_{\max}+1}, \dots, \theta_i^0, \theta_i^{+1}, \dots, \theta_i^{q_{\max}} \right).$$

Element θ_i^q is the incremental private benefit foregone by selling one unit of the security given current position q , where positive (negative) q indicates a long (short) position. Similarly, θ_i^{q+1} is the marginal private gain from buying an additional unit given current net position q .

We generate Θ_i from a set of $2q_{\max}$ values drawn independently from a Gaussian distribution. Let $\hat{\theta} \sim \mathcal{N}(0, \sigma_{PV}^2)$ denote one of these drawn values. To ensure that the valuation reflects diminishing marginal utility, that is, $\theta^{q'} \geq \theta^q$ for all $q' \leq q$, we sort the $\hat{\theta}$ in non-increasing order and set the θ_i^q to respective values in the sorted list.

Background trader i 's valuation v for the security at time t is based on its current position q_t and the value of the global fundamental at time T , the end of the trading horizon:

$$v_i(t) = r_T + \begin{cases} \theta_i^{q_t+1} & \text{if buying 1 unit} \\ \theta_i^{q_t} & \text{if selling 1 unit.} \end{cases}$$

To be clear, a trader does not know the final fundamental value r_T at any time $t < T$, though it can form an estimate based on its actual observations, as described below (2).

For a single-quantity limit order transacting at time t and price p , a trader obtains surplus:

$$\begin{cases} v_i(t) - p & \text{for buy transactions, or} \\ p - v_i(t) & \text{for sell transactions.} \end{cases}$$

Since the price and fundamental terms cancel out in exchange, the total surplus achieved when agent B buys from agent S is $\theta_B^{q(B)+1} - \theta_S^{q(S)}$, where $q(i)$ denotes the pre-trade position of agent i .

The total payoff to an agent over a game instance equals the net cash taken in from trades, plus the value of final holdings q_T , accounting for both common (fundamental) and private components of valuation:

$$c_T + q_T r_T + \begin{cases} \sum_{j=1}^{q_T} \theta_i^j & \text{if } q_T > 0 \\ 0 & \text{if } q_T = 0 \\ \sum_{j=q_T+1}^0 -\theta_i^j & \text{if } q_T < 0, \end{cases} \quad (1)$$

where c_T is the agent's cash holdings at time T .

4.2 Background Trading Strategies

There is an extensive literature on autonomous bidding strategies for CDAs (Friedman, 1993; Das, Hanson, Kephart, & Tesauro, 2001; Wellman, 2011). In this study, we consider trading strategies in the so-called *Zero Intelligence* (ZI) family (Gode & Sunder, 1993).

The background traders arrive at the market according to an independent Poisson process per trader, with rate λ_{BG} . On arrival, they are assigned to buy or sell with equal probability, and accordingly submit an order to buy or sell a single unit. (A trader is randomly reassigned to buy or to sell each time it arrives.) Background traders subsequently reenter the market, with time between entries distributed exponentially at the same rate λ_{BG} —in other words, each trader reenters the market by an independent Poisson process, just like its arrival process. Agents may trade once per

arrival and any number of times overall, as long as their net positions do not exceed q_{\max} (either long or short). Background traders are notified of all transactions and current price quotes with zero delay, and may use this information in computing their bids.

The randomized agent arrival process serves several useful functions in a financial market model (Goettler et al., 2009). It provides each agent a noisy signal of the security's true value at random intervals, reflecting the impossibility of an agent detecting every favorable trading opportunity. It presents the agent with a choice between trading immediately at a known price or waiting for an uncertain later arrival. In aggregate, it generates a stochastic sequence of orders to the limit order book. Relative rates of arrival for background traders and MMs (λ_{BG} versus λ_{MM}) further serve to capture speed differences between these two categories of trader.

A ZI trader assesses its valuation $v_i(t)$ at the time of market entry t , using an estimate \hat{r}_t of the terminal fundamental r_T . The estimate is based on the current fundamental, r_t , adjusted to account for mean reversion:

$$\hat{r}_t = (1 - (1 - \kappa)^{T-t}) \bar{r} + (1 - \kappa)^{T-t} r_t. \quad (2)$$

The ZI agent then submits a bid shaded from this estimate by a random offset—the degree of surplus it demands from the trade. The amount of shading is drawn uniformly from range $[R_{\min}, R_{\max}]$. Specifically, a ZI trader i arriving at time t with current position q submits a limit order for a single unit of the security at price

$$p_i \sim \begin{cases} \mathcal{U} [\hat{r}_t + \theta_i^{q+1} - R_{\max}, \hat{r}_t + \theta_i^{q+1} - R_{\min}] & \text{if buying} \\ \mathcal{U} [\hat{r}_t + \theta_i^q + R_{\min}, \hat{r}_t + \theta_i^q + R_{\max}] & \text{if selling.} \end{cases}$$

We extend ZI by including a threshold parameter $\eta \in [0, 1]$, whereby if the agent can achieve a fraction η of its requested surplus at the current price quote, it submits a marketable order that simply takes that quote. Technically, the agent still posts a limit order, but at a price that ensures the order will transact immediately rather than get added to the order book. Setting $\eta = 1$ is equivalent to the strategy without employing the threshold.

Note that although the ZI strategy as originally defined by Gode and Sunder (1993) was not tuned to its environment, we implicitly assume that a community of ZI agents would eventually select strategy parameters R_{\min} , R_{\max} , and η to reach a Nash equilibrium in their environment. Each agent will play a fixed parameterization of ZI, but the community of agents will play a Nash equilibrium mixed strategy over ZI parameters.

In our model, background traders are permitted to reenter the market. Upon each reentry, the trader withdraws its previous order (if not transacted yet) before executing its extended ZI strategy described above.

4.3 Market Maker Strategies

Much of the prior work on MM strategies treats the market maker as a dealer (Glosten & Milgrom, 1985; Das, 2005), which must take one side of each trade. In our model, however, all trades execute through the CDA order book, therefore the MM submits limit orders just as background traders do. We consider a family of MM strategies that submit at time t a *ladder* of unit-quantity buy and sell orders, composed of K rungs spaced ξ ticks apart:

$$\begin{cases} [S_t, S_t + \xi, S_t + 2\xi, \dots, S_t + (K - 1)\xi] & \text{for sell orders} \\ [B_t - (K - 1)\xi, \dots, B_t - 2\xi, B_t - \xi, B_t] & \text{for buy orders} \end{cases}$$

with $S_t > B_t$ and $K, \xi > 0$. Each MM arrives at time 0 and reenters the market according to its own, independent Poisson process with rate λ_{MM} . On reentry at time t , the MM observes the current fundamental r_t , which it may use in determining its ladder of buy and sell orders. It cancels any standing orders remaining from its previous ladder when submitting a new ladder.

Like the background traders, the MM values its inventory at the end of the trading horizon at the global fundamental r_T . The MM's total profit is defined by the sum of trading cash flow plus the value of holdings. The total payoff for an MM is thus the same as for a background trader (1), but without the private value term.

To avoid crossing the current BID - ASK quote, the MM truncates its ladder. Specifically, if $BID_t > S_t$ (or similarly, $B_t > ASK_t$), the agent cuts the ladder off at the rung that is at or above (below) the current BID (ASK) price. The truncated ladder is:

$$\begin{cases} [S_t + (K - x)\xi, \dots, S_t + K\xi] & \text{if } BID_t > S_t \\ [B_t - K\xi, \dots, B_t - (K - x)\xi] & \text{if } B_t > ASK_t, \end{cases}$$

where $x > 0$ specifies the rung immediately above BID (for sell orders) or below ASK (for buy orders). That is, x satisfies the condition $S_t + (K - x - 1)\xi < BID_t < S_t + (K - x)\xi$ for sell orders in the ladder, and $B_t - (K - x)\xi < ASK_t < B_t - (K - x - 1)\xi$ for buy orders.

The MM uses its observation of the current fundamental r_t to inform its ladder construction. Specifically, the MM strategies we implement compute an estimate \hat{r}_t of the terminal fundamental r_T via (2), and center the ladder around this estimate. The spread ω is set by a strategy parameter. The central ladder prices are:

$$S_t = \hat{r}_t + \frac{1}{2}\omega, \quad B_t = \hat{r}_t - \frac{1}{2}\omega.$$

A more adaptive market maker strategy would adjust its spread ω in response to changing conditions. We experimented with simple adaptive strategies, and did not find any that improved over those with fixed spread parameter. Nevertheless, it could well be that more sophisticated adaptation strategies or other variations could outperform the MM strategies employed here.

5. Empirical Game-Theoretic Analysis

We have described strategies for the roles of background trader and MM, each with tunable parameters. Evaluating the effect of market making for all combinations of strategy choices would be infeasible; moreover, the various strategic contexts are not equally relevant. Generally speaking, we are most interested in the effect of market making when all agents are doing their best to generate profit. In other words, we wish to evaluate the impact of MM *in equilibrium*—where both the background traders and MM are adopting the best strategies, given the environment and other agent strategy selections.

We qualify our equilibrium analysis in two ways. First, we consider only a restricted set of available strategy choices, defined by selected parameterized versions of the strategies introduced above. These strategies were selected through a manual heuristic process (described in Section 5.2), which sought to produce high-fitness strategies with broad coverage of the strategy space. Second, we determine equilibria among these strategies through a simulation-based process, known as *empirical game-theoretic analysis* (EGTA) (Wellman, 2016). In EGTA, we use systematic simulation of strategy profiles in a specified environment to induce a game model of that environment. For the present

study, we simulate an instance of the financial market described in Section 4, using an extension of the discrete-event market simulation system developed for our previous study of latency arbitrage (Wah & Wellman, 2013). We generate data for various combinations of the strategies introduced in Sections 4.2 and 4.3, each sampled over many runs (at least 20,000 per profile, often many more) to account for stochastic effects (valuation schedules, trajectories of the market fundamental, agent arrival patterns). From these data we estimate game payoffs and derive equilibria with respect to the strategy space explored. We then take these equilibria as the basis for evaluating MM welfare effects.

The experiments conducted for the present study supersede those reported in our prior paper (Wah & Wellman, 2015). The present results incorporate an expanded strategy set (Table 1) and subtle changes to the background-trader arrival process (Section 4.2). We also more thoroughly sample the profile space, covering more profiles and with more simulations per profile. The results are qualitatively consistent with our previous findings, though with a more ambiguous relationship between trading horizon and MM impact on surplus gains. Our expanded strategy set includes background traders who persist in strict shading indefinitely, which affords greater scope for MM benefit even over long trading horizons.

5.1 Environment Settings

We evaluate the performance of background traders and the MM within 30 parametrically distinct environments. For each environment, we analyze two empirical games that differ in whether or not an MM is present; in some environments, we also analyze games with two or four MMs present. In all settings, there are $N \in \{25, 66\}$ background traders. Each simulation run lasts T time steps, for $T \in \{1, 4, 12, 24\} \times 10^3$. Simulations with lower T represent market settings where traders have a greater desire for trading immediacy. If present, the MM in each environment enters the market at the start of the simulation and reenters with rate $\lambda_{MM} = 0.005$, or approximately once every 200 time steps. The global fundamental has a mean value $\bar{r} = 10^5$ and mean-reversion parameter $\kappa = 0.05$. The minimum tick size p_{ts} is fixed at 1. The maximum number of units the background trader can be long or short at any time is $q_{\max} = 10$.

The environments differ in number of background traders (N), background-trader reentry rate (λ_{BG}), fundamental shock variance (σ_s^2), and time horizon (T). Environments with a low background-trader arrival rate λ_{BG} represent market settings where MMs have a speed advantage relative to other traders. The configurations of parameter settings for $N \in \{25, 66\}$ background traders and $T \in \{1, 4, 12, 24\} \times 10^3$ are as follows.

A $\lambda_{BG} = 0.0005, \sigma_s^2 = 1 \times 10^6, \sigma_{PV}^2 = 5 \times 10^6$

B $\lambda_{BG} = 0.005, \sigma_s^2 = 1 \times 10^6, \sigma_{PV}^2 = 5 \times 10^6$

C $\lambda_{BG} = 0.005, \sigma_s^2 = 5 \times 10^5, \sigma_{PV}^2 = 5 \times 10^6$

In addition, we considered two variations on environment A that increase and decrease (resp.) the private value variance (σ_{PV}^2) by a factor of two. Modulating this parameter is tantamount to changing the weighting in importance between private and common components of background-trader valuation. For these variants, we investigated configurations with $N = 25$ and $T \in \{1, 4, 12\} \times 10^3$.

D $\lambda_{BG} = 0.0005, \sigma_s^2 = 1 \times 10^6, \sigma_{PV}^2 = 1 \times 10^7$

$$\mathbf{E} \lambda_{BG} = 0.0005, \sigma_s^2 = 1 \times 10^6, \sigma_{PV}^2 = 2.5 \times 10^6$$

We describe each environment by its configuration label, followed by time horizon (in thousands). For example, B12 is the environment labeled \mathbf{B} above with $T = 12000$.

5.2 EGTA Process

We model our market as a *role-symmetric game*, in which players are partitioned into roles, each role with its own strategy set and payoff function. The payoff for playing a strategy in a particular role depends on the number of other agents playing each strategy in each role, but not on how the strategies are mapped to players within the roles. The two roles in our model are background trader (25 or 66 players) and market maker (0, 1, 2, or 4 players). The background-trader strategy set \mathcal{S} comprises thirteen parameterized versions of the ZI strategy, as shown in Table 1. The MM strategy set includes the seven MM strategies shown in Table 4. The payoff function for each agent is defined by (1).

We constructed the strategy set for each role through a heuristic, iterative process, starting from a few instances of parameterized strategy families from prior work. We started by running pilot simulations using an initial set of parameterized strategies. Based on observations of which strategies from the initial set were in the support (i.e., played with positive probability) of Nash equilibria in different environments, we augmented the strategy set with new strategies that appeared promising. For example, if the initial ZI strategy with the lowest R_{\max} appeared with high probability in many equilibria, we might introduce a new ZI strategy with a lower R_{\max} value. In this way, we iteratively adjusted the strategy set until it appeared to provide ample opportunities for agents to adapt to the different market conditions of our various environments.

Even if we exploit symmetry, the game size (number of distinct strategy profiles) grows exponentially with the number of players and strategies, so it is computationally infeasible to sample every profile in games with many players and strategies. For example, in a game with $N = 25$ background traders, no MMs, and $|\mathcal{S}| = 13$ background-trader strategies, there are $\binom{N+|\mathcal{S}|-1}{N} = \binom{25+13-1}{25} \approx 1.9 \times 10^9$ possible symmetric strategy profiles—far too many to evaluate via simulation. We therefore apply aggregation to approximate a many-player game as a game with fewer players, which we call a *reduced game*.

The technique we use, called *deviation-preserving reduction* (DPR) (Wiedenbeck & Wellman, 2012), defines reduced-game payoffs in terms of payoffs in the original game as follows. Consider an N -player symmetric game, reduced to a k -player game, for $k < N$. The payoff for playing strategy s_1 in the reduced game, with other agents playing strategies (s_2, \dots, s_k) , is given by the payoff of playing s_1 in the original N -player game when the other $N - 1$ agents are evenly divided ($\frac{N-1}{k-1}$ each) among strategies s_2, \dots, s_k . The reduction is termed “deviation-preserving” because the payoff to one player switching strategies in a reduced profile is the same with respect to a deviation in the corresponding original profile.⁴

To see how DPR reduces the number of samples required, note that if we reduce a game of $N = 25$ players to $k = 5$, with $|\mathcal{S}| = 13$ strategies, we now require only $\binom{k+|\mathcal{S}|-1}{k} = \binom{5+13-1}{5} = 6188$ (reduced-game) profiles. And if we limit the search space to four strategies, there are only $\binom{5+4-1}{5} = 56$ profiles.

4. Although DPR provides no general guarantees, we have found in several settings that equilibrium profiles in the reduced game often have low regret in the original game (Wiedenbeck & Wellman, 2012).

We chose values for N in this study to facilitate DPR by ensuring that the required aggregations come out as integers: the approximation of an $(N, 1)$ -size game (i.e., N background traders and 1 MM) by a $(k, 1)$ -player reduced game works best when k divides N and $k - 1$ divides $N - 1$. Specifically, we use simulation data from the $(66, 1)$ -agent environments to estimate reduced $(6, 1)$ -player games, where six players represent the 66 background traders in the simulated environment. We similarly estimate $(5, 1)$ -player games from the $(25, 1)$ -agent cases.

For an example of DPR in action, consider a game with $N = 25$ background traders and no MM. We can map this to a reduced game of $k = 5$ background traders. The focal player in the reduced game is said to control one agent in the original game, and the other four players each control $\frac{N-1}{k-1} = 6$ original-game agents. The reduced-game payoff function maps each strategy $s \in \mathcal{S}$ for the focal player and vector of realizable counts c_i of how many of the other four agents play each strategy $i \in \mathcal{S}$, to the payoff for playing s in this setting. The realizable count vectors have $c_i \in \{0, \dots, k-1\}$ for all i , and $\sum_{i=1}^{|\mathcal{S}|} c_i = k-1$. The empirical payoff for (s, \vec{c}) is estimated by averaging the results from simulations of the associated profile where $1 + c_s \times \frac{N-1}{k-1}$ play strategy s , and for each $i \neq s$, $c_i \times \frac{N-1}{k-1}$ play strategy i .

DPR works similarly in a game with MMs. For example, from an original game with $N = 25$ background traders and two MMs, DPR profiles are constructed both from the perspective of a background trader as focal player, and with an MM as focal. In either case, the reduced game has $k = 5$ background traders and two MMs. If a background trader is the focal player in the reduced game, it controls one background trader in the original; the other four control $\frac{N-1}{k-1} = 6$ background traders each in the original game, and each reduced-game MM controls one original-game MM. If an MM is the focal player in the reduced game, it controls one MM in the original game; the other reduced-game MM controls the other original-game MM; and each of the five background traders in the reduced game controls $\frac{N}{k} = 5$ original-game background traders. This also demonstrates why it is useful for k to divide N and for $(k - 1)$ to divide $(N - 1)$.

We iteratively apply EGTA to guide our exploration of the strategy space. We start by simulating all the role-symmetric pure profiles, where a single strategy is shared by all players in a role. Exploration then spreads to the explored profiles' neighbors, that is, those profiles related by a single-player deviation (one player switching strategies) to some profile explored already. The goal in this process is to identify Nash equilibria, and we focus our search on *role-symmetric Nash equilibria* (RSNE). A role-symmetric profile assigns a mixed strategy (probability of playing each action) to every role, such that all agents of a role have the same mixed strategy. The profile is an RSNE if no agent can benefit in expectation by unilaterally switching to a different strategy.

As the observed payoffs from our simulator are incrementally added to the empirical game's payoff matrix, we compute mixed-strategy equilibria for each completed subgame. A *subgame* is a mapping from each role to a nonempty subset of that role's strategy set. A subgame is *complete* if we have obtained at least some minimum number of samples (e.g., 20,000) for every DPR strategy profile that uses only strategies contained in the subgame.

Given a complete subgame, we can compute its role-symmetric Nash equilibria. The equilibria of the complete subgames are initially *candidates* for equilibria in the full game, which we can either *refute* by finding a beneficial deviation outside the subgame's strategy set, or *confirm* by examining all deviations without finding any beneficial deviation. It is feasible to evaluate all possible deviations for a subgame, as these are bounded in number by the size of the subgame times the size of the strategy set. To find the value for a single agent of deviating to pure strategy s while other

agents play profile \vec{c} , it is sufficient to sample the DPR profiles for the subgame including strategies in the support of \vec{c} , along with strategy s .

We continue to refine the empirical game with additional simulations, until the following conditions are met:

1. at least one equilibrium is confirmed,
2. all non-confirmed candidates are refuted (up to a threshold support size—in this study three background strategies and unlimited MMs), and
3. for all refuted candidates (up to the threshold support size), we have explored subgames formed by adding the best response to the candidate’s support.

In this study, we successfully find at least one and at most four non-trivial RSNEs for each game evaluated, with support sizes up to four for background traders and up to two for MMs.

We utilize the EGTAOnline infrastructure (Cassell & Wellman, 2013) for conducting and managing our experiments, and run our simulations on a high-performance computing cluster at the University of Michigan. The process accumulates a dataset of profile simulation results, which we use to estimate payoff values for strategy profiles in the game.

For all the games we model, there exists a trivial pure RSNE in which all agents play a “NOOP” strategy that refrains from bidding. This exists because if none of the other agents (background traders or MM) submit limit orders, then there is nobody to trade with and there will be no transactions regardless of the strategy the subject agent employs. In our discussion below, we ignore this degenerate equilibrium, which obviously has payoff zero for all agents.

To provide a benchmark for efficiency, we calculate the social optimum based on the trader population and valuation distribution used in our environments A, B, and C⁵ (i.e., $N \in \{25, 66\}$ background traders with parameters $q_{\max} = 10$ and $\sigma_{PV}^2 = 5 \times 10^6$). We determine the optimum for a particular draw of N valuation vectors by treating each as a demand curve and finding a uniform competitive equilibrium price. This is conveniently implemented in our simulation environment, where valuation vector Θ_i is represented by a background trader i , who submits q_{\max} single-unit sell orders at prices $\bar{r} + \theta_i^s$, $s \in \{-q_{\max} + 1, \dots, 0\}$, and q_{\max} single-unit buy orders at prices $\bar{r} + \theta_i^b$, $b \in \{+1, \dots, q_{\max}\}$. A call market computes a uniform clearing price to match supply and demand, which defines the optimal allocation for the sample. From 20,000 samples, we find a mean social welfare of 44155 and 16306 for 66 and 25 background traders, respectively. Figure 2 presents histograms of trades per background trader in the social optima.

5.3 Game without Market Making

The empirical games without MM cover 14 background-trader strategies: 13 versions of ZI (see Table 1), and (implicitly) the no-trade strategy NOOP. We identified 1–4 ZI equilibria for each of our 30 environments (see Tables 2 and 3). For each equilibrium, we estimated background-trader surplus by sampling 2,500 profiles according to the equilibrium mixture, running 25–100 simulations per sampled profile (at least 62,500 simulations in total) and then recording the aggregate surplus.

5. Environments D and E employ different weighting between private and common value, and hence have different social optima.

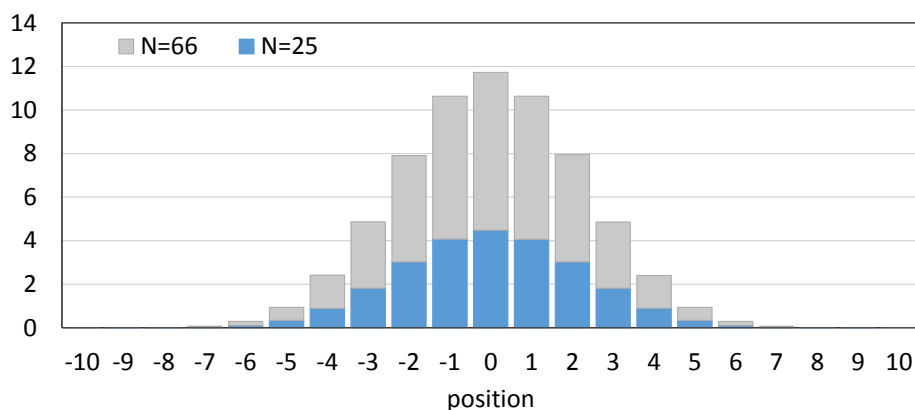


Figure 2: Distributions of the net position (equivalently, number of units traded) of N background traders in socially optimal allocations for environments A, B, and C. The histograms (shown superimposed) are compiled from 20,000 samples.

Table 1: ZI strategy combinations included in empirical game-theoretic analysis.

R_{\min}	R_{\max}	η
0	65	0.8
0	125	0.8
0	125	1
0	250	0.8
0	250	1
0	500	1
0	1000	0.8
0	1000	1
0	1500	0.6
0	2500	1
250	500	1
500	1000	0.4
1000	2000	0.4

5.4 Game with Market Making

Our games with MM include the 14 background-trader strategies from the no-MM treatment above, plus seven strategies for the MM role. The MMs employed in our game analysis are as described in Section 4.3, with $K = 100$ rungs spaced $\xi \in \{25, 50, 100\}$ units apart. Each MM strategy type employs a fixed spread $\omega \in \{64, 128, 256, 512, 1024\}$. The set of all MM strategies employed is in Table 4. The equilibria found are presented in Tables 5 and 6. Background-trader surplus and MM profit are estimated for each equilibrium based on the sampling method described for the no-MM game above.

Table 2: Symmetric equilibria for games without market makers, $N = 66$, calculated from the 6-player DPR approximation. Each row of the table describes one equilibrium found and its average values for total surplus and two strategy parameters: R_{mid} (the midpoint of ZI range $[R_{\text{min}}, R_{\text{max}}]$) and threshold η . Values presented are the average over strategies in the profile, weighted by mixture probabilities. Surplus values are means from thousands of simulations, where strategies are randomly sampled from the equilibrium mixed-strategy profile.

Env	Surplus	R_{mid}	η
A1	3712	750	0.4
A1	4439	374	0.980
A4	16578	340	0.977
A4	16551	353	1
A12	33741	267	0.955
A24	42361	125	0.8
A24	42413	107	0.941
A24	41563	229	0.972
B1	29150	441	0.894
B4	40392	411	0.961
B12	40102	494	0.810
B24	40170	497	0.806
B24	40100	677	0.576
C1	30803	250	1
C1	29726	375	1
C4	41130	500	0.8
C4	39901	390	0.976
C12	41410	446	0.923
C24	43021	501	0.799
C24	41531	416	0.939

5.5 Comparison of Market Performance

Our findings with regard to the central question in this paper are presented in Figures 3 and 4. For each environment, we compare equilibrium outcomes, with and without an MM, on two measures: social welfare and background-trader surplus. Since there are often multiple equilibria, the differences are presented as ranges, delimiting the most and least favorable comparisons.

In the scenarios with 66 background traders (Figure 3(a)), the change in overall welfare is generally positive, with four environments (B4, C1, C4, and C24) providing small exceptions. The change in background-trader surplus, in contrast, varies widely across environments, with multiple examples of both positive and negative changes. The effect is strongly negative in the A environments with longer trading horizons, which may be explained by the significant information advantage of MMs over background traders due to their disparate reentry rates ($\lambda_{MM} = 0.005$ versus $\lambda_r = 0.0005$). For environments B and C with $T \geq 4000$, the total social welfare without MM is

Table 3: Symmetric equilibria for games without market makers, $N = 25$, calculated from the 5-player DPR approximation. Data are presented as in Table 2.

Env	Surplus	R_{mid}	η
A1	1041	750	0.404
A1	1351	371	1
A4	5616	350	0.986
A12	11669	335	1
A24	13697	375	1
A24	15162	218	0.949
A24	15543	117	0.826
B1	8752	750	0.4
B4	14041	517	0.773
B12	14256	553	0.715
B24	14478	556	0.710
C1	10379	375	1
C4	14225	476	0.838
C12	14617	441	0.894
C24	14618	490	0.816
D1	2120	366	0.994
D1	1926	690	0.496
D4	7894	716	0.454
D12	17334	392	0.972
E1	1055	125	0.884
E4	4110	178	1.000
E12	8696	160	0.999
E12	8891	125	0.886

Table 4: MM strategy parameter combinations explored.

K	ξ	ω
100	25	256
100	50	64
100	50	128
100	50	256
100	50	512
100	50	1024
100	100	512

over 89% of the socially efficient outcome of 44155. That is, the ZI background traders in these environments extract a high fraction of the potential surplus in the market on their own. Intuitively, given sufficient time for reentry (as governed by horizon T and reentry rate λ_r), agents with pri-

Table 5: Role-symmetric equilibria for games with one market maker, $N = 66$, calculated from the (6, 1)-player DPR approximation. Each row of the table describes one equilibrium found and its average values for background-trader surplus, MM profit, and four strategy parameters: R_{mid} (the midpoint of ZI range $[R_{\text{min}}, R_{\text{max}}]$), threshold η , MM spread ω , and rung size ξ . Values presented are the average over strategies in the profile, weighted by mixture probabilities.

Env	Surplus	Profit	R_{mid}	η	ω	ξ
A1	4545	461	238	0.933	512	61
A1	4553	485	205	1.000	512	100
A1	4465	382	298	0.943	512	50
A4	16503	1890	139	0.834	256	50
A4	16768	1704	135	1.000	256	25
A4	16517	1953	132	0.810	256	25
A12	32765	3279	135	0.815	256	48
A24	39367	3668	110	0.976	256	50
A24	39180	3912	105	0.800	256	25
A24	38817	4249	118	0.802	256	25
B1	29238	21	433	0.907	931	59
B4	40042	196	431	0.942	512	100
B12	40575	113	400	0.974	512	100
B12	42304	820	492	0.806	491	50
B24	41631	1196	550	0.725	256	50
B24	42379	1253	589	0.857	256	50
B24	40693	374	480	0.864	512	100
B24	41150	2451	807	0.645	256	25
C1	29507	302	375	1.000	512	50
C4	39669	878	421	0.954	256	50
C4	42248	626	250	1.000	256	25
C12	41658	2003	500	0.864	256	50
C12	40836	1233	431	0.911	256	50
C12	42037	1572	455	0.976	256	50
C24	42912	489	488	0.925	256	50
C24	42964	702	418	0.887	256	25
C24	42236	37	440	0.897	256	25


vate values on the correct side of competitive prices will eventually trade, and any inefficient trades can effectively be reversed. When the background traders have sufficient time to reach efficient outcomes, the MM may provide little benefit to overall welfare, and its profits tend to come out of background-trader surplus. Accordingly, we observe that the MM degrades investor surplus for some or all equilibria in four of these six cases.

The trading horizon T reflects whatever might limit an investor's patience (liquidity needs, portfolio hedging, cost of monitoring, etc.). By curbing agents' ability to find efficient trades, the time

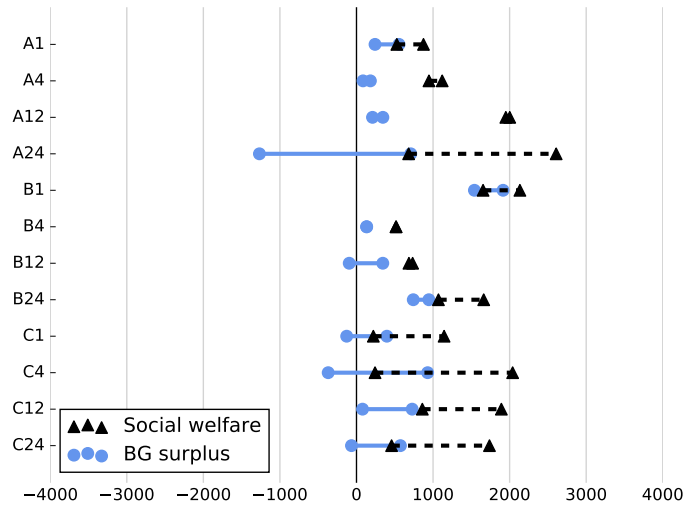
Table 6: Role-symmetric equilibria for games with one market maker, $N = 25$, calculated from the $(5, 1)$ -player DPR approximation. Data are presented as in Table 5.

Env	Surplus	Profit	R_{mid}	η	ω	ξ
A1	1593	285	164	0.863	512	50
A1	1596	320	115	0.940	512	100
A4	5711	852	188	0.896	512	100
A4	5773	961	125	0.956	512	100
A4	5701	874	186	0.916	512	50
A4	5797	934	117	1.000	512	100
A12	11975	1696	88	0.883	256	25
A12	11907	1728	111	0.803	256	25
A12	12014	1606	103	1.000	256	47
A12	11879	1785	63	0.944	256	50
A24	14406	1899	57	0.800	256	50
A24	14275	1949	112	0.868	256	25
B1	10666	220	329	1.000	512	100
B1	10292	113	446	0.886	512	100
B4	14174	384	413	0.927	512	100
B12	14600	341	444	0.889	512	67
B12	14161	828	643	0.656	512	100
B24	15423	716	552	0.824	256	50
B24	15219	330	424	0.994	512	100
C1	10251	348	375	1.000	512	100
C1	10776	747	238	0.981	512	50
C4	13856	612	415	0.937	512	50
C4	15155	1110	215	1.000	256	36
C12	15344	1166	225	0.996	256	50
C12	15236	1183	492	0.988	256	28
C12	14904	1087	461	0.863	256	25
C12	14696	781	419	0.930	512	50
C24	15191	1163	502	0.814	256	50
C24	14552	524	388	0.979	512	50
D1	2099	349	352	0.966	1024	50
D4	8266	1038	316	0.966	512	100
D4	8162	1128	332	0.971	512	54
D4	8283	1522	115	0.945	512	50
D12	16608	2309	213	0.941	512	53
D12	16518	2036	265	1.000	512	100
D12	16443	1964	309	1.000	512	50
E1	1105	158	115	1.000	512	50
E1	1144	159	125	0.805	420	41
E1	1105	160	133	0.806	512	50
E4	4165	668	91	0.888	256	25
E4	4098	599	144	0.830	256	50
E4	4164	712	63	0.920	256	50
E12	8296	1125	125	0.840	256	50

constraint limits their ability to extract all potential surplus solely by trading with each other. This problem is exacerbated in a thin market, where agents encounter fewer potential counterparties per unit time. Both factors increase the likelihood that agents trade inefficiently, as they lack sufficient



(a) $N = 66$



(b) $N = 25$

Figure 3: The effect of presence of a single MM on background-trader surplus and social welfare in equilibrium, across environments A, B, and C. Differences are presented as ranges, reflecting the multiplicity of equilibria found in some environments. The left point of each range is the minimum gain (in some cases a loss), that is, the lowest value observed with an equilibrium with MM minus the highest value observed in any equilibrium without MM. The right point is the maximum improvement observed: the difference between the highest value with a MM and the lowest without MM.

time and opportunity to reverse poor transactions. In such scenarios, the MM can boost not only overall welfare but also background-trader surplus by facilitating trade among impatient investors arriving at different times.

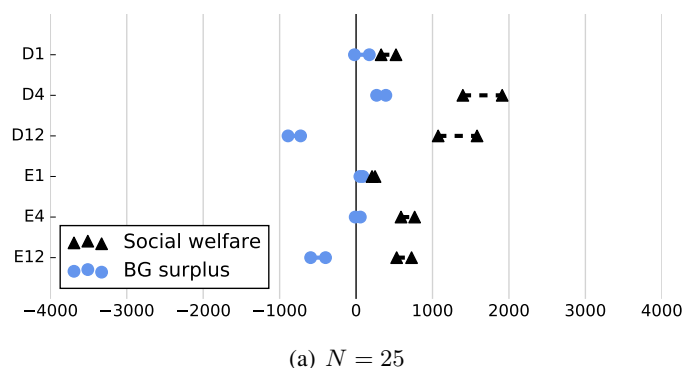


Figure 4: The effect of presence of a single MM on background-trader surplus and social welfare in equilibrium, for environments D and E. Data are presented as in Figure 3.

In our study, for markets populated by 25 background traders (Figures 3(b) and 4), the market maker improves welfare in all eighteen test environments. It improves background-trader surplus unambiguously in nine, and with a range partly on the positive side in seven more. Two more cases (D12 and E12) exhibit negative effects. Comparing environments A, D, and E, we do not observe any qualitative impact of changing the weighting between private and common value.

We observe that background traders are prone to shade less (i.e., midpoint R_{mid} of the ZI bid range is lower) when MM is present. The constant presence of a market maker shields the background traders from adverse selection by later-arriving background traders, thus alleviating the pressure on them to bid defensively. This effect is particularly strong for thin markets ($N = 25$), where the risk of adverse selection is greater. These are also the environments where MM tends to improve background-trader surplus. This indicates that the MM facilitates optimal allocations: with MM present, background investors can demand less surplus per trade, yet still achieve greater payoff than without the market maker.

We also find that MM spread ω tends to be larger for environments with shorter trading horizons, as we would expect when traders are more impatient. As evidence for this claim, in the single-MM equilibria with $N = 25$, mean MM spread ω weakly decreases with increasing duration T , in 41 of 43 cases. In single-MM equilibria with $N = 66$, 26 of 27 show weakly decreasing MM spread ω with increasing duration.

Finally, we evaluate liquidity for the maximum-welfare RSNE (Figure 5), with and without MM, by sampling results from profiles at the RSNE proportions. We measure liquidity via the *BID-ASK* spread (narrower spreads reflect greater liquidity) and background-trader execution time (interval between order submission and transaction). In general, both spreads and execution times drop with MM, which is indicative of the liquidity-provisioning capacity of the MM. In every environment except B4, $N = 66$, the median spread from local simulation is lower in each single-MM equilibrium than in any no-MM equilibrium. In the thinner markets, spreads without MM are significantly wider than in the thicker markets, as would be expected. The presence of the MM serves to narrow spreads considerably, nearly down to the levels present in the more populous environments. In every environment except B12 and C1 with $N = 66$, mean execution time is higher in equilibria without MM than in any equilibrium with one MM. The fact that the liquidity proxy

measures improve with MM in environments where background-trader surplus does not, however, underscores that these measures are not adequate substitutes for direct evaluation of investor welfare (Wellman & Wah, 2017).

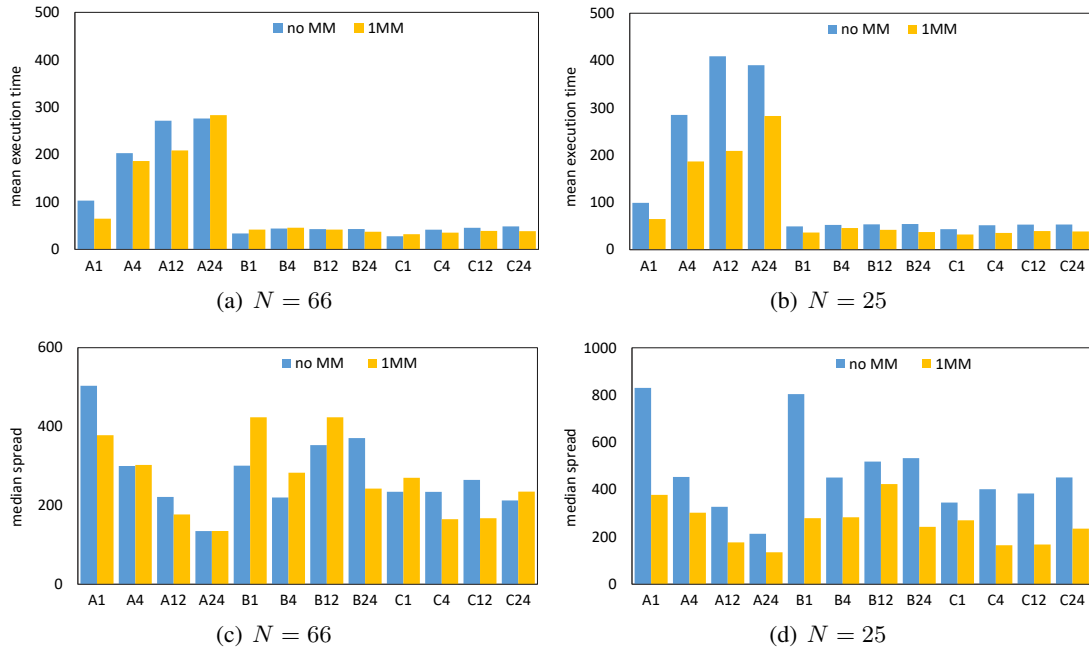


Figure 5: Comparison of background-trader execution time (Figures 5(a) and 5(b)) and median spread (Figures 5(c) and 5(d)) for the maximum-welfare RSNE in environments A–C, with and without MM. Mixed-strategy RSNE are approximated by profiles with trader population proportions corresponding to the strategy probabilities.




Figure 6: Comparison of background-trader execution time (Figure 6(a)) and median spread (Figure 6(b)) for the maximum-welfare RSNE in environments D and E, with and without MM. Mixed-strategy RSNE are approximated by profiles with trader population proportions corresponding to the strategy probabilities.

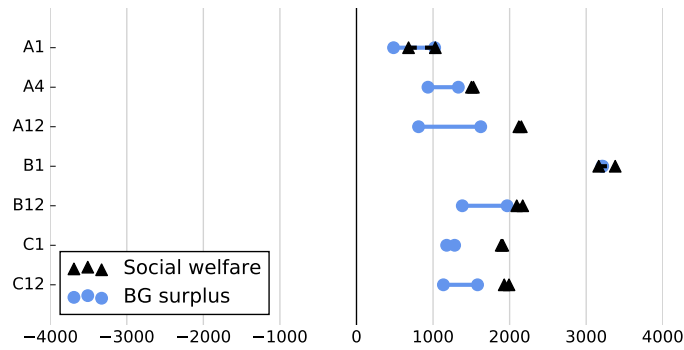
Table 7: Role-symmetric equilibria for games with two market makers, $N = 25$, calculated from the $(5, 2)$ -player DPR approximation. Profits shown are total for the two MMs. Data are presented as in Table 5.

Env	surplus	profit	R_{mid}	η	ω	ξ
A1	1835	216	63	0.800	256	25
A1	1860	172	96	0.800	256	48
A1	2062	10	73	0.800	128	50
A4	6549	568	52	0.800	128	50
A4	6946	196	47	0.851	64	50
A12	12480	1310	70	0.862	128	50
A12	13292	530	83	0.995	64	50
A12	13228	566	87	0.919	64	50
B1	11957	174	125	0.917	256	50
B1	11968	-52	179	0.886	248	24
B12	15638	712	532	0.774	256	25
B12	15953	410	387	1.000	256	50
B12	15971	406	403	1.000	256	50
B12	16224	204	293	0.968	256	25
C1	11639	630	110	0.967	256	25
C1	11660	612	114	0.968	256	25
C1	11556	732	120	0.809	256	25
C12	16198	410	330	1.000	128	50
C12	15752	822	387	0.716	128	50
C12	15806	740	488	0.809	128	50
C12	15803	764	447	0.800	128	50

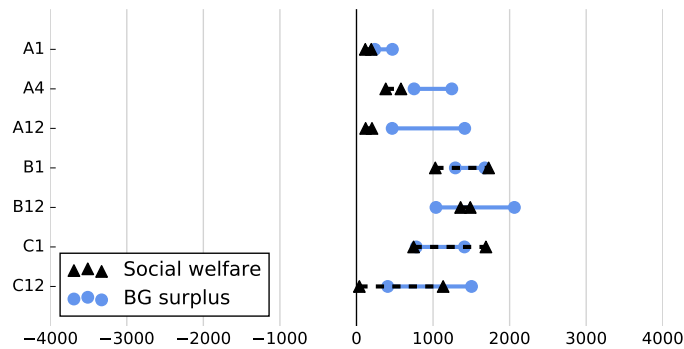
5.6 Competition among Market Makers

Our games with multiple MMs have either two or four MMs, and 25 background traders. MMs in these games choose from 6–7 available strategies from Table 4. Our search for equilibria in these games is as described above. The equilibria for two- and four-MM settings are presented in Tables 7 and 8, respectively. We use our usual sampling method to estimate the background-trader surplus and MM profit for each equilibrium profile.

Figure 7 presents the difference in social welfare and background-trader surplus between two-MM equilibria and zero- or one-MM equilibria, in seven environments. As for Figure 3, each row shows the range of differences in social welfare or background-trader surplus, comparing equilibria in settings with different numbers of MMs present. In each environment we tested, every equilibrium in the two-MM setting had greater social welfare and greater background-trader surplus than any equilibrium in the corresponding zero-MM or one-MM setting. Competition between oligopolistic MMs tends to benefit background traders and overall efficiency, relative to not having an MM or to having one monopolist MM.



(a) $N = 25$, two MM vs. zero MM



(b) $N = 25$, two MM vs. one MM

Figure 7: The effect of two MMs on background-trader surplus and social welfare in equilibrium, compared to corresponding environments without MM (Figure 7(a)), or with one MM (Figure 7(b)). Data are presented as in Figure 3.

We evaluated two environments with 25 background traders and four MMs, employing a $(5, 2)$ -player DPR approximation. As shown in Table 8, background-trader surplus in every equilibrium found is higher than in any equilibrium for corresponding settings with zero or one MM. Background-trader surplus also increases as we increase from two to four MMs for the A1 environment, but is ambiguous in A4. There is no clear trend in social welfare as the number of MMs increases from two to four, as the greater background-trader surplus with four MMs is offset by lower profits for each MM. Indeed, each four-MM equilibrium we found has lower MM profit, both in total and per MM, than any corresponding equilibrium with one or two MMs.⁶ Overall, as the number of MMs present increases beyond two, the effect of MM competition exhibits diminishing returns.

The spread setting ω of the MM equilibrium strategy is strictly lower in each equilibrium with multiple MMs than in any of the corresponding equilibria with one MM. With some competition, the MMs accept smaller expected gains from each trade by setting a narrower spread, in exchange for an increased likelihood of trading.

6. The negative profits exhibited for A1 environments reflect DPR approximation errors. The reduced-game payoffs for these equilibria were zero or slightly positive. Here, as for all our results, we derive equilibria using DPR but evaluate their surplus and welfare with respect to simulations of the solution profile as an original-game mixture.

Table 8: Role-symmetric equilibria for games with four market makers, $N = 25$, calculated from the $(5, 2)$ -player DPR approximation. Profits shown are total for the four MMs. Data are presented as in Table 5.

Env	surplus	profit	R_{mid}	η	ω	ξ
A1	2090	-4	52	0.821	128	50
A1	2097	-20	63	0.800	143	52
A1	2086	-36	103	0.886	145	52
A4	6664	452	63	0.853	128	50

In summary, environments with multiple MMs yield greater social welfare and background-trader surplus than environments with one or fewer MMs. Competition drives the MMs to provide cheaper liquidity in terms of reduced spreads, and as might be expected, makes them less profitable.

6. Conclusions

Market makers are generally considered to serve a valuable function in continuous market mechanisms by providing liquidity to bridge ebbs and flows of trader orders. The precise impact of this behavior, however, depends on specific features of market environments and trading strategies. We conducted a systematic agent-based simulation study to compare several parameterized environments with and without market-maker agents. We modeled a single security traded in a CDA populated by multiple background traders, and we characterized the strategic play in the induced empirical game model. This enabled us to compare outcomes *in equilibrium*, that is, allowing the background traders and market makers to strategically react to each other's presence.

Our analysis demonstrates the generally beneficial effects of market making on efficiency, and shows that whether these benefits accrue to background investors depends on market characteristics. Specifically, we find a tendency of a monopolist MM to improve the welfare of impatient investors (those in thin markets or with relatively few opportunities to trade with each other), but not in general. In markets with multiple MMs, we found larger and more consistent benefits to background traders. Competition among MMs leads them to apply narrower spreads, which provides better liquidity to the investors at some sacrifice of profit per MM.

Our study has several limitations, which must be taken into account in assessing our conclusions. First, our methods involve sampling, approximation, and limited search, all of which bear on the accuracy of equilibrium determinations. Sampling error is mitigated through the large number of simulation runs we gather over a breadth of environments and profiles, so it is not a fundamental concern for our conclusions here. The player reduction method we employ (DPR) has been shown to produce good approximate equilibrium estimates on other problems (Wiedenbeck & Wellman, 2012), and for our purposes approximate equilibria provide a sufficient basis for outcome comparison. However, DPR estimates are not guaranteed approximations, and it would be reassuring to confirm their quality in the context of our trading scenario. Even within the DPR game, we are unable to evaluate all profiles and cannot be sure that we have found all equilibria. Our search process attempts to evaluate all promising equilibrium candidates, but identifying these is not guaranteed.

A second area of limitation is the relatively narrow exploration of strategies. Our experience suggests that our equilibration process effectively tunes ZI strategies, and that these strategies exhibit reasonable behavior for this environment. Nevertheless, further investigation may yield improved versions of ZI or other strategies, for example adaptive variants (Cliff, 2009; Vytelingum, Cliff, & Jennings, 2008), that could alter equilibrium findings. Similar improvements may be found on the MM side, for instance with strategies incorporating learning (Abernethy & Kale, 2013).

Finally, our exploration of environments is also far from exhaustive. Whereas covering all plausible environments is infeasible, a broader range of variation on number of players, valuation distributions, and fundamental dynamics could go a long way in illuminating and validating robust conditions for qualitative welfare effects of market making in continuous double auctions.

Acknowledgments

This work significantly extends and supersedes a paper by the same title presented at the *14th International Conference on Autonomous Agents and Multiagent Systems* (Wah & Wellman, 2015). Elaine Wah completed her work on this project while a research assistant at the University of Michigan, prior to employment at IEX. Conclusions and views expressed herein are her own. Zhiyi Zhang's senior honors thesis served as a pilot study of MM strategies. We are also grateful to Uday Rajan for helpful advice on this work, and to anonymous reviewers for constructive feedback on our presentation. This work was supported in part by grant IIS-1421391 from the US National Science Foundation, and by the Microsoft Technology Policy Group.

Appendix A. Equilibria

A.1 Games without MM

Table 9: Symmetric equilibria for games without market makers, $N = 66$, calculated from the 6-player DPR approximation. The numeric column headings give R_{\max} values for the ZI strategies. All employ $R_{\min} = 0$ with the exception of the double star and double dagger (\ddagger) values which use $R_{\min} = \frac{1}{2}R_{\max}$. All employ $\eta = 1$, except for the starred values which use $\eta = 0.8$, the dagger (\dagger) value which uses $\eta = 0.6$, and the double dagger value which uses $\eta = 0.4$. Each row of the table describes the mixture probabilities for strategies for one equilibrium, and corresponds to the matching row in Table 2. The column for $R_{\max} = 2500$ is not listed in the table as this strategy is not played in any of the equilibria found.

Env	65*	125*	125	250*	250	500	500**	1000*	1000	1000 \ddagger	1500 \dagger	2000 \ddagger
A1	0	0	0	0	0	0	0	0	0	1	0	0
A1	0	0	0	0	0	0.106	0.861	0	0	0.033	0	0
A4	0	0.113	0	0	0	0	0.887	0	0	0	0	0
A4	0	0	0	0	0.029	0.115	0.856	0	0	0	0	0
A12	0	0	0	0.223	0.209	0	0.568	0	0	0	0	0
A24	0	0	0	1	0	0	0	0	0	0	0	0
A24	0	0.296	0	0	0.704	0	0	0	0	0	0	0
A24	0	0	0	0.138	0.030	0.832	0	0	0	0	0	0
B1	0	0	0	0	0	0	0.824	0	0	0.176	0	0
B4	0	0	0	0	0	0	0.714	0.193	0.092	0	0	0
B12	0	0	0	0	0	0	0.683	0	0	0.317	0	0
B24	0	0	0	0	0	0	0.676	0	0	0.324	0	0
B24	0	0	0	0	0	0	0	0	0.293	0.707	0	0
C1	0	0	0	0	0	1	0	0	0	0	0	0
C1	0	0	0	0	0	0	1	0	0	0	0	0
C4	0	0	0	0	0	0	0	1	0	0	0	0
C4	0	0	0	0	0	0	0.960	0	0	0.040	0	0
C12	0	0	0	0	0	0.011	0.796	0	0	0	0.193	0
C24	0	0	0	0	0	0	0	0.995	0	0	0.005	0
C24	0	0	0	0	0	0	0.714	0.267	0	0	0.019	0

Table 10: Symmetric equilibria for games without market makers, $N = 25$, calculated from the 5-player DPR approximation. Data are presented as in Table 9. Each row corresponds to the matching row in Table 3. The column for $R_{\max} = 2500$ is not listed in the table as this strategy is not played in any of the equilibria found.

Env	65*	125*	125	250*	250	500	500**	1000*	1000	1000 [‡]	1500 [‡]	2000 [‡]
A1	0	0	0	0	0	0	0	0	0	0.978	0.022	0
A1	0	0	0	0	0	0.035	0.965	0	0	0	0	0
A4	0.072	0	0	0	0	0	0.928	0	0	0	0	0
A12	0	0	0	0	0	0.318	0.682	0	0	0	0	0
A24	0	0	0	0	0	0	1	0	0	0	0	0
A24	0	0	0	0.256	0	0.744	0	0	0	0	0	0
A24	0	0	0.132	0.868	0	0	0	0	0	0	0	0
B1	0	0	0	0	0	0	0	0	0	1	0	0
B4	0	0	0	0	0	0	0.621	0	0	0.379	0	0
B12	0	0	0	0	0	0	0.525	0	0	0.475	0	0
B24	0	0	0	0	0	0	0.490	0.039	0	0.470	0	0
C1	0	0	0	0	0	0	1	0	0	0	0	0
C4	0	0	0	0	0	0	0.730	0	0	0.270	0	0
C12	0	0	0	0	0	0	0.823	0	0	0.177	0	0
C24	0	0	0	0	0	0	0.694	0	0	0.306	0	0
D1	0	0.029	0	0	0	0	0.971	0	0	0	0	0
D1	0	0	0	0	0	0	0.160	0	0	0.840	0	0
D4	0	0	0	0	0	0	0.090	0	0	0.910	0	0
D12	0	0	0	0	0	0	0.954	0	0	0.046	0	0
E1	0	0	0	0.581	0.419	0	0	0	0	0	0	0
E4	0	0	0	0	0.688	0.200	0.112	0	0	0	0	0
E12	0	0	0	0	0.717	0.282	0	0	0	0	0	0
E12	0	0	0	0.572	0.428	0	0	0	0	0	0	0

A.2 Games with a Single MM

Table 11: Role-symmetric equilibria for games with a market maker, $N = 66$, based on (6, 1)-player DPR approximation. Column headings give R_{\max} values for the ZI strategies as in Table 9. Each row of the table corresponds to the matching row in Table 5 and describes one equilibrium found. The column for R_{\max} value of 2000^\ddagger is not listed in the table as this strategy is not played in any of the equilibria found.

Env	Background-trader R_{\max}											
	65*	125*	125	250*	250	500	500**	1000*	1000	1000 [‡]	1500 [†]	2500
A1	0	0	0	0.337	0.076	0.273	0.314	0	0	0	0	0
A1	0	0	0.116	0	0.185	0.699	0	0	0	0	0	0
A1	0.061	0	0	0.225	0	0	0.714	0	0	0	0	0
A4	0.034	0.072	0	0.724	0	0.171	0	0	0	0	0	0
A4	0	0	0	0	0.917	0.083	0	0	0	0	0	0
A4	0.058	0	0	0.893	0	0	0.049	0	0	0	0	0
A12	0	0	0	0.924	0	0.076	0	0	0	0	0	0
A24	0	0.122	0.115	0	0.763	0	0	0	0	0	0	0
A24	0	0.327	0	0.673	0	0	0	0	0	0	0	0
A24	0.070	0	0.012	0.918	0	0	0	0	0	0	0	0
B1	0	0	0	0	0	0	0.845	0	0	0.155	0	0
B4	0	0	0	0	0	0	0.843	0	0.012	0	0.145	0
B12	0	0	0	0	0	0	0.934	0	0	0	0.066	0
B12	0	0	0	0	0	0.032	0	0.968	0	0	0	0
B24	0	0	0	0	0	0	0	0.779	0.022	0.199	0	0
B24	0	0	0	0	0	0	0	0	0.642	0	0.357	0
B24	0	0	0	0	0	0	0.616	0	0.157	0.227	0	0
B24	0	0	0	0	0	0	0	0	0	0	0.887	0.113
C1	0	0	0	0	0	0	1	0	0	0	0	0
C4	0	0	0	0	0	0	0.634	0.230	0.136	0	0	0
C4	0	0	0	0	0	1	0	0	0	0	0	0
C12	0	0	0	0	0	0	0	0.678	0.322	0	0	0
C12	0	0	0	0	0	0	0.554	0.446	0	0	0	0
C12	0	0	0	0.119	0	0	0	0	0.881	0	0	0
C24	0	0	0	0	0	0	0.093	0.377	0.529	0	0	0
C24	0	0	0	0	0	0.327	0	0.564	0.109	0	0	0
C24	0	0	0	0	0	0	0.483	0.517	0	0	0	0

Table 12: Role-symmetric equilibria for games with a market maker, $N = 66$, based on $(6, 1)$ -player DPR approximation. Column headings give ω values for MM strategies. All MM strategies use $K = 100$ and $\xi = 50$ except those with subscripts indicating the ξ used. Each row of the table corresponds to the matching row in Table 5 and describes one equilibrium found. The columns for MM ω values of 64 and 128 are not listed in the table as these strategies are not played in any of the equilibria found.

Env	Market-maker ω				
	256 ₂₅	256	512	1024	512 ₁₀₀
A1	0	0	0.787	0	0.213
A1	0	0	0	0	1
A1	0	0	1	0	0
A4	0	1	0	0	0
A4	1	0	0	0	0
A4	1	0	0	0	0
A12	0.091	0.908	0	0	0
A24	0	1	0	0	0
A24	1	0	0	0	0
A24	1	0	0	0	0
B1	0	0	0	0.819	0.181
B4	0	0	0	0	1
B12	0	0	0	0	1
B12	0	0.082	0.918	0	0
B24	0	1	0	0	0
B24	0	1	0	0	0
B24	0	0	0	0	1
B24	1	0	0	0	0
C1	0	0	1	0	0
C4	0	1	0	0	0
C4	1	0	0	0	0
C12	0	1	0	0	0
C12	0	1	0	0	0
C12	0	1	0	0	0
C24	0	1	0	0	0
C24	1	0	0	0	0
C24	1	0	0	0	0

Table 13: Role-symmetric equilibria for games with a market maker, $N = 25$, calculated from the $(5, 1)$ -player DPR approximation. Each row of the table corresponds to the matching row in Table 6 and describes one equilibrium found. Data are presented as in Table 11, but with columns for R_{\max} values of 2000[‡] and 2500 excluded from the table as these strategies are not played in any of the equilibria found.

Env	Background-trader R_{\max}										
	65*	125*	125	250*	250	500	500**	1000*	1000	1000 [‡]	1500 [‡]
A1	0	0	0	0.687	0	0.313	0	0	0	0	0
A1	0	0	0.156	0.300	0.544	0	0	0	0	0	0
A4	0	0	0	0.519	0	0.462	0.019	0	0	0	0
A4	0	0	0	0.221	0.779	0	0	0	0	0	0
A4	0	0	0	0.419	0.127	0.420	0.034	0	0	0	0
A4	0	0	0.130	0	0.870	0	0	0	0	0	0
A12	0	0.587	0	0	0.413	0	0	0	0	0	0
A12	0.149	0	0	0.835	0.017	0	0	0	0	0	0
A12	0	0	0.349	0	0.651	0	0	0	0	0	0
A12	0.280	0	0.673	0	0	0.047	0	0	0	0	0
A24	0.170	0.830	0	0	0	0	0	0	0	0	0
A24	0.145	0	0	0.515	0.340	0	0	0	0	0	0
B1	0	0	0	0	0	0.367	0.633	0	0	0	0
B1	0	0	0	0	0	0	0.810	0	0	0.190	0
B4	0	0	0	0	0	0.058	0.821	0	0	0.121	0
B12	0	0	0	0	0	0	0.754	0.092	0	0.153	0
B12	0	0	0	0	0	0	0	0	0.427	0.573	0
B24	0	0	0	0	0	0	0	0.256	0.537	0.207	0
B24	0	0	0	0	0	0	0.610	0.031	0.359	0	0
C1	0	0	0	0	0	0	1	0	0	0	0
C1	0	0	0	0.096	0	0.904	0	0	0	0	0
C4	0	0	0	0	0	0	0.894	0	0	0.106	0
C4	0	0	0	0	0.281	0.719	0	0	0	0	0
C12	0	0	0	0	0.239	0.742	0	0.019	0	0	0
C12	0	0	0.034	0	0	0	0	0	0.936	0	0.030
C12	0	0	0	0	0	0	0.315	0.685	0	0	0
C12	0	0	0	0	0	0	0.649	0.351	0	0	0
C24	0	0	0	0	0	0	0.155	0.761	0	0	0.084
C24	0	0	0	0	0	0	0.897	0.103	0	0	0
D1	0	0	0	0.117	0	0	0.831	0.052	0	0	0
D4	0.172	0	0	0	0	0	0.829	0	0	0	0
D4	0	0	0.020	0.146	0	0	0.834	0	0	0	0
D4	0	0.165	0	0.110	0.726	0	0	0	0	0	0
D12	0	0	0	0.297	0	0.703	0	0	0	0	0
D12	0	0	0.196	0	0.193	0	0.610	0	0	0	0
D12	0	0	0	0	0.231	0.064	0.705	0	0	0	0
E1	0	0	0.153	0	0.847	0	0	0	0	0	0
E1	0	0	0	0.975	0.025	0	0	0	0	0	0
E1	0	0	0	0.968	0	0	0.032	0	0	0	0
E4	0	0.562	0	0	0.430	0.008	0	0	0	0	0
E4	0	0	0	0.849	0	0.151	0	0	0	0	0
E4	0	0.400	0.600	0	0	0	0	0	0	0	0
E12	0	0	0	0.801	0.199	0	0	0	0	0	0

Table 14: Role-symmetric equilibria for games with a market maker, $N = 25$, calculated from the (5, 1)-player DPR approximation. Each row of the table corresponds to the matching row in Table 6 and describes one equilibrium found. Data are presented as in Table 12, but with MM ω values of 64 and 128 excluded from the table as these strategies are not played in any of the equilibria found.

Env	Market-maker ω				
	256 ₂₅	256	512	1024	512 ₁₀₀
A1	0	0	1	0	0
A1	0	0	0	0	1
A4	0	0	0	0	1
A4	0	0	0	0	1
A4	0	0	1	0	0
A4	0	0	0	0	1
A12	1	0	0	0	0
A12	1	0	0	0	0
A12	0.139	0.861	0	0	0
A12	0	1	0	0	0
A24	0	1	0	0	0
A24	1	0	0	0	0
B1	0	0	0	0	1
B1	0	0	0	0	1
B4	0	0	0	0	1
B12	0	0	0.669	0	0.331
B12	0	0	0	0	1
B24	0	1	0	0	0
B24	0	0	0	0	1
C1	0	0	0	0	1
C1	0	0	1	0	0
C4	0	0	1	0	0
C4	0.540	0.460	0	0	0
C12	0	1	0	0	0
C12	0.879	0.121	0	0	0
C12	1	0	0	0	0
C12	0	0	1	0	0
C24	0	1	0	0	0
C24	0	0	1	0	0
D1	0	0	0	1	0
D4	0	0	0	0	1
D4	0	0	0.913	0	0.086
D4	0	0	1	0	0
D12	0	0	0.942	0	0.058
D12	0	0	0	0	1
D12	0	0	1	0	0
E1	0	0	1	0	0
E1	0.359	0	0.641	0	0
E1	0	0	1	0	0
E4	1	0	0	0	0
E4	0	1	0	0	0
E4	0	1	0	0	0
E12	0	1	0	0	0

A.3 Games with Multiple MM

Table 15: Role-symmetric equilibria for games with two market makers, $N = 25$, from the $(5, 2)$ -player DPR approximation. Data are presented as in Table 11. The columns for some background-trader strategies are not listed, because these strategies are not played in any of the equilibria found. Each row corresponds to the matching row in Table 7.

Env	Background-trader R_{\max}									
	65*	125*	125	250*	250	500	1000*	1000	1000 [†]	1500 [†]
A1	0	1	0	0	0	0	0	0	0	1
A1	0.318	0	0	0.682	0	0	0	0	0	0.072
A1	0	0.828	0	0.172	0	0	0	0	0	0
A4	0.790	0	0	0.210	0	0	0	0	0	0
A4	0.502	0.241	0.257	0	0	0	0	0	0	0
A12	0.414	0.274	0	0	0.312	0	0	0	0	0
A12	0	0.026	0.651	0	0.323	0	0	0	0	0
A12	0	0.406	0.207	0	0.388	0	0	0	0	0
B1	0	0	0	0.415	0.585	0	0	0	0	0
B1	0	0	0	0.570	0	0.430	0	0	0	0.969
B12	0	0	0	0	0	0	0.872	0	0.129	1
B12	0	0	0	0	0.302	0	0	0.698	0	0
B12	0	0	0	0	0.115	0.215	0	0.670	0	0
B12	0	0	0	0.161	0	0.586	0	0.253	0	1
C1	0.163	0	0	0	0.837	0	0	0	0	1
C1	0	0	0.173	0.160	0.668	0	0	0	0	1
C1	0	0.076	0	0.880	0.044	0	0	0	0	1
C12	0	0	0	0	0	0.680	0	0.320	0	0
C12	0	0	0	0.581	0	0	0	0	0.419	0
C12	0	0	0	0	0	0.047	0.953	0	0	0
C12	0	0.106	0	0.018	0	0	0.876	0	0	0

Table 16: Role-symmetric equilibria for games with two market makers, $N = 25$, from the $(5, 2)$ -player DPR approximation. Data are presented as in Table 12. The columns for some MM strategies are not listed, because these strategies are not played in any of the equilibria found. Each row corresponds to the matching row in Table 7.

Env	Market-maker ω			
	256 ₂₅	64	128	256
A1	0	0	0	0
A1	0	0	0.928	0
A1	0	1	0	0
A4	0	1	0	0
A4	1	0	0	0
A12	0	1	0	0
A12	1	0	0	0
A12	1	0	0	0
B1	0	0	1	0
B1	0	0	0	0.031
B12	0	0	0	0
B12	0	0	1	0
B12	0	0	1	0
B12	0	0	0	0
C1	0	0	0	0
C1	0	0	0	0
C1	0	0	0	0
C12	0	1	0	0
C12	0	1	0	0
C12	0	1	0	0
C12	0	1	0	0

Table 17: Role-symmetric equilibria for games with four market makers, $N = 25$, calculated from the $(5, 2)$ -player DPR approximation. Data are presented as in Tables 11 and 12. Some strategies are not listed in the table as these strategies are not played in any of the equilibria found. Each row corresponds to the matching row in Table 8.

Env	Background-trader R_{\max}					Market-maker ω		
	65*	125*	125	250*	250	128	512 ₁₀₀	NOOP
A1	0.341	0.555	0.104	0	0	1	0	0
A1	0	1	0	0	0	0.957	0.039	0.004
A1	0	0	0.346	0.571	0.084	0.952	0.045	0.003
A4	0	0.734	0.266	0	0	1	0	0

A.4 Surplus Comparison, with and without MM

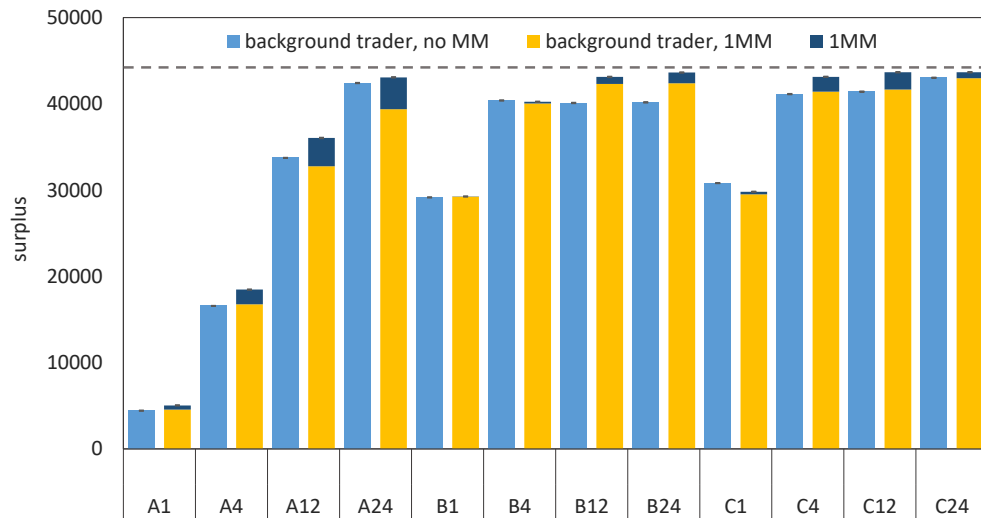


Figure 8: Comparison of background-trader surplus (with and without market makers) and MM profit for $N = 66$. The dotted line is the optimal social welfare available (44155). Error bars indicate the 95% confidence interval for total welfare in the maximum-welfare role-symmetric Nash equilibrium in each environment, with and without MM.




Figure 9: Comparison of background-trader surplus (with and without market makers) and MM profit for $N = 25$. The dotted line is the optimal social welfare available (16306). Error bars indicate the 95% confidence interval for total welfare in the maximum-welfare role-symmetric Nash equilibrium in each environment, with and without MM.

References

- Abernethy, J., Chen, Y., & Wortman Vaughan, J. (2011). An optimization-based framework for automated market-making. In *12th ACM Conference on Electronic Commerce*, pp. 297–306.
- Abernethy, J., & Kale, S. (2013). Adaptive market making via online learning. In *Advances in Neural Information Processing Systems*, pp. 2058–2066.
- Amihud, Y., & Mendelson, H. (1980). Dealership market: Market-making with inventory. *Journal of Financial Economics*, 8(1), 31–53.
- Bernhardt, D., & Hughson, E. (1997). Splitting orders. *Review of Financial Studies*, 10(1), 69–101.
- Bessembinder, H., Hao, J., & Lemmon, M. L. (2011). Why designate market makers? Affirmative obligations and market quality. Tech. rep. 989061, SSRN.
- Bessembinder, H., Hao, J., & Zheng, K. (2015). Market making contracts, firm value, and the IPO decision. *Journal of Finance*, 70(5), 1997–2028.
- Biais, B., Glosten, L., & Spatt, C. (2005). Market microstructure: A survey of microfoundations, empirical results, and policy implications. *Journal of Financial Markets*, 8(2), 217–264.
- Biais, B., Martimort, D., & Rochet, J.-C. (2000). Competing mechanisms in a common value environment. *Econometrica*, 68(4), 799–837.
- Brusco, S., & Jackson, M. O. (1999). The optimal design of a market. *Journal of Economic Theory*, 88(1), 1–39.
- Cassell, B.-A., & Wellman, M. P. (2013). EGTAOnline: An experiment manager for simulation-based game studies. In *Multi-Agent-Based Simulation XIII*, Vol. 7838 of *Lecture Notes in Artificial Intelligence*. Springer.
- Chakraborty, T., & Kearns, M. (2011). Market making and mean reversion. In *11th ACM Conference on Electronic Commerce*, pp. 307–314.
- Chan, N. T., & Shelton, C. (2001). An electronic market-maker. Tech. rep. AI Memo 2001-005, Massachusetts Institute of Technology.
- Chen, Y., & Pennock, D. M. (2007). A utility framework for bounded-loss market makers. In *23rd Conference on Uncertainty in Artificial Intelligence*, pp. 49–56.
- Chowdhry, B., & Nanda, V. (1991). Multimarket trading and market liquidity. *Review of Financial Studies*, 4(3), 483–511.
- Cliff, D. (2009). ZIP60: Further explorations in the evolutionary design of online auction market mechanisms. *IEEE Transactions on Evolutionary Computation*, 13(1), 3–18.
- Conrad, J., Wahal, S., & Xiang, J. (2015). High-frequency quoting, trading, and the efficiency of prices. *Journal of Financial Economics*, 116(2), 271–291.
- Darley, V., Outkin, A., Plate, T., & Gao, F. (2000). Sixteenths or pennies? Observations from a simulation of the NASDAQ stock market. In *IEEE/IAFE/INFORMS Conference on Computational Intelligence for Financial Engineering*, pp. 151–154.
- Das, R., Hanson, J. E., Kephart, J. O., & Tesauro, G. (2001). Agent-human interactions in the continuous double auction. In *17th International Joint Conference on Artificial Intelligence*, pp. 1169–1176.

- Das, S. (2005). A learning market-maker in the Glosten-Milgrom model. *Quantitative Finance*, 5(2), 169–180.
- Das, S. (2008). The effects of market-making on price dynamics. In *7th International Joint Conference on Autonomous Agents and Multiagent Systems*, pp. 887–894.
- Das, S., & Magdon-Ismail, M. (2008). Adapting to a market shock: Optimal sequential market-making. In *Advances in Neural Information Processing Systems*, pp. 361–368.
- Dennert, J. (1993). Price competition between market makers. *Review of Economic Studies*, 60(3), 735–751.
- Eldor, R., Hauser, S., Pilo, B., & Shurki, I. (2006). The contribution of market makers to liquidity and efficiency of options trading in electronic markets. *Journal of Banking and Finance*, 30(7), 2025–2040.
- Feng, Y., Yu, R., & Stone, P. (2004). Two stock-trading agents: Market making and technical analysis. In *Agent-Mediated Electronic Commerce V. Designing Mechanisms and Systems*, pp. 18–36. Springer.
- Frey, S., & Grammig, J. (2006). Liquidity supply and adverse selection in a pure limit order book market. *Empirical Economics*, 30(4), 1007–1033.
- Friedman, D. (1993). The double auction market institution: A survey. In Friedman, D., & Rust, J. (Eds.), *The Double Auction Market: Institutions, Theories, and Evidence*, pp. 3–25. Addison-Wesley.
- Garman, M. B. (1976). Market microstructure. *Journal of Financial Economics*, 3(3), 257–275.
- Glosten, L. R. (1994). Is the electronic open limit order book inevitable?. *Journal of Finance*, 49(4), 1127–1161.
- Glosten, L. R., & Milgrom, P. R. (1985). Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics*, 14(1), 71–100.
- Gode, D. K., & Sunder, S. (1993). Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy*, 101(1), 119–137.
- Goettler, R. L., Parlour, C. A., & Rajan, U. (2009). Informed traders and limit order markets. *Journal of Financial Economics*, 93(1), 67–87.
- Grossman, S. J., & Miller, M. H. (1988). Liquidity and market structure. *Journal of Finance*, 43(3), 617–633.
- Hanson, R. (2007). Logarithmic market scoring rules for modular combinatorial information aggregation. *Journal of Prediction Markets*, 1(1), 3–15.
- Hasbrouck, J., & Sofianos, G. (1993). The trades of market makers: An empirical analysis of NYSE specialists. *Journal of Finance*, 48(5), 1565–1593.
- Huang, J., & Wang, J. (2010). Market liquidity, asset prices, and welfare. *Journal of Financial Economics*, 95(1), 107–127.
- Huang, R. D., & Stoll, H. R. (1996). Dealer versus auction markets: A paired comparison of execution costs on NASDAQ and the NYSE. *Journal of Financial Economics*, 41(3), 313–357.

- Jumadinova, J., & Dasgupta, P. (2010). A comparison of different automated market-maker strategies. In *12th Workshop on Agent-Mediated Electronic Commerce*, pp. 141–154.
- Kyle, A. S. (1985). Continuous auctions and insider trading. *Econometrica*, 53(6), 1315–1335.
- Leach, J. C., & Madhavan, A. N. (1992). Intertemporal price discovery by market makers: Active versus passive learning. *Journal of Financial Intermediation*, 2(2), 207–235.
- LeBaron, B. (2006). Agent-based computational finance. In Tesfatsion, L., & Judd, K. L. (Eds.), *Handbook of Agent-Based Computational Economics*, pp. 1187–1233. Elsevier.
- Madhavan, A. (2000). Market microstructure: A survey. *Journal of Financial Markets*, 3(3), 205–258.
- Manaster, S., & Mann, S. C. (1996). Life in the pits: Competitive market making and inventory control. *Review of Financial Studies*, 9(3), 953–975.
- Menkveld, A. J. (2013). High frequency trading and the new market makers. *Journal of Financial Markets*, 16(4), 712–740.
- Menkveld, A. J., & Wang, T. (2013). How do designated market makers create value for small-caps?. *Journal of Financial Markets*, 16(3), 571–603.
- O’Hara, M. (1995). *Market Microstructure Theory*, Vol. 108. Blackwell Cambridge.
- O’Hara, M., & Oldfield, G. (1986). The microeconomics of market making. *Journal of Financial and Quantitative Analysis*, 21(4), 361–376.
- Saar, G. (2010). Specialist markets. In *Encyclopedia of Quantitative Finance*. Wiley Online Library.
- Sandås, P. (2001). Adverse selection and competitive market making: Empirical evidence from a limit order market. *Review of Financial Studies*, 14(3), 705–734.
- Schwartz, R. A., & Peng, L. (2013). Market makers. In *Encyclopedia of Finance*, pp. 487–489. Springer.
- Seppi, D. J. (1997). Liquidity provision with limit orders and a strategic specialist. *Review of Financial Studies*, 10(1), 103–150.
- Strategic Reasoning Group (2016). MarketSim. <https://github.com/egtaonline/market-sim/tree/marketsim1>.
- Vytelingum, P., Cliff, D., & Jennings, N. R. (2008). Strategic bidding in continuous double auctions. *Artificial Intelligence*, 172(14), 1700–1729.
- Wah, E., & Wellman, M. P. (2013). Latency arbitrage, market fragmentation, and efficiency: A two-market model. In *14th ACM Conference on Electronic Commerce*, pp. 855–872.
- Wah, E., & Wellman, M. P. (2015). Welfare effects of market making in continuous double auctions. In *14th International Conference on Autonomous Agents and Multiagent Systems*, pp. 57–66.
- Wellman, M. P. (2011). *Trading Agents*. Morgan & Claypool.
- Wellman, M. P. (2016). Putting the agent in agent-based modeling. *Autonomous Agents and Multi-Agent Systems*, 30(6), 1175–1189.
- Wellman, M. P., & Wah, E. (2017). Strategic agent-based modeling of financial markets. *Russell Sage Foundation Journal of the Social Sciences*, 3(1), 104–119.

- Wiedenbeck, B., & Wellman, M. P. (2012). Scaling simulation-based game analysis through deviation-preserving reduction. In *11th International Conference on Autonomous Agents and Multiagent Systems*, pp. 931–938.
- Zhan, W., & Friedman, D. (2007). Markups in double auction markets. *Journal of Economic Dynamics and Control*, 31(9), 2984–3005.