

Welfare Effects of Market Making in Continuous Double Auctions

Elaine Wah
ewah@umich.edu

Michael P. Wellman
wellman@umich.edu

Computer Science & Engineering
University of Michigan

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 not only total surplus across a variety of environments, but also background-trader surplus in thin markets with impatient investors, with urgency captured by a limited trading horizon. Comparison across environments reveals factors that influence the existence and magnitude of benefits provided by the market maker function.

Categories and Subject Descriptors

J.4 [Social and Behavioral Sciences]: Economics

Keywords

Market maker; agent-based simulation; allocative efficiency

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 [43].

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 [31]. 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 [41]. With the transition to electronic markets, pure limit-order markets are becoming predominant [20, 23], thus 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 [24, 33]; other prior research has investigated the effects of MM on liquidity (e.g., as measured by price spreads) [18] and price discovery [34]. 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 affirmative MM obligation due to adverse selection [4, 5], the cost structure of market participation in supplying liquidity [30], and trading mechanisms to incentivize market making [7].

In this study, 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 CDA mechanism in a market environment comprising a single market maker and multiple background traders. 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 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 single MM has no private value, and thus aims to profit by 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 [25]. For the MM, we consider heuristic strategies loosely based on that defined by Chakraborty and Kearns [9]. 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 and evidence for a general benefit for background investors in thin markets with limited trading horizons. We derive empirical equilibria with and without mar-

Appears in: *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2015)*, Bordini, Elkind, Weiss, Yolum (eds.), May 4–8, 2015, Istanbul, Turkey.

Copyright © 2015, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

ket making in a number of environments, finding that the mix of background-trader strategies in equilibrium varies depending on the presence and strategy choice of the MM. In all of our environments, market making is profitable in equilibrium, and in all but two equilibrium comparisons, the presence of MM increases overall welfare (background trader surplus combined with MM profit). In markets with impatient investors, we find typically that the background traders themselves are better off (in expectation), that is, welfare is improved net of MM profits. However, if the background traders have ample time horizons, then MM profits often exceed any improvement in overall efficiency, so the background investors are net worse off with the MM 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 strategies, and Section 5 describes our MM strategies. In Section 6 we present our empirical game-theoretic analysis, and we conclude in Section 7.

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 .




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

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,

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

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 profitable trade with this buyer, but at a price quite 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.

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. In other words, liquidity has economic value only if 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. 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 [6, 36, 39]. Early work focused on dealer markets, in which a monopolistic MM (the dealer) controls trading by acting as the middleman. Garman [22] presents an explicit formulation of the market maker’s optimization problem. O’Hara and Oldfield [40] and Amihud and Mendelson [3] concentrate on the impact of dealer inventory on spreads, while the seminal model of Glosten and Milgrom [24] frames spread as arising from adverse selection. Others focus on the consequences of informed trading on MM [33, 12, 17], as well as the role of market makers as liquidity providers [27, 44].

Much of the relevant theoretical literature, however, relies on simplifying assumptions of MM behavior and trader interactions

[6]. Empirical studies have provided insight on the effects of market makers in real-world markets [38, 20, 42, 37, 29]. 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 [35]; however, only a handful of ABM finance papers focus on market making [10, 14, 17].

Outside of microstructure, researchers have developed MM strategies for a variety of settings, including prediction markets [28, 11, 2], dealer-mediated markets [16, 32], CDAs [19], and environments where prices are generated exogenously [1]. In this last category, Chakraborty and Kearns [9] demonstrate the profitability of market making given a mean-reverting price series 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, address questions about allocative efficiency in the market. To our knowledge, the literature on welfare effects of MM behavior is quite limited, and existing studies are largely concerned with how adverse selection affects allocative efficiency. For example, Bessembinder et al. [4] 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 [7] 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 [30] 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. [5] 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.

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 . 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 market maker. At any given time, the background investors are restricted to 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.

4.1 Valuation Model

Each background trader has an individual valuation for the security comprised 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 parameter $u_t \sim \mathcal{N}(0, \sigma_s^2)$ 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 et al. [26]. The vector is of size $2q_{\max}$, where $q_{\max} > 0$ is the maximum number of units the agent can be long or short at any time, with

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

Element θ_i^q is the incremental private benefit obtained from 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 $\hat{\theta}$ 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:

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

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 $\theta_B^{q_B(B)+1} - \theta_S^{q(S)}$, where $q(i)$ denotes the pre-trade position of agent i .

4.2 Background Trading Strategies

There is an extensive literature on autonomous bidding strategies for CDAs [21, 48, 15]. In this study, we consider trading strategies in the so-called *Zero Intelligence* (ZI) family [25].

The background traders arrive at the market according to a Poisson process with rate λ_a . On arrival, they are assigned to buy or sell (with equal probability), and accordingly submit an order to buy or sell a single unit. Traders subsequently reenter the market, with time between entries distributed exponentially at rate λ_r . Agents may trade any number of times, as long as their net positions do not exceed q_{\max} (either long or short). At the end of the simulation period, background traders liquidate their accumulated inventory at r_T , the end-time fundamental. Background traders are notified of all transactions and current price quotes with zero delay, and may use this information in computing their bids.

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:²

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

²We also implemented a version of this strategy that did not account for mean reversion, taking $\hat{r}_t = r_t$. This misestimating agent performed poorly with $\kappa > 0$, so we omit it from further discussion.

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 also consider a threshold parameter $\eta \in [0, 1]$, whereby if the agent could achieve a fraction η of its requested surplus at the current price quote, it would simply take that quote rather than posting a limit order to the book. Setting $\eta = 1$ is equivalent to the strategy without employing the threshold.

In our model, background traders are permitted to reenter the market, thus we denote agents employing this strategy *ZI with Reentry* (ZIR). Upon each reentry, the ZIR trader withdraws its previous order (if not transacted yet) before executing the strategy described above.

5. MARKET MAKER STRATEGIES

Much of the prior work on MM strategies treats the market maker as a dealer [24, 16]. In our model, 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 single-quantity buy and sell orders, comprised of K rungs spaced ξ ticks apart:

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

with $S_t > B_t$ and $K, \xi > 0$. The MM arrives at time 0 and reenters the market according to a 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 liquidates its inventory at the end of the trading horizon. The liquidation price is the global fundamental value r_T . The MM's total profit is defined by the sum of trading cash flow plus liquidation proceeds.

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.

Within this family of MM strategies, we define the *Fundamental Market Maker* (FUNDMM), which uses its observation of the current fundamental r_t to inform its ladder construction. FUNDMM computes an estimate \hat{r}_t of the terminal fundamental r_T via (1), and centers its ladder around this estimate. It determines its spread ω either via the current quote (i.e., $\omega_t = ASK_t - BID_t$), or by a fixed value specified as a strategy parameter. The central ladder prices are:

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

6. EMPIRICAL GAME-THEORETIC ANALYSIS

We have described a variety of 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, that is, 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. Second, we determine equilibria among these strategies through a simulation-based process, known as *empirical game-theoretic analysis* (EGTA) [47]. 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 [46]. We generate data for various combinations of the strategies introduced in Sections 4.2 and 5, each sampled over many runs (all environments in this study averaged at least 241 per profile, most over 5000) to account for stochastic effects (valuation schedules, trajectories of the market fundamental, agent arrival patterns). From this 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.

6.1 Environment Settings

We evaluate the performance of background traders and the MM within thirty parametrically distinct environments. For each environment, we analyze two empirical games that differ in whether the MM is present. In all settings, there are $N \in \{25, 66\}$ background traders who arrive at a rate $\lambda_a = 0.075$. Each simulation run lasts T time steps. 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. If we interpret the tick size as a thousandth of a dollar (\$0.001), then the mean fundamental is \$100. The variance for the private value vector is $\sigma_{PV}^2 = 5 \times 10^6$. 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_r), fundamental shock variance (σ_s^2), and time horizon (T). The configurations of parameter settings for $N \in \{25, 66\}$ background traders are as follows.

A $\lambda_r = 0.0005, \sigma_s^2 = 1 \times 10^6,$
 $T \in \{1000, 2000, 4000, 12000, 60000\}$

B $\lambda_r = 0.005, \sigma_s^2 = 1 \times 10^6,$
 $T \in \{1000, 2000, 4000, 12000, 120000\}$

C $\lambda_r = 0.005, \sigma_s^2 = 5 \times 10^5,$
 $T \in \{1000, 2000, 4000, 12000, 120000\}$

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

6.2 EGTA Process

We model our market as a *role-symmetric game*, in which players are partitioned into roles, each with a specified strategy set. The payoff for playing a strategy in a particular role depends on the number of other agents playing each strategy in this role and the others, but not on how the strategies are mapped to players within the roles. Our two roles are background trader (25 or 66 players) and market maker (one player).

Even exploiting symmetry, game size grows exponentially in players and strategies, so it is computationally infeasible to analyze games with this many traders. We therefore apply aggregation to approximate the many-player games as games with fewer players. The specific technique we employ, called *deviation-preserving reduction* (DPR) [49], defines reduced-game payoffs in terms of payoffs in the full game as follows. Consider an N -player symmetric game, reduced to a k -player game. 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 full 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 .

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.

We iteratively apply EGTA to guide our exploration of the strategy space. Exploration starts with the most promising profiles and spreads through their neighbors, that is, those profiles related by single-agent deviations. The goal in this process is to identify Nash equilibria, and we focus our search on role-symmetric (pure or mixed) Nash equilibria (RSNE). As the observed payoffs from our simulator for a given profile are incrementally added, we analyze each successive intermediate game model by computing (mixed) equilibria for each complete subgame, defined as a set of strategies for which we have simulated all profiles. These represent candidates for equilibria of the full game, which we can refute by finding a beneficial deviation outside the strategy set, or confirm by examining all deviations without refuting. We continue to refine the empirical game with additional simulations until at least one equilibrium is confirmed and all non-confirmed candidates are refuted. In this study, we successfully found at least one and at most six non-trivial RSNEs for each game evaluated, with support sizes (i.e., numbers of strategies played with positive probability) up to five for background traders and up to three for MMs.

We utilize the EGTAOnline infrastructure [8] for conducting and managing our experiments, and our simulations are run on the high-performance computing cluster at the Center for Advanced Computing 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 (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.




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

6.3 Game without Market Making

The empirical games without MM cover 13 background-trader strategies: 12 versions of ZIR³ (see Table 1), and the no-trade strategy NOOP. We identified at least 1–3 ZIR equilibria⁴ for each of our thirty environments (see Tables 2 and 3). For each equilibrium, we calculated background-trader surplus by sampling 10,000 full-game profiles according to the equilibrium mixture, running one simulation per sampled profile and then recording the aggregate surplus.

6.4 Game with Market Making

Our games with MM include the 13 background-trader strategies from the no-MM treatment above, plus eight strategies for the MM role. The MMs employed in our game analysis are all parametric variants of FUNDMM (Section 5), with $K = 100$ rungs spaced ξ units apart. One variant sets its spread using the current price quote at reentry, whereas the others employ a constant spread $\omega \in \{64, 128, 256, 512, 1024, 2048\}$. Rung size ξ is 100, plus two variants with $\xi = 50$ for $\omega \in \{256, 512\}$. The equilibria found are presented in Tables 4 and 5. Background-trader surplus and MM profit are estimated for each equilibrium based on the sampling method described for the no-MM game above.

³Not every ZIR strategy appears in every environment; for instance, in the $N = 25$ case, agents playing strategies with $R_{\min} \geq 2500$ never trade, so we exclude such strategy variants from the analysis of those environments.

⁴We report only one equilibrium when multiple close equilibria (for which the mixture probabilities are nearly identical) exist. The full details of all equilibria in games both with and without MM are available in an online appendix (<http://hdl.handle.net/2027.42/110647>).

R_{\min}	0	0	0	0	0	0	500	0	0	2500	5000	5000
R_{\max}	125	125	250	500	1000	1000	1000	1500	2500	10000	10000	15000
η	0.8	1	1	1	0.8	1	0.4	0.6	1	1	1	1

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

Env	surplus	R_{mid}	η
A1	12997	416	0.878
A1	11958	586	0.721
A2	16501	377	0.878
A4	22384	339	0.909
A12	34853	283	0.930
A12	36668	161	0.919
A12	34865	338	0.930
A60	43343	300	1
B1	27387	616	0.683
B2	35648	572	0.706
B4	39180	633	0.647
B12	42635	393	0.885
B120	41076	746	0.598
B120	41469	630	0.689
B120	41066	665	0.554
C1	29433	379	0.959
C2	37533	443	0.845
C4	40926	469	0.842
C12	43001	436	0.932
C120	43308	350	0.948

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 ZIR range $[R_{\min}, R_{\max}]$) and threshold η . Values presented are the average over strategies in the profile, weighted by mixture probabilities.

6.5 Comparison of Market Performance

Our findings with regard to the central question in this paper are presented in Figure 3. 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 only two environments (A60 and B120) 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 A60, B12, B120, C12, and C120, the total social welfare, regardless of whether MM is present, is very close to the socially efficient outcome of 44155. That is, the ZIR background traders in these environments extract nearly all 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 private values on the right 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

Env	surplus	R_{mid}	η
A1	4110	634	0.678
A2	5103	677	0.576
A4	8543	125	1
A4	8237	285	1
A12	13639	125	1
A60	15798	349	0.980
A60	15910	347	0.935
A60	15629	412	0.844
B1	9591	684	0.532
B2	12792	583	0.604
B4	13823	650	0.593
B12	15673	500	1
B12	15627	521	0.909
B12	15675	445	0.885
B120	15658	747	0.551
C1	9831	656	0.566
C1	10010	626	0.543
C2	12678	626	0.599
C2	13058	553	0.637
C4	13967	646	0.566
C12	14688	658	0.622
C12	15792	500	0.872
C120	16047	500	0.8
C120	15012	677	0.541

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

outcomes, the MM can provide little benefit to overall welfare, and its profits tend to come out of background-trader surplus.

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 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 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. In our study, for markets populated by 25 background traders (Figure 3(b)), the market maker significantly improves both welfare and background-trader surplus for $T \in \{1000, 2000\}$. As for $N = 66$, the $N = 25$ background traders in environments with sufficiently long T get close to the social optimum without MM. The presence of the MM further improves overall social welfare in these environments up to the optimal value, enabling extraction of all surplus in the market at some loss to the background traders.

For both population sizes, background traders generally shade less when MM is present. In particular, for the $N = 25$ environments in which background-trader surplus improves with MM, the midpoint R_{mid} of the ZIR bid range is significantly lower in

Env	surplus	profit	R_{mid}	η	ω	ξ
A1	13379	1490	154	0.946	256	50
A2	17002	1558	151	0.935	256	50
A4	22520	2000	181	1	256	50
A12	34826	2745	182	1	256	50
A60	41049	2250	500	0.848	<i>var</i>	100
B1	29065	223	487	0.788	512	100
B2	36679	541	453	0.875	512	100
B2	37424	623	376	0.924	512	50
B4	40733	528	500	0.889	256	100
B12	42265	938	356	0.569	256	100
B12	42413	744	452	1	256	100
B120	40969	1837	673	0.607	256	100
B120	42282	1414	589	0.700	256	50
B120	42374	1272	621	0.682	311	61
B120	42648	994	621	0.682	312	61
B120	42769	1070	500	0.800	256	50
B120	42189	1268	556	0.735	128	100
C1	30049	794	311	0.962	512	100
C2	38551	1633	202	1	256	50
C4	39991	1403	500	0.903	256	79
C12	41488	2000	407	0.844	256	100
C120	41732	1702	455	0.909	256	54

Table 4: Role-symmetric equilibria for games with market makers, $N = 66$, 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 ZIR 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. Spread *var* indicates context-varying spread based on current quotes, so an average is unavailable.

Env	surplus	profit	R_{mid}	η	ω	ξ
A1	4704	890	100	0.919	512	100
A2	6318	770	154	1	256	100
A4	8119	893	214	1	512	50
A12	12556	1861	109	1	256	50
A60	14335	2141	310	0.909	253	51
B1	11305	651	274	0.933	512	72
B2	13232	881	402	0.818	512	100
B4	14066	704	531	0.794	512	100
B12	15143	945	399	0.854	512	50
B12	15038	1277	504	0.833	256	100
B120	15234	1309	506	0.756	257	50
C1	10901	779	300	0.960	512	100
C1	11247	935	217	1	512	50
C2	13858	1054	247	1	512	50
C4	15319	911	223	1	256	88
C12	15030	1350	500	0.916	256	100
C120	15352	1313	511	0.904	256	50

Table 5: Role-symmetric equilibria for games with market makers, $N = 25$, calculated from the (5,1)-player DPR approximation. Data presented is as for Table 4.

equilibria with MM. 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.

Finally, we evaluate liquidity for the maximum-welfare RSNE (Figure 4), 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 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 significantly narrow spreads nearly down to the levels present in the more populous environments. 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.

7. CONCLUSIONS

In this study, we employed a simulation-based approach to compare several parameterized environments with and without a market maker. 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 others' 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 strong tendency of MMs to improve the welfare of impatient investors (those in environments with short time horizons), especially in thin markets, but no such benefit in environments comprising more patient investors.

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 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 [49], 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. Further investigation may yield improved versions of ZIR or other strategies (for example adaptive variants [13, 45]) that could alter equilibrium findings. Similar improvements may be found on the MM side, for instance with strategies incorporating learning [1].

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.

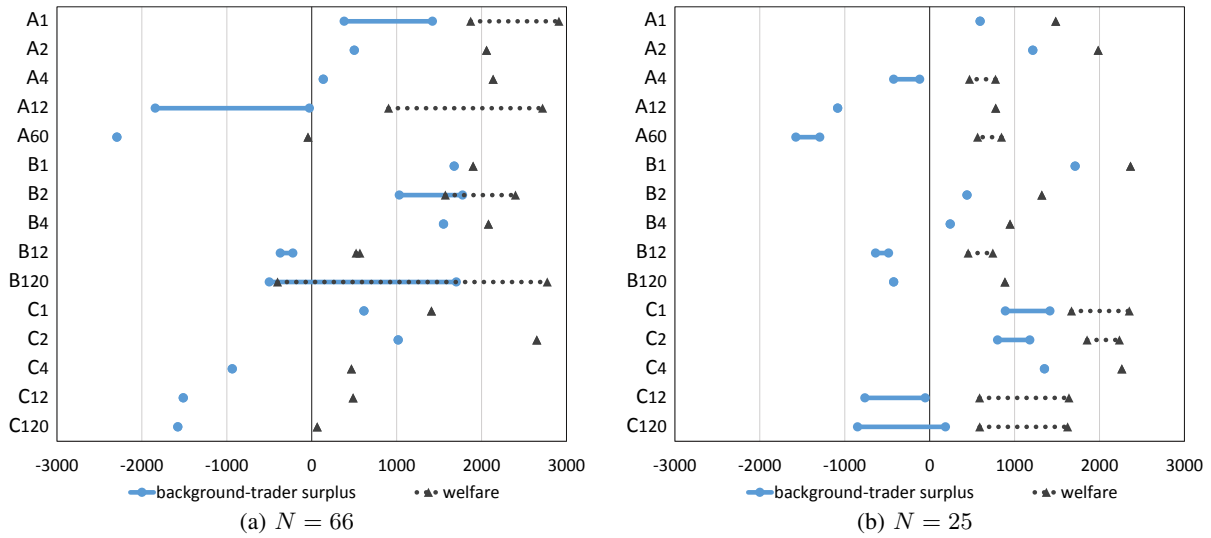


Figure 3: The effect of MM presence on background-trader surplus and social welfare in equilibrium, across all environments. Differences are presented as ranges, reflecting the multiplicity of equilibria found in some environments. The left point of each range is the minimum gain (in some cases a loss), 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 MM and the lowest without MM.



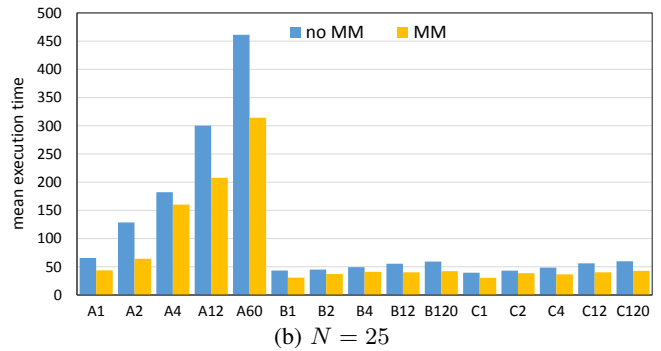



Figure 4: Comparison of background-trader execution time (Figures 4(a) and 4(b)) and median spread (Figures 4(c) and 4(d)) for the maximum-welfare RSNE in each environment, with and without MM. Mixed-strategy RSNE are approximated by profiles with trader population proportions corresponding to the strategy probabilities. Each bar is compiled from 10,000 samples.

Acknowledgments

Zhiyi Zhang’s senior honors thesis served as a pilot study of MM strategies. Mason Wright improved our background-trader strategies. We are also grateful to Uday Rajan for helpful advice on

this work, and to the 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.

References

- [1] J. Abernethy and S. Kale. Adaptive market making via online learning. In *Advances in Neural Information Processing Systems*, pages 2058–2066, 2013.
- [2] J. Abernethy, Y. Chen, and J. Wortman Vaughan. An optimization-based framework for automated market-making. In *12th ACM Conference on Electronic Commerce*, pages 297–306, 2011.
- [3] Y. Amihud and H. Mendelson. Dealership market: Market-making with inventory. *Journal of Financial Economics*, 8(1):31–53, 1980.
- [4] H. Bessembinder, J. Hao, and M. L. Lemmon. Why designate market makers? Affirmative obligations and market quality. Technical Report 989061, SSRN, 2011.
- [5] H. Bessembinder, J. Hao, and K. Zheng. Market making contracts, firm value, and the IPO decision. Technical Report 2130875, SSRN, 2014.
- [6] B. Biais, L. Glosten, and C. Spatt. Market microstructure: A survey of microfoundations, empirical results, and policy implications. *Journal of Financial Markets*, 8(2):217–264, 2005.
- [7] S. Brusco and M. O. Jackson. The optimal design of a market. *Journal of Economic Theory*, 88(1):1–39, 1999.
- [8] B.-A. Cassell and M. P. Wellman. EGTAOnline: An experiment manager for simulation-based game studies. In *Multi-Agent-Based Simulation XIII*, volume 7838 of *Lecture Notes in Artificial Intelligence*. Springer, 2013.
- [9] T. Chakraborty and M. Kearns. Market making and mean reversion. In *11th ACM Conference on Electronic Commerce*, pages 307–314, 2011.
- [10] N. T. Chan and C. Shelton. An electronic market-maker. Technical Report AI Memo 2001-005, Massachusetts Institute of Technology, 2001.
- [11] Y. Chen and D. M. Pennock. A utility framework for bounded-loss market makers. In *23rd Conference on Uncertainty in Artificial Intelligence*, pages 49–56, 2007.
- [12] B. Chowdhry and V. Nanda. Multimarket trading and market liquidity. *Review of Financial Studies*, 4(3):483–511, 1991.
- [13] D. Cliff. ZIP60: Further explorations in the evolutionary design of online auction market mechanisms. *IEEE Transactions on Evolutionary Computation*, 13(1):3–18, 2009.
- [14] V. Darley, A. Outkin, T. Plate, and F. Gao. Sixteenths or pennies? Observations from a simulation of the NASDAQ stock market. In *IEEE/IAFE/INFORMS Conference on Computational Intelligence for Financial Engineering*, pages 151–154, 2000.
- [15] R. Das, J. E. Hanson, J. O. Kephart, and G. Tesauro. Agent-human interactions in the continuous double auction. In *17th International Joint Conference on Artificial Intelligence*, pages 1169–1176, 2001.
- [16] S. Das. A learning market-maker in the Glosten–Milgrom model. *Quantitative Finance*, 5(2):169–180, 2005.
- [17] S. Das. The effects of market-making on price dynamics. In *7th International Joint Conference on Autonomous Agents and Multiagent Systems*, volume 2, pages 887–894, 2008.
- [18] S. Das and M. Magdon-Ismael. Adapting to a market shock: Optimal sequential market-making. In *Advances in Neural Information Processing Systems*, pages 361–368, 2008.
- [19] Y. Feng, R. Yu, and P. Stone. Two stock-trading agents: Market making and technical analysis. In *Agent-Mediated Electronic Commerce V. Designing Mechanisms and Systems*, pages 18–36. Springer, 2004.
- [20] S. Frey and J. Grammig. Liquidity supply and adverse selection in a pure limit order book market. *Empirical Economics*, 30(4):1007–1033, 2006.
- [21] D. Friedman. The double auction market institution: A survey. In D. Friedman and J. Rust, editors, *The Double Auction Market: Institutions, Theories, and Evidence*, pages 3–25. Addison-Wesley, 1993.
- [22] M. B. Garman. Market microstructure. *Journal of Financial Economics*, 3(3):257–275, 1976.
- [23] L. R. Glosten. Is the electronic open limit order book inevitable? *Journal of Finance*, 49(4):1127–1161, 1994.
- [24] L. R. Glosten and P. R. Milgrom. Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics*, 14(1):71–100, 1985.
- [25] D. K. Gode and S. Sunder. Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy*, 101(1):119–137, 1993.
- [26] R. L. Goettler, C. A. Parlour, and U. Rajan. Informed traders and limit order markets. *Journal of Financial Economics*, 93(1):67–87, 2009.
- [27] S. J. Grossman and M. H. Miller. Liquidity and market structure. *Journal of Finance*, 43(3):617–633, 1988.
- [28] R. Hanson. Logarithmic market scoring rules for modular combinatorial information aggregation. *Journal of Prediction Markets*, 1(1):3–15, 2007.
- [29] J. Hasbrouck and G. Sofianos. The trades of market makers: An empirical analysis of NYSE specialists. *Journal of Finance*, 48(5):1565–1593, 1993.
- [30] J. Huang and J. Wang. Market liquidity, asset prices, and welfare. *Journal of Financial Economics*, 95(1):107–127, 2010.
- [31] R. D. Huang and H. R. Stoll. Dealer versus auction markets: A paired comparison of execution costs on NASDAQ and the NYSE. *Journal of Financial Economics*, 41(3):313–357, 1996.
- [32] J. Jumadinova and P. Dasgupta. A comparison of different automated market-maker strategies. In *12th Workshop on Agent-Mediated Electronic Commerce*, pages 141–154, 2010.
- [33] A. S. Kyle. Continuous auctions and insider trading. *Econometrica*, 53(6):1315–1335, 1985.

- [34] J. C. Leach and A. N. Madhavan. Intertemporal price discovery by market makers: Active versus passive learning. *Journal of Financial Intermediation*, 2(2):207–235, 1992.
- [35] B. LeBaron. Agent-based computational finance. In L. Tesfatsion and K. L. Judd, editors, *Handbook of Agent-Based Computational Economics*, pages 1187–1233. Elsevier, 2006.
- [36] A. Madhavan. Market microstructure: A survey. *Journal of Financial Markets*, 3(3):205–258, 2000.
- [37] S. Manaster and S. C. Mann. Life in the pits: Competitive market making and inventory control. *Review of Financial Studies*, 9(3):953–975, 1996.
- [38] A. J. Menkveld. High frequency trading and the new market makers. *Journal of Financial Markets*, 16(4):712–740, November 2013.
- [39] M. O’Hara. *Market Microstructure Theory*, volume 108. Blackwell Cambridge, 1995.
- [40] M. O’Hara and G. Oldfield. The microeconomics of market making. *Journal of Financial and Quantitative Analysis*, 21(4):361–376, 1986.
- [41] G. Saar. Specialist markets. In *Encyclopedia of Quantitative Finance*. Wiley Online Library, 2010.
- [42] P. Sandås. Adverse selection and competitive market making: Empirical evidence from a limit order market. *Review of Financial Studies*, 14(3):705–734, 2001.
- [43] R. A. Schwartz and L. Peng. Market makers. In *Encyclopedia of Finance*, pages 487–489. Springer, 2013.
- [44] D. J. Seppi. Liquidity provision with limit orders and a strategic specialist. *Review of Financial Studies*, 10(1): 103–150, 1997.
- [45] P. Vytelingum, D. Cliff, and N. R. Jennings. Strategic bidding in continuous double auctions. *Artificial Intelligence*, 172(14):1700–1729, 2008.
- [46] E. Wah and M. P. Wellman. Latency arbitrage, market fragmentation, and efficiency: A two-market model. In *14th ACM Conference on Electronic Commerce*, pages 855–872, 2013.
- [47] M. P. Wellman. Methods for empirical game-theoretic analysis (Extended abstract). In *21st National Conference on Artificial Intelligence*, pages 1552–1555, 2006.
- [48] M. P. Wellman. *Trading Agents*. Morgan & Claypool, 2011.
- [49] B. Wiedenbeck and M. P. Wellman. Scaling simulation-based game analysis through deviation-preserving reduction. In *11th International Conference on Autonomous Agents and Multiagent Systems*, pages 931–938, 2012.