Welfare Effects of Market Making in Continuous Double Auctions: Extended Abstract*

Elaine Wah, Mason Wright, Michael P. Wellman
Computer Science & Engineering
University of Michigan
{ewah,masondw,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 a variety of 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. Introducing additional market makers increases these benefits, as competition drives market makers to provide liquidity at lower price spreads.

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 and 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 and 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, the two 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 becoming predominant [Frey and Grammig, 2006], 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 trade-off and its implications for MM strategies [Glosten and Milgrom, 1985; Kyle, 1985]; other prior research investigates the effects of MM on liquidity (e.g., as measured by price spreads) [Das and Magdon-Ismail, 2008] and price discovery [Leach and Madhavan, 1992]. Although liquidity and price discovery are generally expected to improve 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 *et al.*, 2011; 2015], the cost structure of market participation in supplying liquidity [Huang and Wang, 2010], and trading mechanisms to incentivize market making [Brusco and Jackson, 1999].

The impact of MMs on the welfare of other traders may depend on whether a market has just one monopolistic MM or a set of competitive MMs. 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 and Hughson, 1997]; where background traders can deal with each MM separately [Dennert, 1993]; or, most similar to this study, where MMs compete in the same market through a common limit order book [Biais *et al.*, 2000; Glosten and Milgrom, 1985].

*Extracts from a report (in journal submission) that extends and supersedes a paper by the same title presented at the *14th Int'l Conference on Autonomous Agents and Multiagent Systems* [Wah and Wellman, 2015].

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 the CDA mechanism in a market environment comprising multiple background traders, and in some cases one or more MMs. 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 the market according to a stochastic arrival process, each time to offer to buy or sell a single unit of the security. MMs in our model have no private value, and thus aim to profit by maintaining buy and sell offers with a positive price spread.

To compare outcomes with and without market making, we search for strategy configurations where traders best-respond to the environment and other-agent behavior. As it appears intractable to compute an analytic game-theoretic solution for this model, 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 and 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. We derive empirical equilibria with and without market making in 21 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 three equilibrium settings, the presence of MM increases overall welfare (background-trader surplus combined with MM profit).

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 the majority of environments tested, and tended to do so particularly in relatively thin markets. For impatient investors with relatively infrequent trading opportunities, the MM was more beneficial the shorter the trading horizon.

In contrast, 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 bid–ask 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 MMs in alleviating allocative inefficiencies. Section 3 discusses the market environment. Section 4 presents our results. We conclude in Section 5.

2 Motivating Example

We illustrate the problem of allocative inefficiency in CDAs, and the influence of market makers, with the following simple example. Suppose we have 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 underperforming by Δ.

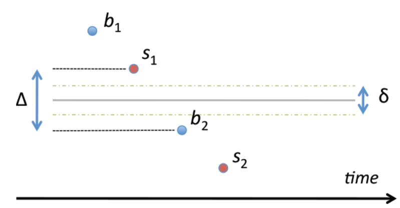

Figure 1: A sequence of CDA orders leading to a suboptimal allocation.

Now suppose there is an MM 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* prices[1] $[b_2, s_1]$, then the efficient allocation is achieved for *any* sequence. In such cases, the MM accrues profit 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 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]

[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.

[2] A modest amount of bid shading can also prevent inefficient

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 and 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. 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 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. 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.

Each background trader has an individual valuation for the security composed of private and common components. There is an extensive literature on autonomous bidding strategies for CDAs [Das *et al.*, 2001; Friedman, 1993; Wellman, 2011]. In this study, we consider trading strategies in the so-called *Zero Intelligence* (ZI) family [Gode and Sunder, 1993]. The background traders arrive at the market according to an independent Poisson process per trader, with rate λ_{BG}. They subsequently reenter the market, with time between entries distributed exponentially at the same rate. At the end of the simulation period, background traders liquidate their accumulated inventory.

In our model, 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, composed of K rungs spaced ξ ticks apart. Each MM arrives at time 0 and reenters the market according its own, independent Poisson process with rate λ_{MM}. 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 MM's total profit is defined by the sum of trading cash flow plus liquidation proceeds.

4 Empirical Game-Theoretic Analysis

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 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) [Wellman, 2006]. 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 3, using an extension of the discrete-event market simulator developed for our previous study of latency arbitrage [Wah and Wellman, 2013].

We generate data for various combinations of the strategies introduced above, each sampled over many runs 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.

We evaluate the performance of background traders and the MM within 21 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 $T \in \{1, 4, 12, 24\} \times 10^3$.

trades, and indeed equilibrium shading strategies often lead to more efficient outcomes than truthful bidding in CDAs [Zhan and Friedman, 2007].

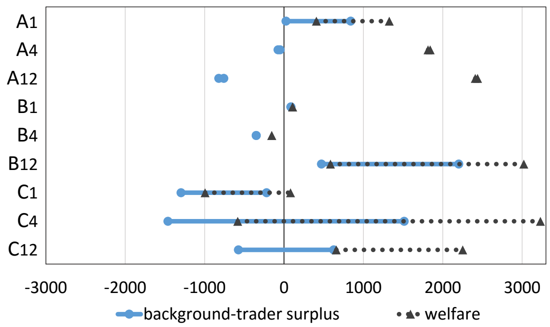

(a) $N = 66$

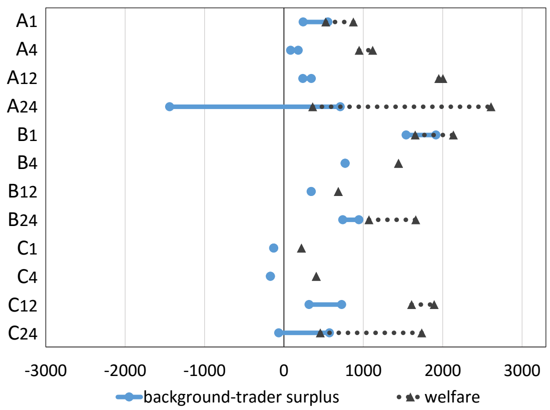

(b) $N = 25$

Figure 2: The effect of presence of a single MM on background-trader surplus and social welfare in equilibrium, across all environments. The left point of each range is the minimum gain (or 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.

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). 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 $\lambda_{BG} = 0.0005$, $\sigma_s^2 = 1 \times 10^6$

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

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

We describe each environment by its configuration label, followed by time horizon (in thousands).

Our central findings are presented in Figure 2. 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.

The trading horizon T reflects whatever might limit an investor's patience (liquidity needs, portfolio hedging, cost of monitoring, etc.). 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 three of these four cases. 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. We also find evidence 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 MM.

To evaluate the effects of competition among MMs, we add a second MM to seven of our $N = 25$ environments. In each case, every equilibrium in the two-MM setting has 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. We increase the number of MMs to four in two of these environments, finding that background-trader surplus in every equilibrium is higher than in any equilibrium for corresponding settings with zero, one, or two MMs. This suggests that as the number of MMs present increases beyond two, the benefit of MM competition to background traders continues to increase. There is no clear trend in social welfare as the number of MMs increases from two to four, however, as the greater background-trader surplus with four MMs is offset by lower profits for each MM.

5 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.*

Our analysis demonstrates the generally beneficial effects of MMs on efficiency, and shows that whether these benefits accrue to background investors depends on market characteristics. We find a tendency of a monopolist MM to improve the welfare of impatient investors (those in thin markets or with fewer opportunities to trade), 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 MM profit.

References

[Bernhardt and Hughson, 1997] Dan Bernhardt and Eric Hughson. Splitting orders. *Review of Financial Studies*, 10(1):69–101, 1997.

[Bessembinder *et al.*, 2011] Hendrik Bessembinder, Jia Hao, and Michael L. Lemmon. Why designate market makers? Affirmative obligations and market quality. Technical Report 989061, SSRN, 2011.

[Bessembinder *et al.*, 2015] Hendrik Bessembinder, Jia Hao, and Kuncheng Zheng. Market making contracts, firm value, and the IPO decision. *Journal of Finance*, 70(5):1997–2028, 2015.

[Biais *et al.*, 2000] Bruno Biais, David Martimort, and Jean-Charles Rochet. Competing mechanisms in a common value environment. *Econometrica*, 68(4):799–837, 2000.

[Brusco and Jackson, 1999] Sandro Brusco and Matthew O. Jackson. The optimal design of a market. *Journal of Economic Theory*, 88(1):1–39, 1999.

[Chakraborty and Kearns, 2011] Tanmoy Chakraborty and Michael Kearns. Market making and mean reversion. In *11th ACM Conference on Electronic Commerce*, pages 307–314, 2011.

[Das and Magdon-Ismail, 2008] Sanmay Das and Malik Magdon-Ismail. Adapting to a market shock: Optimal sequential market-making. In *Advances in Neural Information Processing Systems*, pages 361–368, 2008.

[Das *et al.*, 2001] Rajarshi Das, James E. Hanson, Jeffrey O. Kephart, and Gerald Tesauro. Agent-human interactions in the continuous double auction. In *17th International Joint Conference on Artificial Intelligence*, pages 1169–1176, 2001.

[Dennert, 1993] Jürgen Dennert. Price competition between market makers. *Review of Economic Studies*, 60(3):735–751, 1993.

[Frey and Grammig, 2006] Stefan Frey and Joachim Grammig. Liquidity supply and adverse selection in a pure limit order book market. *Empirical Economics*, 30(4):1007–1033, 2006.

[Friedman, 1993] Daniel Friedman. The double auction market institution: A survey. In Daniel Friedman and John Rust, editors, *The Double Auction Market: Institutions, Theories, and Evidence*, pages 3–25. Addison-Wesley, 1993.

[Glosten and Milgrom, 1985] Lawrence R. Glosten and Paul R. Milgrom. Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics*, 14(1):71–100, 1985.

[Gode and Sunder, 1993] Dhananjay K. Gode and Shyam 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.

[Huang and Stoll, 1996] Roger D. Huang and Hans 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.

[Huang and Wang, 2010] Jennifer Huang and Jiang Wang. Market liquidity, asset prices, and welfare. *Journal of Financial Economics*, 95(1):107–127, 2010.

[Kyle, 1985] Albert S. Kyle. Continuous auctions and insider trading. *Econometrica*, 53(6):1315–1335, 1985.

[Leach and Madhavan, 1992] J. Chris Leach and Ananth N. Madhavan. Intertemporal price discovery by market makers: Active versus passive learning. *Journal of Financial Intermediation*, 2(2):207–235, 1992.

[Saar, 2010] Gideon Saar. Specialist markets. In *Encyclopedia of Quantitative Finance*. Wiley Online Library, 2010.

[Schwartz and Peng, 2013] Robert A. Schwartz and Lin Peng. Market makers. In *Encyclopedia of Finance*, pages 487–489. Springer, 2013.

[Wah and Wellman, 2013] Elaine Wah and Michael P. Wellman. Latency arbitrage, market fragmentation, and efficiency: A two-market model. In *14th ACM Conference on Electronic Commerce*, pages 855–872, 2013.

[Wah and Wellman, 2015] Elaine Wah and Michael P. Wellman. Welfare effects of market making in continuous double auctions. In *14th International Conference on Autonomous Agents and Multiagent Systems*, pages 57–66, 2015.

[Wellman, 2006] Michael P. Wellman. Methods for empirical game-theoretic analysis (Extended abstract). In *21st National Conference on Artificial Intelligence*, pages 1552–1555, 2006.

[Wellman, 2011] Michael P. Wellman. *Trading Agents*. Morgan & Claypool, 2011.

[Zhan and Friedman, 2007] Wenjie Zhan and Daniel Friedman. Markups in double auction markets. *Journal of Economic Dynamics and Control*, 31(9):2984–3005, 2007.