An Agent-Based Model of Financial Benchmark Manipulation

Megan Shearer [1] **Gabriel Rauterberg** [2] **Michael P. Wellman** [1]

Abstract

Financial benchmarks estimate market values or reference rates used in a wide variety of contexts, but are often calculated from data generated by parties who have incentives to manipulate these benchmarks. Since the the London Interbank Offered Rate (LIBOR) scandal in 2011, market participants, scholars, and regulators have scrutinized financial benchmarks and the ability of traders to manipulate them. We study the impact on market quality and microstructure of manipulating transaction-based benchmarks in a simulated market environment. Our market consists of a single benchmark manipulator with external holdings dependent on the benchmark, and numerous background traders unaffected by the benchmark. All market participants use zero-intelligence trading strategies. When these agents trade under equilibrium settings in our market environment with and without benchmark manipulation, we find that the total surplus of all market participants who are trading increases with manipulation. However, the aggregated market surplus decreases for all trading agents, and the market surplus of the manipulator decreases, so the manipulator's surplus from the benchmark significantly increases. This entails under natural assumptions that the market and any third parties invested in the opposite side of the benchmark from the manipulator are negatively impacted by this manipulation.

1. Introduction

Financial benchmarks play a pervasive role in modern commerce and finance. A benchmark is a numerical estimate of some market value, such as the price of an asset. Benchmarks are employed by market participants for various purposes, including as reference measures for asset values (e.g., the S&P 500), interest rates (LIBOR), and market volatility (VIX); to define derivative instruments; or as price terms in contracts (Gellasch & Nagy, 2019). Benchmarks in the form of reference measures can provide a concise reflection of market realities, thereby assisting decision-making in the real economy. As such, accurate benchmark prices represent a positive externality from functional financial markets (Bond et al., 2012). Their use in financial instruments and contracts also provide a valuable function in commerce and risk management.

Given their role in market decisions and contracts, some entities may have strong incentives to manipulate benchmarks. For instance, since 2012, multiple parties have been convicted of manipulating the *London Interbank Offered Rate* (LIBOR), an estimate of the rate at which banks can borrow from each other (McBride, 2016). LIBOR supports more than $300 trillion worth of loans around the world. More recently, in February 2018, there were accusations of manipulation in the Chicago Board Options Exchange (CBOE) *Volatility Index* (VIX), a measure of U.S. stock market volatility based on the cost of buying certain options (Banerji, 2018). LIBOR was particularly vulnerable to manipulation because it is calculated using self-reported data provided by parties with conflicts of interest regarding the benchmark's value (Duffie & Dworczak, 2018; Gellasch & Nagy, 2019). In the wake of the LIBOR scandal, regulators, academics, and market participants lobbied for a transaction-based replacement for LIBOR, such as the Secured Overnight Finance Rate (SOFR) or the U.S. Dollar Intercontinental Exchange (ICE) Bank Yield Index (Duffie & Dworczak, 2018; ICE Benchmark Administration Limited, 2019). Whereas it may be harder to manipulate transaction-based benchmarks, it is still possible, as in the alleged manipulation of the VIX in 2018 and the World Markets/Reuters Closing Spot Rates (WM/R FX rates) in 2014 (Boyle, 2014). Finding accurate and robust transaction-based benchmarks is a daunting but necessary task.

We introduce an agent-based model to shed light on how benchmark manipulation can operate, with the ultimate goal of supporting the design of manipulation-resistant benchmarks. Our model simulates a financial market, showing how a party with vested interest in a transaction-based benchmark can manipulate this benchmark through trading. The benchmark manipulator extends the behavior of

[1]Computer Science & Engineering, University of Michigan, Ann Arbor, USA [2]University of Michigan Law School, University of Michigan, Ann Arbor, USA. Correspondence to: Megan Shearer <shearerj@umich.edu>.

ICML-19 Workshop on AI in Finance, Long Beach, CA, 2019.

a zero intelligence trader (Gode & Sunder, 1993; Farmer et al., 2005), adjusting its offers systematically in order to influence the benchmark in a certain direction. For example, if the manipulator wants to lower the benchmark, the manipulator offers to sell shares at a lower price to shift the benchmark down. The manipulator might take a loss in the market through resulting sales, but may still earn a net profit if it successfully shifts the benchmark enough to impact its external holdings linked to the benchmark.

Prior work has employed theoretical models and historical data to study benchmark manipulation in financial markets (Bariviera et al., 2016; Duffie & Dworczak, 2018; Duffie, 2018; Eisl et al., 2017; Rauch et al., 2013). Using a simulated market allows us to incorporate complex details of *market microstructure*, representing the actual mechanics of trade, interactions among market participants, and the structure of the market. By combining the agent-based model with game-theoretic reasoning, we can also consider the response of strategic agents to the presence of a benchmark manipulator, and consider a wide range of market settings, benchmark designs, and trading strategy options.

We employ a standard market mechanism organized around a *limit order book* for a single security. We assume a benchmark defined by transaction prices on this security. Trading agents may submit buy and sell orders, with orders executing with zero delay when matched, or resting in the order book pending execution against a subsequent order. The market is populated by a single manipulator, along with background agents who do have private reasons to trade the security but no interests dependent on the benchmark. The manipulation activity potentially impacts the background agents through the market for this security, as well as (unmodeled) external parties who do not participate in the market but do have interests dependent on the benchmark.

We determine the impact of benchmark manipulation by comparing strategic equilibria when the manipulator chooses to manipulate and to not manipulate. We find in particular settings that manipulation is profitable overall to the manipulator. The manipulation activity itself is costly, in that the manipulator must take trading losses to move the benchmark. The background traders actually benefit from the manipulation, as their aggregate gains from trading increase. The external parties dependent on the opposite side of the benchmark are the real losers from the manipulation, with their losses captured in part by the manipulator and in part by the background agents whose trading is effectively subsidized.

This paper is organized as follows. Following a discussion of related work in the next section, we describe the market environment in Section 3. Section 4 discusses the benchmark manipulator and the trading strategy it employs in this paper. Section 5 presents the results with and without benchmark manipulation. Section 6 analyzes current and possible policies around financial benchmarks. We conclude in Section 7.

2. Related Work

Our market model is based on the models of Wah et al. (2017) and Wang et al. (2018). We use a similar market structure, as well as similar zero-intelligence (ZI) background agents (Gode & Sunder, 1993; Farmer et al., 2005). However, these studies analyze questions that are unrelated to benchmark manipulation.

The majority of prior work on benchmark manipulation is either theoretical or utilizes historical market data. Duffie & Dworczak (2018) construct a theoretical model to analyze the robustness and bias of certain benchmark calculations, and find that *volume-weighted average price* (VWAP) results in the most robust and unbiased benchmark estimations. Another theoretical study constructs construct new market models that prove to lead to more robust and unbiased benchmark calculations, providing insight into the quality of benchmark design and market structure (Duffie, 2018). Prior work which use historical market data detects and investigates previous instances of manipulation on LIBOR. Bariviera et al. (2016) and Eisl et al. (2017) use historical data to find instances of manipulations and provide suggestions for more robust benchmarks and regulation. Rauch et al. (2013) also use historical data to find instances of benchmark manipulation in LIBOR and investigate which banks were potentially involved in the 2011 scandal. There is also an extensive policy discussion around reforms of LIBOR and other financial benchmarks (Duffie & Stein, 2015; Gellasch & Nagy, 2019; IOSCO, 2013; Verstein, 2015).

To our knowledge, little to no prior work addresses benchmark manipulation in a simulated environment. Analyzing benchmark manipulation in a simulated market, rather than a theoretical model or historical data, has the benefit of controlling traders' intent to manipulate the benchmark. It also adds additional realism to the investigation; prior theoretical work features markets without any microstructure, that is, without any details regarding the mechanics of trade. Our approach allows us to study the impact of the manipulation on the market, other agents, and the benchmark.

3. Market Environment

3.1. Market Mechanism

To determine the effects of benchmark manipulation, we construct a simple continuous double auction (CDA) where one security is traded. Our market model is similar to that of Wah et al. (2017) and Wang et al. (2018). Prices are discrete and integer multiples of the tick size. Time is also discrete with the finite time horizon T. We denote the fundamental value of the underlying security at time t by r_t. This fundamental varies throughout the simulation by a stochastic mean-reverting process, more formally:

$$r_t = \max\{0, \kappa + (1-\kappa)r_{t-1} + u_t\bar{r}\}, t \in [0,T]; r_0 = \bar{r}, \quad (1)$$

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

Agents trade on the single security by submitting *limit orders*, which are orders that specify a limit price at which they are willing to transact. We only consider single-unit trades. The market mechanism tracks the orders which have not executed yet through a limit order book of the aggregated resting orders.

3.2. Benchmark

After the termination of the market at time T, the benchmark β_T is calculated. This benchmark is an estimate of the current value of the security, which can later be used in decision-making, for example an interest rate for a contract. There are numerous ways to calculate a benchmark. In this study, we use *volume-weighted average price* (VWAP) of all transactions executed during the trading horizon as the benchmark. We use VWAP because prior work found this to be a fairly robust and unbiased benchmark (Duffie & Dworczak, 2018).

3.3. Agents in the Market

Our market is populated by numerous background agents and one benchmark manipulator. The background traders with private values depict investors with preferences on holding a long or short position in the underlying security. The benchmark manipulator also holds private values, but represents an investor who tries to maximize combined profits within the market and externally through the benchmark.

A private value vector Θ_i captures the position preference of background trader i. The vector Θ_i has length $2q_{\max}$, where $q_{\max}$ is the maximum number of units an agent can be long or short at any time. Element θ^i_{q+1} represents the marginal gain from buying an additional unit given the current net position q. We produce Θ_i from a set of $2q_{\max}$ values independently drawn from $\mathcal{N}(0, \sigma_{PV}^2)$. Next, we sort the elements in Θ_i in order of diminishing marginal utility, so that $\theta^{q'} > \theta^q$, for all $q' < q$. An agent's valuation for a unit of the security at time t is the sum of its private value at the current position q_t and the global fundamental at time T, more formally:

$$v_i(t) = \begin{cases} \hat{r}_t + \theta_i^{q+1} & \text{buying,} \\ \hat{r}_t - \theta_i^{q} & \text{selling.} \end{cases}$$

A background trader's final surplus is the final valuation of its holdings at time T. More formally, the final valuation of background trader i with final holdings H:

$$V_i = \begin{cases} r_T H + \sum_{k=1}^{k=H} \theta_i^k & \text{long positions } H > 0, \\ r_T H - \sum_{k=H+1}^{k=0} \theta_i^k & \text{short positions } H < 0. \end{cases}$$

The background agents arrive to the market by a Poisson process with rate λ_a. On each entry, the trader observes a new, unique noisy fundamental observation $o_t = r_t + n_t$ with the observation noise following $n_t \sim \mathcal{N}(0, \sigma_n^2)$. This noisy observation of the fundamental attempts to capture varying viewpoints of different market participants on the traded security. Since each agent has incomplete information about the true fundamental, so agents may benefit from considering market information from the aggregate observations of other traders. When an agent arrives at the market, it will withdraw any previous orders which have not executed in order to better react to its new observation of the fundamental. An agent then submits a new single-unit limit order to buy or sell with equal probability.

3.4. Background Trading Strategies

At the ending of a trading period, agents evaluate their holdings of the security based on their estimate from noisy observations of the final fundamental value. Given a new noisy observation o_t, an agent estimates the current fundamental by updating its Bayesian posterior mean $\tilde{r}_t$ and variance $\tilde{\sigma}_t^2$. We let t' represent the agent's preceding arrival time. The previous posteriors, $\tilde{r}_{t'}$ and $\tilde{\sigma}_{t'}^2$, are first updated by taking account of mean reversion for the interval since before the arrival ($\delta = t - t'$):

$$\tilde{r}_{t'} \leftarrow (1 - (1-\kappa)^\delta)\bar{r} + (1-\kappa)^\delta \tilde{r}_{t'};$$

$$\tilde{\sigma}_{t'}^2 \leftarrow (1-\kappa)^{2\delta}\tilde{\sigma}_{t'}^2 + \frac{1-(1-\kappa)^{2\delta}}{1-(1-\kappa)^2}\sigma_s^2.$$

The estimates for t are given by:

$$\tilde{r}_t = \frac{\sigma_n^2}{\sigma_n^2 + \tilde{\sigma}_{t'}^2}\tilde{r}_{t'} + \frac{\tilde{\sigma}_{t'}^2}{\sigma_n^2 + \tilde{\sigma}_{t'}^2} o_t; \tilde{\sigma}_t^2 = \frac{\sigma_n^2 \tilde{\sigma}_{t'}^2}{\sigma_n^2 + \tilde{\sigma}_{t'}^2}.$$

An agent calculates an estimate $\hat{r}_t$ at time t of the terminal fundamental r_T based on the posterior estimate of $\tilde{r}_t$ by adjusting for mean reversion, more formally:

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

We consider a version *zero intelligence* (ZI) for a strategy of background traders (Gode & Sunder, 1993; Farmer et al., 2005). A ZI agent decides a limit-order price by shading its evaluation of the security with a random offset uniformly drawn from $[R_{\min}, R_{\max}]$. More formally, a ZI trader i

arriving at time t with position q generates a limit price $p_i^{\mathbf{ZI}}(t)$ by:

$$p_i^{\mathbf{ZI}}(t) \sim \begin{cases} U[v_i(t) - R_{\max}, v_i(t) - R_{\min}] & \text{buying,} \\ U[v_i(t) + R_{\min}, v_i(t) + R_{\max}] & \text{selling.} \end{cases}$$

A ZI trader utilizes a strategic surplus threshold parameter $\eta \in [0, 1]$ to consider the current visible quoted price. If a ZI agent could gain a fraction η of its desired surplus by accepting the most competitive visible order, then it will take that quote by submitting a limit-order at the same price.

4. Benchmark Manipulation Strategies

We evaluate a benchmark manipulator who submits orders to a CDA like the background traders, and attempts to maximize its total profit between its market transactions and the benchmark. We assume that a benchmark manipulator has some holdings ψ_i in the benchmark it attempts to manipulate. The total surplus of a benchmark manipulator at the end of the time horizon T becomes:

$$V_i^{\mathbf{BMM}} = V_i^{\mathbf{ZI}} + \psi_i \beta_T \tag{3}$$

If $\psi_i < 0$, then the manipulator tries to lower the benchmark to maximize its total surplus, and the manipulator likely holds a short position at time T. While if $\psi_i > 0$, then the manipulator attempts to raise the benchmark and likely holds a long position at time T.

4.1. Manipulation with Zero-Intelligence

The first manipulation strategy we consider behaves similarly to the ZI background trader. However, the benchmark manipulator offsets the price determined by a ZI agent by $s(\psi_i)\varepsilon_i$, where $s(\psi_i)$ is the sign of the trader's benchmark holdings and ε is the amount the manipulator decides to offset the price. More formally, a manipulator with zero-intelligence i arriving at time t with position q generates a limit price $p_i(t)$ by:

$$p_i^{\mathbf{BMM}}(t) = p_i^{\mathbf{ZI}}(t) + s(\psi_i)\varepsilon \tag{4}$$

This manipulator sometimes utilizes the strategic surplus threshold parameter $\eta \in [0, 1]$ to consider the current visible quoted price. A benchmark manipulator uses η in the same manner as a ZI agent if $p_i^{\mathbf{BMM}}(t) \leq v_i(t)$ when it buys, and if $p_i^{\mathbf{BMM}}(t) \geq v_i(t)$ when it sells. However, if the $p_i^{\mathbf{BMM}}(t) > v_i(t)$ when the manipulator submits a buy limit-order, then it transacts with the most competitive visible order within $[v_i(t), p_i^{\mathbf{BMM}}(t)]$. Likewise, if the manipulator submits a sell limit-order and $p_i^{\mathbf{BMM}}(t) < v_i(t)$, then it transacts with the most competitive visible order within $[p_i^{\mathbf{BMM}}(t), v_i(t)]$.

5. Empirical Game-Theoretic Analysis

We propose a set of strategies for the background agents and benchmark manipulator, then find the combination of strategies that agents utilize in equilibria. To determine the impact of benchmark manipulation under equilibrium settings, we implement agent-based simulations and game-theoretic analysis.

5.1. Market Environment Settings

We explore a variety of market environments to analyze the robustness of our results. Following Wang et al. (2018), these environments vary by market shock σ_s^2 and observation noise σ_n^2. The variance of market shock dictates fluctuations to the true value of the fundamental, so higher shock variance leads to higher price volatility. The variance of observation noise controls the accuracy of agents' information on the fundamental. Thus, higher observation variance leads to less accurate fundamental information for agents. We consider three market environments, where the first has low shock and high observation noise (**LSHN**) with $\sigma_s^2 = 10^5$ and $\sigma_n^2 = 10^9$. The second consists of medium shock and medium observation noise (**MSMN**) with $\sigma_s^2 = 5 \times 10^5$ and $\sigma_n^2 = 10^6$. Lastly, the third contains high shock and low observation noise (**HSLN**) with $\sigma_s^2 = 10^6$ and $\sigma_n^2 = 10^3$.

The market is populated with ten background agents and one manipulator. The global fundamental time series of the market is produced by equation 1, with fundamental mean $\bar{r} = 10^5$ and mean reversion $\kappa = 0.05$. The finite time horizon of the market lasts $T = 10,000$ time steps. All agents arrive to the market according to a Poisson distribution with rate $\lambda_a = 0.005$. The maximum number of units all agents can hold at any time is $q_{\max} = 10$. Lastly, the private value variance is $\sigma_{PV}^2 = 5 \times 10^6$.

Each agent in our market can employ multiple trading strategies to maximize its profits. Table 1 specifies the strategies of the background traders and benchmark manipulator. The background agents can strategically choose values for parameters $R_{\min}$, $R_{\max}$, and η. The benchmark manipulator can also strategically choose values for these parameters as well as the price offset ε, which determines how much the manipulator shifts its order price to impact the benchmark. Our benchmark manipulator explores five strategies, but the manipulator chooses strategy $\mathbf{BMM_4}$ in an equilibrium of each market setting we explore in this paper.

The benchmark manipulator's surplus depends on its benchmark holdings $\psi_i \in \{-40, 40\}$. We denote an environment when the manipulator hopes to shift the benchmark down ($s(\psi_i) = -1$) as having a downward impact (**DI**). Whereas when the manipulator shifts the benchmark up ($s(\psi_i) = 1$), we label the environment as possessing an upward impact (**UI**).

Table 1. Strategies employed by the background traders ($\mathbf{ZI}$) and benchmark manipulator ($\mathbf{BMM}$). Strategies $\mathbf{ZI_4}$ and $\mathbf{BMM_4}$ are the same but are utilized by different roles.

Strategy	$\mathbf{ZI_1}$	$\mathbf{ZI_2}$	$\mathbf{ZI_3}$	$\mathbf{ZI_4}$	$\mathbf{ZI_5}$	$\mathbf{BMM_1}$	$\mathbf{BMM_2}$	$\mathbf{BMM_3}$	$\mathbf{BMM_4}$	$\mathbf{BMM_5}$
$R_{\min}$	0	0	0	0	0	0	0	0	0	0
$R_{\max}$	1000	1000	1000	1000	1000	1000	1000	1000	1000	1000
η	0.1	0.3	0.5	0.8	1.0	0.8	0.8	0.8	0.8	0.8
ε	0	0	0	0	0	0	500	1000	2500	5000

5.2. EGTA Process

Empirical game-theoretic analysis (EGTA) is a method for finding equilibria in games by a heuristic strategy space and simulated payoff data (Wellman, 2016). We use EGTA to find equilibria under different market settings. EGTA uses an iterative process to detect potential equilibria in subgames, incrementally add strategies and confirm or refute these potential equilibria by examining deviations until quiescence. Various prior studies on multi-agent systems have utilized EGTA, particularly when a market environment is complex and applying standard analytic methods is hard (Brinkman & Wellman, 2017; Wah et al., 2016).

We model our market as a role-symmetric game, which consists of players divided into roles that each have a designated strategy set. Our market splits players into two roles: background traders and a single benchmark manipulator. *Deviation-preserving reduction* (DPR) approximates many-player games as fewer player games through aggregation (Wiedenbeck & Wellman, 2012). We utilize DPR because game size grows exponentially in players and strategies. DPR preserves payoffs from single-player deviations and has been shown to generate good approximations in multiple settings.

While using DPR, we select ten background traders and one benchmark manipulator. Observing ten background traders ensures that the required aggregations from DPR come out as integers. Our reduced game becomes three background traders and one benchmark manipulator. In this setting one background agent deviates to a new strategy while the other nine background agents are further reduced to three. We sample at least 50,000 simulation runs for a specified strategy profile of each game to reduce sampling error resulting from stochastic market features.

5.3. Impact of Benchmark Manipulation

This paper focuses on how benchmark manipulation affects market microstructure. We use EGTA to analyze this impact on the different market environments we explore. Specifically, we calculate the market surplus and total surplus of the benchmark manipulator where total surplus aggregates the surplus from market trading (i.e., market surplus) and surplus from the benchmark holdings. We also find the surplus of the background traders. The total surplus and market surplus are the same for the background traders because we assume they are indifferent to the final benchmark calculation. Lastly, we study the market surplus, which is the aggregate surplus of the background traders and the market surplus of the benchmark manipulator, as well as the total surplus of the system, which we define as the aggregate surplus of the background traders and the total surplus of the benchmark manipulator.

Figure 1 depicts the total surplus and market surplus of the benchmark manipulator, respectively. The total surplus of the benchmark manipulator increases when it manipulates the benchmark. Of course, this agent would not manipulate the benchmark in an equilibrium setting if it did not increase its total payoff. The benchmark manipulator's market surplus actually decreases when it manipulates the benchmark. This happens because in order to manipulate the benchmark, this agent must trade at prices it does not believe reflect the fundamental. However, it is worthwhile to the agent to endure the decrease in market surplus because its profits from the benchmark more than cover the loss.

Figure 2 shows the surplus of the background agents with and without benchmark manipulation. The background agents profit from benchmark manipulation, which intuitively makes sense because the benchmark manipulator trades at worse prices to shift the benchmark. Therefore, background traders profit from transacting with the manipulator. Background traders, as an aggregate set, have no incentives to mitigate this manipulation by avoiding transactions with this manipulator, because they actually profit from manipulation.

Figure 3 shows the aggregate total surplus and aggregate market surplus, respectively. The aggregated total surplus, which we find by summing the total surplus of the benchmark manipulator and background traders, increases with benchmark manipulation. The aggregated market surplus we find by summing the market surplus of the benchmark manipulator and background traders, and this value decreases with market manipulation. To understand why the aggregate total surplus increases, but the aggregate market surplus decreases, recall that in equilibrium, the manipu-

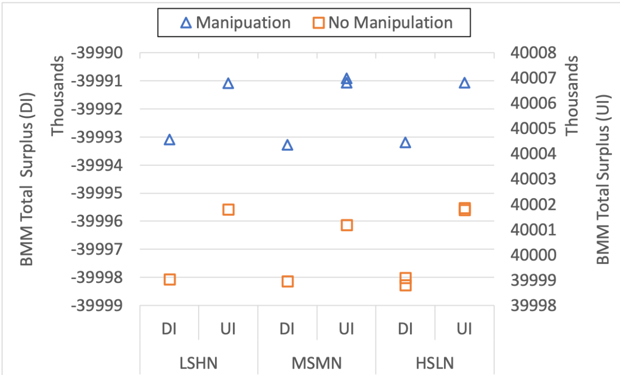

(a) Total (market plus benchmark) surplus of the benchmark manipulator with and without manipulation. The primary y-axis shows total surplus when the benchmark manipulator tries to shift the benchmark down (**DI**), while secondary y-axis shows the surplus when the manipulator tries to shift the benchmark up (**UI**).

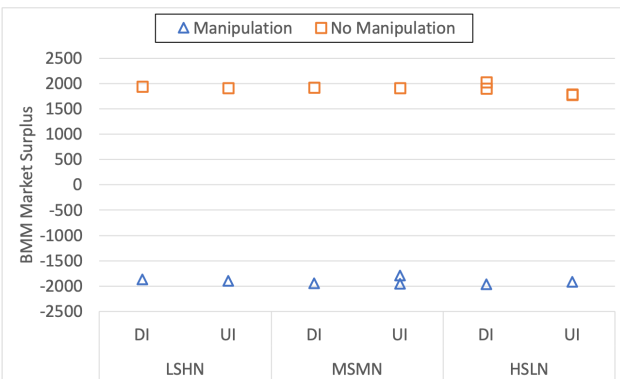

(b) Market surplus of the benchmark manipulator with and without manipulation.

Figure 1. With both figures, the x-axis represents different market environments with varying fundamental shock, observation variance, and benchmark impact.

lator only manipulates when gains from the manipulator's benchmark holdings exceed market losses. To further explore the consequences of benchmark manipulation, we consider a simplified example where the manipulator is involved in a *benchmark contract*, or a contract dependent on the benchmark, with one other party. We assume that both the benchmark contract and the market are *zero-sum games*, which are games where the total surplus is zero. In this toy example, the counterparty to the manipulator in the benchmark contract loses precisely what the benchmark manipulator gains from the benchmark contract. Though to achieve the gain in surplus from the benchmark contract, the manipulator loses by trading in the market to move the benchmark down, and the background traders gain from this loss. However, the manipulator subsidizes the background traders to take positions that are not in their own interest absent the subsidy. In other words, the allocation in the

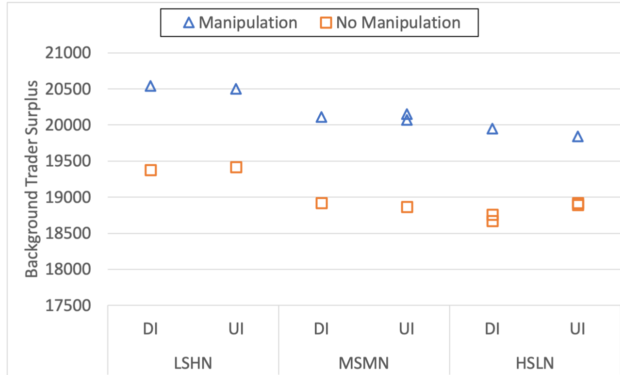

Figure 2. The aggregate total surplus of the ten background agents with and without manipulation. The total and market surplus is the same for background agents. The x-axis represents different market environments with varying fundamental shock, observation variance, and benchmark impact.

manipulated market is less efficient than in the manipulated one, which is why the aggregated market surplus decreases. Also, in equilibrium, this manipulation necessarily impacts the benchmark enough that the manipulator's gain from the benchmark exceeds its losses from trading in the market. The background traders gain the manipulator's loss from the market, so the manipulator's resulting gain from the manipulation exceeds that of the background traders.

Our toy example only considers how benchmark manipulation affects a few parties directly involved in the manipulation, benchmark contract, and calculation of the benchmark. This reduced problem does not take into account the broader equilibrium effects of benchmark manipulation and potentially the ultimate harm of this manipulation. Manipulation reduces the usefulness of the benchmark in refining financial exposures. There also exists potential negative effects of mispricing assets in the real and financial economy. For example, a mispriced LIBOR rate could result in banks lending to customers at rates that do not appropriately reflect their risk.

6. Policy Analysis

Following the LIBOR scandal, regulators investigated other benchmarks that had allegedly been manipulated and imposed some of the largest penalties ever paid by financial institutions. For instance, US and UK regulators fined six banks over $5 billion for manipulation of foreign exchange rates (Chon, 2015). Given the important role of benchmarks as financial infrastructure, regulators also turned to potential policy measures to deter manipulation. The International Organization of Securities Commissions published its Principles for Financial Benchmarks (IOSCO, 2013) and the

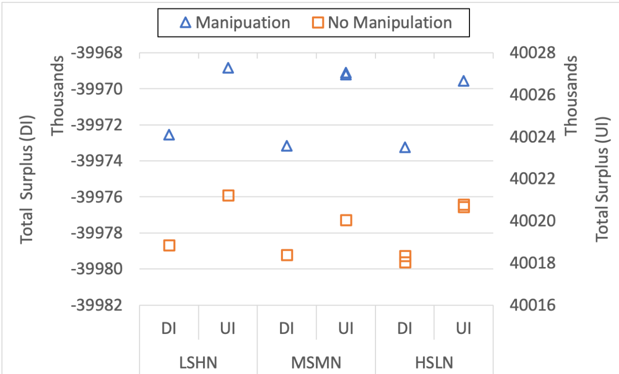

(a) Aggregate total surplus of all agents with and without manipulation. The primary y-axis shows total surplus when the benchmark manipulator tries to shift the benchmark down (**DI**); the secondary y-axis shows the surplus when the manipulator tries to shift the benchmark up (**UI**).

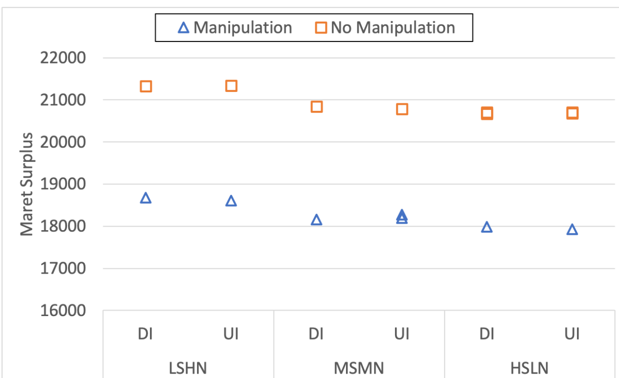

(b) Aggregate market surplus of all agents with and without manipulation.

Figure 3. With both figures, the x-axis represents different market environments with varying fundamental shock, observation variance, and benchmark impact.

European Union adopted its Benchmarks Regulation. Both documents stress the governance obligations of benchmark administrators, the quality of benchmark data, and most relevantly, robust methodological design of benchmarks. Nonetheless, regulators have neither suggested, nor mandated benchmark design features at a microstructure level of granularity. International regulators' interest in developing best practices for benchmark methodology means there should be substantial interest in results along the lines developed here.

The role of regulation is also important because we should expect markets to fail to produce optimal benchmarks themselves. In general, index providers do not operate in fully competitive markets or internalize anything like the full costs and benefits of the indices they produce. There are several reasons for this: First, indices are subject to familiar network effects that can cause suboptimal indices to gain a significant degree of lock-in and thus render the index provider a quasi-monopolist. Second, many widely used benchmarks are produced as a side-effect of other economic activity and thus do not provide their administrators with a robust revenue stream, notwithstanding that the benchmark can have significant effects on the welfare of counterparties (Rauterberg & Verstein, 2013). To illustrate, LIBOR originally arose to serve as a reference rate for banks' own lending activities, but came to play a pivotal role in the enormous interest rate derivatives market, without generating any direct revenue for the LIBOR panel banks. As a result of these forces, administrators' private incentives to ensure optimal benchmark design are frequently weaker than what would be socially desirable.

7. Conclusion

We analyze agents in a simulated CDA market where one security is traded to determine the impact of financial benchmark manipulation on market microstructure and welfare. The financial benchmark is calculated at the conclusion of a market by taking the volume weighted average price of all transactions executed during the market period. We simulated one benchmark manipulator who wishes to shift the benchmark calculation to increase its surplus from the benchmark, and multiple background traders who are indifferent to the benchmark. We find that the surplus of both the benchmark manipulator and background traders increases. Thus, background traders have no incentives to help mitigate this type of manipulation. Though the aggregate surplus of the market participants increases, the aggregate total surplus generated from the market decreases when the benchmark is manipulated, so the general welfare of the market decreases. Given that the surplus of all market participants increases, third parties invested in the opposite direction of the manipulator bear a large portion of the manipulation costs.

One limitation of this study is that we explore only one strategy type for the manipulator. The single strategy type is also based on a ZI trading strategy, so it is not necessarily the most informed or effective manipulation strategy. It would be interesting to develop and analyzing a more sophisticated benchmark manipulator that took market state into account when submitting orders to maximize the impact of its manipulation.

Another limitation of this study is that we consider only one simple benchmark calculation. Many financial benchmarks used to estimate asset values in real financial markets are derived from far more complex calculations. It is possible that these benchmarks are more (or less) difficult to manipulate than VWAP, so it would be worthwhile to explore the possibilities of manipulating benchmarks similar to them.

Lastly, some others limitations might affect the outcome of our equilibrium analysis. We try to mitigate the effects of sampling error by running a large number of simulations on the equilibria strategies we found. The game reduction method, DPR, we use to find equilibria has been shown to find good approximations of reduced games, but these are still only an estimation. It could be beneficial to run analysis on the entire strategy space, though this greatly increases the run-time to find equilibria. Even with these limitations, this study provides several insights into how financial benchmark manipulation can impact market microstructure.

References

Banerji, G. Regulator looks into alleged manipulation of VIX, Wall Street's 'fear index'. *Wall Street Journal*, 2018.

Bariviera, A. F., Guercio, B., Martinez, L. B., and Rosso, O. A. Libor at crossroads: Stochastic switching detection using information theory quantifiers. *Chaos, Solitons & Fractals*, 88:172–182, 2016.

Bond, P., Edmans, A., and Goldstein, I. The real effects of financial markets. *Annual Review of Financial Economics*, 4(1):339–360, 2012.

Boyle, C. Forex manipulation: How it worked, 2014. CNBC.

Brinkman, E. and Wellman, M. P. Empirical mechanism design for optimizing clearing interval in frequent call markets. In *18th ACM Conference on Economics and Computation*, pp. 205–221, 2017.

Chon, G. Six banks fined $5.6bn over rigging of foreign exchange markets, 2015. Financial Times.

Duffie, D. Compression auctions with an application to LIBOR-SOFR swap conversion. Working Paper 3727, Stanford Graduate School of Business, 2018.

Duffie, D. and Dworczak, P. Robust benchmark design. (3175), 2018.

Duffie, D. and Stein, J. C. Reforming LIBOR and other financial benchmarks. *Journal of Economic Perspectives*, 29(2):191–212, 2015.

Eisl, A., Jankowitsch, R., and Subrahmanyam, M. G. The manipulation potential of Libor and Euribor. *European Financial Management*, 23:604–647, 2017.

Farmer, J. D., Patelli, P., and Zovko, I. I. The predictive power of zero intelligence in financial markets. *Proceedings of the National Academy of Sciences*, 102:2254–2259, 2005.

Gellasch, T. and Nagy, C. Benchmark-linked investments: Managing risks and conflicts of interest, 2019. Healthy Markets Association.

Gode, D. K. and Sunder, S. Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy*, 101:119–137, 1993.

ICE Benchmark Administration Limited. U.S. Dollar ICE Bank Yield Index, 2019. Intercontinental Exchange.

IOSCO. Principles for financial benchmarks, 2013. The Board of the International Organization of Securities Commissions.

McBride, J. Understanding the Libor Scandal, 2016. Council on Foreign Relations.

Rauch, B., Goettsche, M., and El Mouaaouy, F. LIBOR manipulation: Empirical analysis of financial market benchmarks using Benford's law. *SSRN Electronic Journal*, 2013.

Rauterberg, G. and Verstein, A. Index theory: The law, promise and failure of financial indices. *Yale Journal on Regulation*, 30:1, 2013.

Verstein, A. Benchmark manipulation. *Boston College Law Review*, 56:215–272, 2015.

Wah, E., Lahaie, S., and Pennock, D. M. An empirical game-theoretic analysis of price discovery in prediction markets. In *26th International Joint Conference on Artificial Intelligence*, pp. 510–516, 2016.

Wah, E., Wright, M., and Wellman, M. P. Welfare effects of market making in continuous double auctions. *Journal of Artificial Intelligence Research*, 59:613–650, 2017.

Wang, X., Vorobeychik, Y., and Wellman, M. P. A cloaking mechanism to mitigate market manipulation. In *27th International Joint Conference on Artificial Intelligence*, pp. 541–517, 2018.

Wellman, M. P. Putting the agent in agent-based modeling. *Autonomous Agents and Multi-Agent Systems*, 30:1175–1189, 2016.

Wiedenbeck, B. and Wellman, M. P. Scaling simulation-based game analysis through deviation-preserving reduction. In *11th International Conference on Autonomous Agents and Multiagent Systems*, pp. 931–938, 2012.