Robust Benchmark Design
Online Appendix

Darrell Duffie[⋆] *and* *Piotr Dworczak*[†]

Abstract

This is the online appendix of Robust Benchmark Design, Duffie and Dworczak (2021). Appendix OA.1 discusses non-linear fixings, Appendix OA.2 presents two models of manipulation that micro-found our reduced-form baseline framework, while Appendix OA.3 contains additional examples.

Keywords: financial benchmarks, manipulation, collusion, mechanism design without transfers, volume-weighted average price

JEL codes: G12, G14, G18, G21, G23, D82.

[⋆]Stanford University, Graduate School of Business, and National Bureau of Economic Research. *Email address:* duffie@stanford.edu.

[†]Northwestern University, Department of Economics. *Email address:* piotr.dworczak@northwestern.edu.

OA.1 Non-linear fixings

In this appendix, we briefly address the issue of non-linear fixings, focusing on the most famous example: the median-based fixing. While comparing the statistical efficiency of the optimal linear benchmark to median-based fixings is beyond the scope of this paper, we use our framework to argue that the potential benefit of such benchmarks – reducing the incentives to manipulate (see for example Eisl *et al.*, 2014, and Youle, 2014) – might actually be lower than expected.

First, assuming n is odd for simplicity, define the *median benchmark fixing* $\hat{Y}^{MED}$ as

$$\hat{Y}^{MED} = \hat{X}^{(\frac{n+1}{2})},$$

where $\hat{X}^{(k)}$ is the k-th order statistic of $\hat{X}_1, \ldots, \hat{X}_n$.

Claim OA.1 *Under the benchmark $\hat{Y}^{MED}$, every agent chooses to manipulate.*

The proof is straightforward: An agent can have a non-negligible influence on the benchmark by submitting an arbitrarily small (thus arbitrarily cheap) transaction with a maximally distorted price. While a heavily distorted price is unlikely to be equal to the benchmark, there is a positive probability (bounded away from zero) that it will affect which price becomes the median observation (relative to not manipulating). Thus, even for very small exposure R_i, agent i would want to manipulate in this way.

Of course, this undesirable property of the median benchmark, as defined above, is due to its insensitivity to transaction sizes. The logic behind Claim OA.1 suggests that we should instead consider the *volume-weighted* median benchmark fixing defined by

$$\hat{Y}^{VMED} = \hat{X}^{(j)},$$

where

$$j = \min\left\{k : \sum_{j=1}^{k} \hat{s}^{(j)} \geq \frac{1}{2}\sum_{i=1}^{n} \hat{s}_i\right\},$$

with $\{(\hat{X}^{(j)}, \hat{s}^{(j)})\}_{j=1}^{n}$ ordered by decreasing prices.

Claim OA.2 *Let $p_n = \mathbb{P}\left(\frac{1}{n-1}\sum_{j=1}^{n-1} s_i \leq \mathbb{E}[s_1]\right)$. Then, a necessary condition for truthful reporting to be an equilibrium under $\hat{Y}^{VMED}$ is that*

$$\bar{R} \leq \frac{n-1}{p_n}\gamma\mathbb{E}[s_1].$$

Recall from Proposition 2 that a sufficient condition for truthful reporting under a linear benchmark is $\bar{R} \leq n\gamma\mathbb{E}[s_1]$. Because p_n will typically be close to $1/2$, Claim OA.2 shows that truthful reporting under the volume-weighted median benchmark requires parameter values that are of the same order of magnitude, and can in fact be more difficult to achieve if the sample size is small and the distribution g of unmanipulated transaction sizes has a thick right tail.

To prove Claim OA.2, suppose that everyone but agent i reports truthfully. If agent i's exposure is $\bar{R}$, one possible manipulation is to submit a transaction with distortion z_i and size $s_i = (n-1)E[s_1]$. Then, there is probability at least p_n that the fixing will be equal to agent i's transaction price, so agent i's expected influence on the benchmark is at least $p_n z_i$. The inequality from Claim OA.2 says that the benefit from manipulating, $\bar{R}p_n z_i$ is smaller than the cost of manipulation, $\gamma(n-1)E[s_1]z_i$.

The necessary condition from Claim OA.2 is not sufficient because *(i)* the simple calculation in the proof ignores the fact that the manipulation can also influence the benchmark indirectly, even when agent i's price does not become the benchmark, and *(ii)* there could be other, more subtle, manipulation strategies. Therefore, the scope of the median benchmark to reduce manipulations is quite limited.

Many benchmarks, including Libor, lie in between the volume-weighted average price and the volume-weighted median price by excluding only some "outlier" prices. An even more sophisticated approach would be to compute, for every transaction, the posterior probability that the transaction is manipulated, and to use this information to construct weights. We leave these directions for future research.

OA.2 Models of manipulation

This section presents two stylized models of trading and manipulation that give rise to the functional forms for costs and incentives assumed in Section 2. Apart from providing a microeconomic foundation for our assumptions, these models give more precise meanings to some model parameters.

OA.2.1 Committed quotes and costly search

We first consider a framework in which manipulation is costly because agents are committed to offering execution at the price quotes they submit to the benchmark administrator. In this framework, as is common in some actual benchmark settings, the submitting agents are dealers whose quotes are used to fix the benchmark. This was the case for the main industry benchmark for interest rate swaps known as ISDAFIX, whose manipulation[1] triggered more

[1]See "CFTC Orders The Royal Bank of Scotland to Pay $85 Million Penalty for Attempted Manipulation of U.S. Dollar ISDAFIX Benchmark Swap Rates."

than $600 million in fines for several dealers, Deutsche Bank, Goldman Sachs, Royal Bank of Scotland, Citibank, and Barclays, and to a more robust benchmark design, as outlined by Aquilina, Ibikunle, Mollica, Pirrone and Steffen (2018).

Manipulation consists in quoting a price that is an overestimate or underestimate of the true value of the asset to the dealer. If the values for the asset are highly correlated among market participants, then a mispriced quote is likely to be executed by a different investor, yielding a loss to the quoting bank. In an instance of manipulation of ISDAFIX by Deutsche Bank Securities Inc., the CFTC found[2] that "DBSI Swap traders would tell the Swaps Broker their need for a certain swap level at 11:00 a.m. or their need to have the level moved up or down. On at least one occasion, the Swaps Broker expressed the need to know how much 'ammo' certain DBSI traders had to use in order to move the screen at 11:00 a.m." The "ammo" presumably refers to losses that the DBSI would incur from trades at manipulated quotes.

The probability of an execution at a distorted quote depends both on the degree of distortion and on the transparency of the market. If quotes are public (as would be the case in a centralized limit order book), a significantly distorted quote would be executed with a probability close to one. If the market is more opaque or less active, and especially if quotes are revealed to traders only upon request (as in bilateral over-the-counter markets and on multilateral request-for-quote platforms), then the probability of incurring a loss by offering a distorted quote would be lower.

In our model, dealer i chooses $\hat{s}_i \in [0, \bar{s}]$ and $\hat{z}_i \in \{-\bar{z}, 0, \bar{z}\}$, for some constant $\bar{z} > 0$ which we could set to σ_z to match the notation from the baseline model. The variable X_i is interpreted as the actual per-unit value of the asset to dealer i. The dealer commits to trade up to $\hat{s}_i$ units at a price $\hat{X}_i = X_i + \hat{z}_i$, where the pair $(\hat{X}_i, \hat{s}_i)$ is used as a benchmark submission. For simplicity, we set the bid-ask spread to zero, that is, $\hat{X}_i$ is both a bid and an ask. We assume that Y has unbounded support, while ϵ_i has a symmetric distribution on an interval $[-\bar{\epsilon}, \bar{\epsilon}]$, for some $\bar{\epsilon} \leq \bar{z}/2$. This captures the idea that the distortion in prices due to manipulation is larger than the distortion due to idiosyncratic differences in the value of the asset to different traders.

We adopt a stylized search protocol to determine the probability that a committed quote is executed. Before observing its manipulation incentive type R_i, dealer i chooses a search intensity $\lambda_i \in [0, 1]$, paying a cost $c(\lambda_i) = \frac{1}{2}\bar{c}\lambda_i^2$. Here, λ_i is the probability that the dealer will be allowed to trade at the committed quotes of some other (randomly chosen) dealer j. We assume that each dealer is contacted at most once.[3] Upon contacting j, dealer i maximizes

[2]See CFTC Orders Deutsche Bank Securities Inc. to Pay $70 Million Penalty for Attempted Manipulation of U.S. Dollar ISDAFIX Benchmark Swap Rates, CFTC, February 1, 2018.

[3]Formally, imagine the following iterative procedure. Dealer 1 contacts one of the in dealer

the value of its chosen transaction. Because X_i is the unit value of the asset to dealer i, the resulting payoff of dealer i is

$$\max\left\{\max_{s\le\hat{s}_j}\left(X_i-\hat{X}_j\right)s,\ \max_{s\le\hat{s}_j}\left(\hat{X}_j-X_i\right)s\right\}.$$

Here, dealer i buys or sells the maximum quantity $\hat{s}_j$ to which dealer j has committed, due to linearity in value. The difference between the value X_i and the quote $\hat{X}_j$ determines the direction of trade.

OA.2.1.1 Solution

We focus on symmetric Nash equilibria. Dealer i makes two choices, the search intensity λ_i and the manipulation levels $(\hat{z}_i,\ \hat{s}_i)$. Regarding the first choice, the expected payoff to a dealer conditional on a successful search depends on the probability that other banks choose to manipulate. If p_M denotes the equilibrium probability of manipulation, then that expected payoff is

$$\begin{aligned}\mathbb{E}\left((1-p_M)|X_i-X_j|+\frac{1}{2}p_M|X_i-\bar{z}-X_j|+\frac{1}{2}p_M|X_i+\bar{z}-X_j|\right)\mathbb{E}(\hat{s}_j)&\\ =\left[(1-p_M)\mathbb{E}\left(|\epsilon_i-\epsilon_j|\right)+p_M\bar{z}\right]\mathbb{E}(\hat{s}_j)\equiv\phi.&\end{aligned}$$

The optimal choice of search intensity is thus $\lambda^\star=\min\left\{1,\ \phi\bar{c}^{-1}\right\}.$

As for the choice of manipulation, the dealer can always guarantee a zero payoff by quoting a price equal to the true value X_i, regardless of the size $\hat{s}_i$, by choosing $\hat{z}_i=0$. On the other hand, choosing $\hat{z}_i\in\{-\bar{z},\ \bar{z}\}$ yields a payoff $-\bar{z}\hat{s}_i$ in the event of being contacted by another dealer. The probability of being contacted is

$$\sum_{k=1}^{n-1}\binom{n-1}{k}(\lambda^\star)^k(1-\lambda^\star)^{n-1-k}\frac{k}{n-1}=\lambda^\star.$$

Taking into account the payoff generated by influencing the benchmark, and normalizing the payoff from not manipulating to zero, we see that the payoff from choosing $(\hat{s}_i,\ \hat{z}_i)$ is equal to

$$\left(R_i f(\hat{s}_i)-\lambda^\star\hat{s}_i\right)\hat{z}_i$$

which is exactly the expression assumed in Section 2, when taking $\gamma=\lambda^\star$.

$\{1,...,n\}\setminus\{1\}$ with probability λ_1. If dealer 1 contacts dealer j, then dealer 2 contacts one of the dealers in $\{1,...,n\}\setminus\{2,\ j\}$ with probability λ_2, and so on.

OA.2.1.2 Discussion

Based on the simple model of the previous subsection, the parameter γ can be interpreted as the probability of execution of a manipulated quote. If trade takes place on an active limit order book, then it is natural to assume that the cost $\bar{c}$ of search is nearly zero, and hence that $\gamma = \lambda^\star$ is close to 1. That is, manipulation would almost always yield a trading loss. On the other hand, in an opaque over-the-counter markets, $\bar{c}$ may be relatively large, and hence manipulation is less costly – a manipulated quote might not always be executed. As a consequence, holding the benchmark fixed, the probability of manipulation is higher in an opaque market.

If $\lambda^\star$ is less than one, there is an additional feedback effect between the benchmark fixing and the probability of manipulation. The ex-ante probability p_M of manipulation by any dealer is $1 - H(R_f)$, which is the probability that the dealer's exposure type R exceeds the threshold R_f determined by the weighting function f used in the fixing. If f is changed to reduce manipulation, then R_f goes up and p_M goes down. This, however, implies that the incentive to search is reduced, because the probability of encountering a profitable distorted quote gets smaller. As a consequence, $\lambda^\star$ decreases, and manipulation becomes cheaper. In a sense, the benchmark fixing and the market forces act as substitutes in preventing manipulation when the market is relatively opaque.

This discussion suggests that moving from a centralized to an opaque market may have an ambiguous influence on the shape of the optimal benchmark. On one hand, because a given manipulation of the price is less costly in an opaque market, the fixing that should be chosen in an opaque market would place a relatively smaller weight on small transactions. On the other hand, a fixing that deters manipulation lowers the cost of manipulating through the equilibrium effect on the search intensity of other market participants.

OA.2.2 An auction model

In this subsection, we consider an alternative trading model. When a liquidity shock hits a dealer, it may request quotes from other dealers, as is typical on electronic request-for-quote (RFQ) platforms. We model this as a sealed-bid auction. Absent incentives to manipulate, the dealer will accept the most attractive quote, for example, the lowest ask when it needs to buy the asset. The execution price, along with the corresponding trade volume, is then used to calculate the benchmark fixing. If, however, the dealer wants to inflate the fixing in order to take advantage of a long position in benchmark-linked assets, the dealer has an incentive to trade at the highest ask offered in the auction. This induces a tradeoff between the loss incurred in the auction and the gain associated with distorting the benchmark fixing.

We build a stylized model that aims to capture the main incentives. Dealer i is hit by a liquidity shock δ_i that takes one of the values $\{-\Delta, \Delta\}$ with equal probability, for some $\Delta > 0$. Dealer i then values each unit of the asset at $Y + \delta_i$, for quantities up to s_i. Whenever a dealer is hit by a shock, it requests quotes from two other dealers who have access to an unlimited supply of the asset at the common-value price Y. (The restriction to only two other dealers is not essential for the qualitative results but will yield explicit analytic solutions.) We model the competition between the two quoting dealers as a first-price auction (Bertrand competition). Absent incentives to manipulate, the dealer requesting the quote chooses the more attractive of the quotes, and thus Bertrand forces push the price to Y. However, when the quote-requesting dealer is a manipulator, it chooses the least attractive of the quotes, creating an incentive for dealers to provide quotes further away from the value Y.

OA.2.2.1 Solution

For concreteness, consider the case in which dealer i requests quotes to buy the asset (the opposite case is symmetric). Let p_M be the equilibrium probability that dealer i manipulates by accepting the higher of the quotes, corresponding to the case of a positive exposure R_i. In the unique symmetric equilibrium of the auction, the two dealers that provide quotes randomize their offers according to a continuous distribution function F with support $[Y + \lambda\Delta, Y + \Delta]$, where λ is determined in equilibrium. Following the line of argument in Stahl (1989), this requires each of the two dealers to be indifferent between all per-unit quotes q in the support of F, so that

$$[(1 - p_M)(1 - F(q)) + p_M F(q)]\,(q - Y) = p_M \Delta.$$

Solving, we obtain

$$F(q) = 1 - \frac{p_M}{1 - 2p_M}\frac{Y + \Delta - q}{q - Y},$$

which is a well defined cdf when $p_M < 1/2$. Moreover, we have $\lambda = p_M/(1 - p_M)$. If p_M is small, the quotes are close to Y. When p_M is relatively high (but below 1/2), the quotes are close[4] to $Y + \Delta$. With the above description, we can calculate equilibrium payoffs, and the distribution of transaction data. Let ε_i^k, for $k = 1, 2$, and $i = 1, 2, \ldots, n$, be the profit margin charged by dealer k in the auction requested by dealer i. That is, $Y + \varepsilon_i^1$ and $Y + \varepsilon_i^2$ are the quotes received by dealer i. Normalizing the payoff from not manipulating to zero, we take the cost of manipulation to be equal to the extra profit margin conceded by dealer i through choosing the less attractive quote for $\hat{s}_i$ units of the asset. This concession is $\hat{s}_i \mathbb{E}\left[\max\{\varepsilon_i^1, \varepsilon_i^2\} - \min\{\varepsilon_i^1, \varepsilon_i^2\}\right]$. Taking into account the benefit from influencing the fixing,

[4] We leave out a description of the equilibrium for the case $p_M \geq 1/2$ which is less relevant for our application. In that case, we would observe bids above $Y + \Delta$.

the net expected payoff from manipulation is equal to

$$(R_i f(\hat{s}_i) - \hat{s}_i)\, \mathbb{E} z_i,$$

where $z_i = \left[\max\{\varepsilon_i^1, \varepsilon_i^2\} - \min\{\varepsilon_i^1, \varepsilon_i^2\}\right]$. This setting can therefore be viewed as a version of our basic model for the case $\gamma = 1$.

OA.2.2.2 Discussion

The model of this section endogenizes the noise structure assumed in Section 2. The noise term ϵ_i reflects the dispersion in bids and asks quoted in the auction requested by dealer i. Manipulated transactions are more noisy than unmanipulated transactions because the worst price is further away from the mean Y than the best price. The noise term ϵ_i is $\pm \min\{\varepsilon_i^1, \varepsilon_i^2\}$, with symmetric probability. Manipulated transactions contain an additional noise term $z_i = \max\{\varepsilon_i^1, \varepsilon_i^2\} - \min\{\varepsilon_i^1, \varepsilon_i^2\}$. Thus, we provided a game-theoretic foundation for our assumption that manipulation reduces the signal-to-noise ratio of a benchmark.

In the framework modeled in this section, there is an additional distortionary channel for manipulation, through its impact on the probability distribution of unmanipulated data. When it is more likely that a counterparty in a transaction is a manipulator, a trader might provide a noisy quote, hoping that it will be accepted when the price distortion happens to be of the sign preferred by the manipulator. As a result, even when the quote requester is not a manipulator, and would take the most attractive quote, the distribution of quotes is more dispersed. As the probability p_M of manipulation rises, the probability distribution F of quotes shifts towards quotes further away from the true value Y. Hence the variance of ϵ_i rises, in that $|\epsilon_i|$ is distributed according to the cdf $1 - (1 - F_\varepsilon(\epsilon))^2$, where

$$F_\varepsilon(\epsilon) = 1 - \frac{p_M}{1 - 2p_M} \frac{\Delta - \epsilon}{\epsilon},$$

implying that

$$\sigma_\epsilon^2 = 2\Delta^2 \left(\frac{p_M}{1 - 2p_M}\right)^2 \left[-\log\left(1 + \frac{2p_M - 1}{1 - p_M}\right) + \frac{2p_M - 1}{1 - p_M}\right].$$

The noise level σ_ϵ^2 is increasing in Δ and p_M. In particular, $\lim_{p_M \to 1/2} \sigma_\epsilon^2 = \Delta^2$.

In this auction setting, because manipulation adversely impacts the precision of unmanipulated price signals, the slope of the optimal benchmark weighting function f is lowered in order to mitigate the risk of manipulation. The benchmark designer can affect the distribution of ϵ_i by choosing f so that $p_M = 1 - H(R_f)$ is relatively low. As a result, the probability of manipulation is smaller than in the baseline model in which the distribution of unmanipulated transaction data is exogenous.

OA.3 Additional examples

In this appendix, we present two additional examples that illustrate the results in Section 4. The first example illustrates Proposition 3.

Example OA.1 *Under the parametric assumptions of Example 1, we can numerically compute the optimal manipulation threshold:* $R^\star \approx 2.58$ *achieves the minimum for the benchmark administrator's problem* $\mathcal{P}$*. We can then apply Theorem 1b to compute the optimal fixing. Figure OA.1 presents the optimal weighting function for* $R = 0.5$*,* $R = 2.58$*, and* $R = 5$*. The ex-ante probabilities of manipulation under these target levels are approximately* 0.78, 0.28, *and* 0.08*, respectively.*[5] *Figure OA.1 shows that it is possible for two feasible weighting functions to never cross. If the distribution of sizes* $\hat{s}_i$ *were fixed, this would clearly be impossible because any two such functions could not have the same expectation with respect to the distribution of* $\hat{s}_i$*. However, this is possible when the distribution of* $\hat{s}_i$ *depends on the shape of* f*.*

Fig. OA.1: Optimal weighting functions for Example OA.1

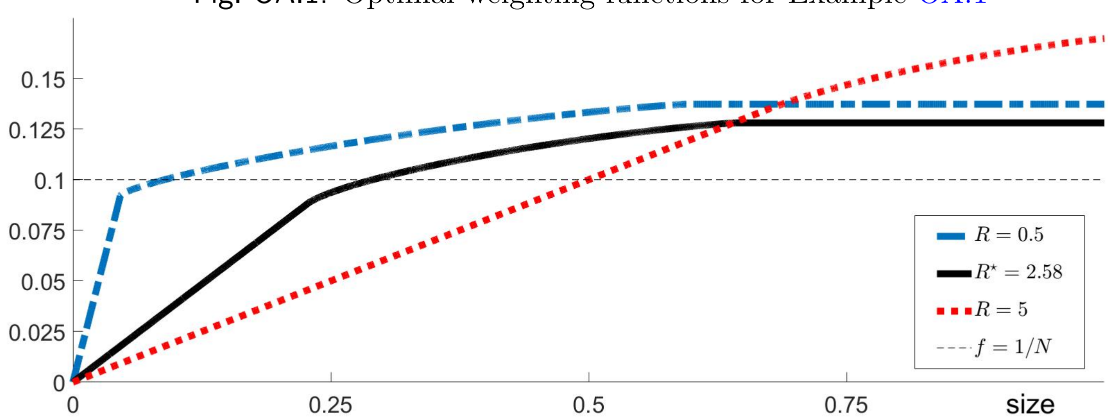

The second example is an illustration of Theorem 2.

Example OA.2 *We adopt the parameters of Example 1 and of Example OA.1.*[6] *The optimal benchmark fixing in the baseline model leads to the threshold* $R^\star \approx 2.58$ *which induces*

[5] Although Figure OA.1 may suggest otherwise, the function corresponding to $R = 5$ has a zero derivative at $s = \bar{s}$. The second derivative gets large close to $s = \bar{s}$, so the first derivative changes rapidly in a small neighborhood of $\bar{s}$. This is the case in which Theorem 1b does not apply and the solution is described by Theorem 1 in Appendix A.1.1.

[6] Because we solve the example numerically, all numerical results reported in this and other examples are approximate.

Fig. OA.2: Optimal weighting functions for Example OA.2

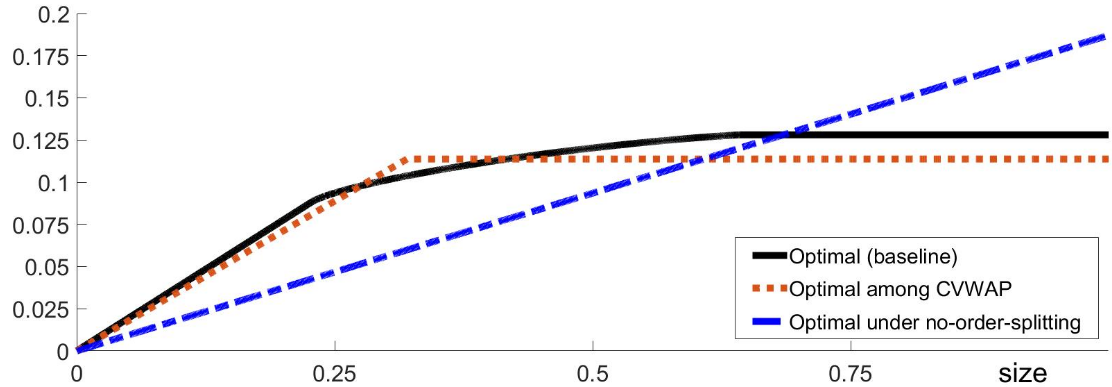

28% *of agents to manipulate. The optimal benchmark that deters order splitting is that which minimizes the probability of manipulation subject to unbiasedness. This yields a manipulation incentive threshold* $\hat{R}$ *of about 5.35, leading to manipulation with a probability of about 7%. The minimized objective function (mean squared error of the estimator) is* 0.142 *in the baseline case, and* 0.19 *when restricted to Condition 1, robustness to order splitting. This sharp increase in benchmark noise is caused by attaching a higher weight to manipulated transactions and inefficiently small weight to small unmanipulated transactions.*

To put this in context, consider the optimal benchmark fixing in the class of capped volume-weighted average price fixings, those with a weighting function that in linear in transaction size s *up to some maximal transaction size, after which the weight remains constant. The best such fixing has a mean squared error of* 0.149 *and induces manipulation by an agent in the event that the agent's manipulation incentive* R *exceeds* 2.81*, which has a probability of about* 24%*. These three weighting functions are depicted in Figure OA.2.*

References

Aquilina, M., Ibikunle, G., Mollica, V., Pirrone, A. and Steffen, T. (2018) The visible hand: Benchmarks, regulation and liquidity, Working paper, University of Edinburgh.

Duffie, D. and Dworczak, P. (2021) Robust benchmark design, *Journal of Financial Economics*.

Eisl, A., Jankowitsch, R. and Subrahmanyam, M. (2014) The manipulation potential of Libor and Euribor, Working paper, New York University.

Stahl, I. D. O. (1989) Oligopolistic pricing with sequential consumer search, *American Economic Review*, **79**, 700–712.

Youle, T. (2014) How much did manipulation distort the Libor?, Working paper, Univeristy of Minnesota.