Manipulation-Free Trading Mechanisms*

Giuliano Graziani† Stefano Lovo‡ Barbara Rindi§

March 26, 2026

Abstract

This paper studies how to design trading mechanisms when traders can manipulate market outcomes either by influencing the price they face or by shaping what other traders learn from market activity. A direct-mechanism benchmark isolates two design requirements: efficient trade requires an endogenous marginal trader, and manipulation-resistance requires individual price-taking. We then develop implementable mechanisms that satisfy these principles under different informational environments. In a private-values setting, a median-price mechanism is weakly dominant-strategy truthful, self-financing, and ex post efficient: the trader who can move the clearing price does not trade. In an interdependent-values setting, one-shot median pricing remains allocatively efficient but is informationally incomplete. We therefore propose a sequential mechanism in which bids are disclosed sequentially. This controlled disclosure improves information aggregation without restoring profitable manipulation. The paper provides a mechanism-design perspective on allocative efficiency, price discovery, and benchmark formation in financial markets, with applications to closing auctions, dark pools, and Limit Order Books.

Keywords: manipulation, spoofing, speed race, auction, interdependent values, dark pools, LOBs

*We thank seminar participants at HEC Paris. Barbara Rindi acknowledges funding from MUR- Prin 2022- Prot. 20227TCX5W, Next Generation EU. Giuliano Graziani acknowledges funding from MUR- Prin 2022- Prot. 20227TCX5W, Next Generation EU.

†Universidad Carlos III de Madrid, giuliano.graziani@uc3m.es

‡HEC Paris, lovo@hec.fr

§Bocconi University, IGIER and Baffi-Carefin, barbara.rindi@unibocconi.it

1 Introduction

Financial markets are designed to allocate assets to those who value them most and to produce prices that aggregate dispersed information about fundamentals (O'Hara, 2003). Market microstructure has made substantial progress in explaining how well these functions are performed in the presence of trading frictions, asymmetric information, and strategic interaction. The usual approach takes the trading protocol as given—a limit order book, a dealer market, a call auction, or a dark pool—and studies how equilibrium behavior unfolds within that institutional environment, or how targeted changes in rules and parameters affect market quality (Glosten and Milgrom, 1985; Kyle, 1985). This perspective is natural when the mechanism itself is treated as a primitive and the main question is how traders respond to it.

This approach becomes less informative when the mechanism is not merely the environment in which strategic behavior takes place, but the source of the distortion itself. The resulting inefficiencies are then not incidental departures from an otherwise well-functioning institution. They are equilibrium outcomes induced by the price-formation and information-release rules. This concern appears across several market mechanisms. In continuous limit order books, the ability to revise and cancel orders creates scope for speed races and for manipulative order activity that moves beliefs without committing to trade (Budish, Cramton, and Shim, 2015; Williams and Skrzypacz, 2020). In dark pools, executions are tied to prices formed in lit markets, creating latency-arbitrage opportunities and cross-venue benchmark-manipulation channels when reference prices are stale or can be strategically influenced (Zhu, 2014; Aquilina, Foley, O'Neill, and Ruf, 2024). In closing auctions investors have an incentive to influence the equilibrium price through strategically chosen orders and quantities (Comerton-Forde and Putniņš, 2011; Jegadeesh and Wu, 2022).

These environments differ in institutional details, but they share two common vulnerabilities. First, the trading rules may allow a participant to affect the price that governs his own allocation or that enters other economically relevant payoffs. Second, the rules may allow a participant to shape what other traders learn from market activity. The first vulnerability compromises price-taking behavior; the second compromises information revelation. Once these distortions are

embedded in the design of the trading protocol, the standard normative microstructure approach reaches its natural limit. Structural fragilities cannot be fully repaired by finer monitoring, greater transparency, or incremental adjustments to existing rules. The question becomes one of mechanism design: how should a market determine who trades, at what price, and with what degree of information disclosure so that trade remains allocatively efficient, prices remain informative, but manipulation is not privately profitable in equilibrium?

We therefore return to first principles. We propose a direct-mechanism benchmark that disciplines our analysis. The benchmark isolates two design principles. Efficient allocation requires an endogenous cutoff separating buyers from sellers: the marginal trader must be determined by the realized profile of values, and in our environment it corresponds to the trader with the median valuation. Incentive compatibility requires individual price-taking: no trader should be able to affect the price at which *he himself* trades. The benchmark delivers these principles cleanly, but it is not directly implementable as a market protocol because it can require the venue to know investors' valuation functions and need not be self-financing. The objective of this paper is to translate these principles into implementable trading rules that respect informational and budget balance constraints.

We begin with a private-information environment in which traders' valuations are not affected by what others know. Each participant holds a private valuation for the traded asset. In this setting, given a finite number of market participants the central problem is direct strategic influence over the trading price. In this setting, we propose the *Median Price Trading* (MPT) mechanism. Each trader submits a single price. Traders who submit prices above the median buy one unit from traders who submit prices below the median, all at the median submitted price; the median trader does not trade. The mechanism endogenously separates buyers from sellers and implements the first-best allocation. The key intuition is mechanism-design in spirit: the only trader who can move the clearing price is precisely the trader whose payoff does not depend on that price, because he does not trade. As a result, unilateral attempts to manipulate the price cannot increase profits. MPT thus implements allocative efficiency while shutting down the most direct price-influence channel.

We then move to a richer environment in which traders' valuations are interdependent, which

introduces the learning mechanism. Each trader's evaluation depends not only on his own signal, but also on other traders' signals, and signals are affiliated in the sense of Milgrom and Weber (1982, 2000). This environment is closer to real-world trading in which order flow and prices convey information. Here, allocative efficiency and price discovery are jointly at stake because the mechanism reallocates assets *and* aggregates information through what it reveals. The forces that motivate the paper become sharper: a trader is able to improve his profit, by influencing others' evaluations before trading.

In this interdependent-values setting, MPT remains strong on the allocation margin and retains the no-manipulation logic. However, MPT reveals only a single statistic—the median price—so most private information remains hidden. A venue can then be allocatively efficient and yet informationally thin, leaving traders exposed to ex-post regret once additional information is revealed (or lack of price discovery if information is not revealed). Improving information aggregation requires that information be released in a controlled and sequential manner, but sequentiality also creates a new danger: belief manipulation becomes tempting precisely because later traders learn from earlier actions.

This motivates our *Sequential Trading Mechanism* (STM). The first round mirrors MPT and endogenously determines who is a buyer and who is a seller. Trade then unfolds through a sequence of second-price buy and procurement auctions. After each round, the buyer and seller exit and their bids are disclosed, while transaction prices are not. The sequencing is the key design choice: a trader's action becomes informative to others only *after* it can no longer affect the terms at which that trader trades. It is not trivial to reconcile the manipulation-free property with information efficiency. The former requires that traders do not care about influencing other traders with their actions; the latter requires traders' actions to be visible to others so that the information implied by those actions can be used by other traders to better assess their valuation of the asset. The mechanism therefore separates information revelation from price manipulation. In equilibrium, traders with the most extreme signals trade earlier (when less information has been revealed) and traders closer to the margin trade later (after more information has been aggregated), generating an endogenous spread that compensates the venue for implementing incentive compatibility under learning.

Our proposed mechanisms clearly show a new trade-off. In one-shot settings, the price cannot be manipulated, there is a fast execution and price impact is zero at the cost, however, of little information. A slower sequential market releases more information and reduces ex-post regret, but it does so at the cost of a larger pre-trade price impact. The paper thus connects the design of auctions to a core market-quality question in finance: how should a venue balance depth, information aggregation, and exposure to manipulation when order flow itself may be informative? This is what we discuss in the last section where we connect our mechanisms to the existing market institutions. For dark pools, the MPT speaks to the problems created by external benchmark pricing, stale references, and cross-venue manipulation. For closing auctions, both mechanisms speak to finite-trader benchmark formation, strategic price influence, and the tension between efficient execution and informative closing prices. For continuous limit order books, they provide a more cautious but still suggestive lesson: when order flow is informative and strategic cancellations or speed races are first-order, auction-based mechanisms may dominate continuous trading, especially in thin markets or during periods of elevated uncertainty. The point is not that one venue should mechanically replace all others. It is that mechanism design can be used to identify which trading rules are structurally aligned with allocative efficiency and price discovery, and which trading rules invite traders to work against both.

Literature review. This paper contributes to three strands of literature. First, it contributes to the literature on manipulation and strategic distortion in financial markets. A large theoretical literature studies how traders can profit by moving prices or by affecting the beliefs of others through their market actions or public messages (Vila, 1989; Allen and Gorton, 1991; Benabou and Laroque, 1992; Van Bommel, 2003; Williams and Skrzypacz, 2020). Empirical work shows that these concerns are not confined to a single venue or asset class. They arise in closing-price manipulation, pump-and-dump schemes, spoofing episodes, and other forms of strategic trading that exploit the institutional details of price formation (Aggarwal and Wu, 2006; Comerton-Forde and Putniņš, 2011; Brogaard, Li, and Yang, 2022; Colla, Distaso, and Vitale, 2024; Tremacoldi-Rossi, 2024). Our contribution is different in emphasis. Rather than treating manipulation primarily as a surveillance or enforcement problem, we study how trading rules

can make manipulation unprofitable in equilibrium. The central object is therefore not a specific manipulative strategy, but the mechanism that gives that strategy value. In that sense, the paper provides general design lessons—and practical guardrails—for trading mechanisms more broadly: it identifies which features of a market's price-formation and information-release rules create vulnerability to strategic price influence.

Second, and following directly from that first contribution, the paper contributes to the debate on alternatives to continuous limit order book markets (Budish et al., 2015; Lehar and Parlour, 2021; Foley, O'Neill, and Putniņš, 2023; Malinova and Park, 2023; Budish, Cramton, Kyle, Lee, and Malec, 2023). Following Budish et al. (2015), a growing literature has studied whether periodic or auction-based trading mechanisms can improve on continuous trading by mitigating speed races and related frictions. Our results speak to that debate through a different mechanism-design channel. We show that auction-based trading can address vulnerabilities that go beyond the speed race itself, namely the ability of traders to exploit the price-formation process or the market's informational channel (Glosten, 1994). In an environment in which valuations depend on both private information and what traders learn from market outcomes, our mechanisms show how price formation can be internalized and how information disclosure can be disciplined so as to reduce manipulation incentives. At the same time, the paper shows that the fast–slow trade-off in sequential markets persists in this richer environment. Consistent with Du and Zhu (2017), slowing the market improves information aggregation but can increase per-round price impact. More broadly, the results suggest that auction-based mechanisms may be particularly valuable where the benefits of immediacy are limited and the risks of manipulation or informational distortion are acute.

Third, the paper contributes to the growing literature that brings mechanism design into market microstructure and exchange design (Du and Zhu, 2017; Andreyanov and Sadzik, 2021; Chen, 2025). Compared to these works, we adopt a classic environment in auction design (interdependent values) to study problems of adverse selection and hence potential market manipulation.[1]

[1]Quoting Du and Zhu (2017), "adverse selection" refers to situations in which different traders hold different pieces of information about the same asset; in our context, "adverse selection" can be read as "interdependent values." Accordingly, modeling interdependent values to study learning dynamics in financial markets is both natural and standard

Therefore, we complement this literature by focusing on an environment in which learning from others' signals is central. Overall, our paper contributes to the broader agenda of integrating market microstructure and mechanism design.

2 General setting and definitions

Let N be an odd integer.[2] In what follows, we use the following notation: given any vector $Y = \{y_1, \ldots, y_N\} \in \mathbb{R}^N$, let $Y_{-i} := \{y_1, y_2, \ldots, y_{i-1}, y_{i+1}, \ldots, y_N\} \in \mathbb{R}^{N-1}$, and for any given integer $l \in \{1, 2, \ldots, N\}$, we will denote with y^l the l-th lowest element of Y. Let $M = (N+1)/2$, so that y^M is the median element of Y, y^1 is the smallest one, and y^N is the largest one. So for example y_i refers to the variable for the generic trader i, whereas y^i is the i-th smallest entry of vector Y. We can now describe the fundamentals of the economy.

There is an odd number N of risk neutral traders. Each trader is interested in trading at most one unit of the asset.

Private signals and valuations Each trader i receives a private signal $x_i \in [\underline{x}, \overline{x}]$ about his valuation for the asset. Let $\mathcal{X} := [\underline{x}, \overline{x}]^N$ and let $X := \{x_1, x_2, \ldots, x_N\} \in \mathcal{X}$ denote a signal profile. Signals are drawn from a continuous distribution on $\mathcal{X}$ such that the probability of receiving two identical signals is nil. Trader i's valuation is denoted $\nu_i = v(x_i, X_{-i})$ where $v : \mathcal{X} \to \mathbb{R}$ is the valuation function, x_i is the trader i's private signal, and X_{-i} is the collection of the other traders' private signals. Given a signal profile $X \in \mathcal{X}$ we denote

$$\nu_i(X) := v(x_i, X_{-i}).$$

We make three assumptions on $v(\cdot)$:

A.1 A trader's valuation is strictly increasing in his signal:

$$\frac{\partial v(x_i, X_{-i})}{\partial x_i} > 0.$$

[2] Appendix D.1 characterizes the equilibria when N is even.

A.2 A trader's valuation is more sensitive to his own signal than to any other trader's signal:

$$\frac{\partial v(x_i, X_{-i})}{\partial x_i} \geq \frac{\partial v(x_i, X_{-i})}{\partial x_j}, \forall j \neq i.$$

A.3 A trader's valuation is symmetric in the other trader's signals in the sense that the order of the remaining signals does not impact the valuation.[3]

This valuation function setting is the same as in Milgrom and Weber (2000). Assumption A.1 that the traders' evaluation function is strictly increasing in their own signal implies a positive private value part, even if potentially small, in the evaluation of each trader. Therefore, while our model covers the pure private value setting, it rules out the common value framework. Yet, due to risk neutrality, the pure common value case implies no gains from trade and therefore the no-trade theorem of Milgrom and Stokey (1982).

Assumptions A.1-A.3 imply that $v(\cdot)$ satisfies the *private signal ordering* as defined below:[4]

Definition 1. *We say that* $v(\cdot)$ *satisfies the* private signal ordering *if for any* $X \in \mathcal{X}$ *and* i,j *such that* $x_i > x_j$, *one has* $\nu_i \geq \nu_j$.

As it is standard in the financial economics literature, we assume that each trader is interested in buying or selling at most one unit of the asset (e.g., Budish et al. (2015), Ernst, Spatt, and Sun (2022)). So we can define a generic trading profile as follows.

Trading profiles We define a trading profile as $\{(q_i, t_i)\}_{i=1}^N$, where $q_i \in \{-1, 0, 1\}$ specifies whether a trader i sells, does not trade or buys, the asset, and $t_i \in \mathbb{R}$ is a transfer of cash made to trader i, where a negative t_i implies a negative cash flow for trader i. We require that the

[3]For example, if there are three traders, $1, 2, 3$, the valuation of the trader 1 is such that $v(x_1, \{x_2, x_3\}) = v(x_1, \{x_3, x_2\})$.

[4]To see this take $x_i > x_j$. Note that $\nu_i = v(x_i, \{x_j, X_{-i,j}\}) = v(x_j, \{x_j, X_{-i,j}\}) + \int_{x_j}^{x_i} \frac{\partial v(z, \{x_j, X_{-i,j}\})}{\partial z} dz$, and similarly $\nu_j = v(x_j, \{x_i, X_{-i,j}\}) = v(x_j, \{x_j, X_{-i,j}\}) + \int_{x_j}^{x_i} \frac{\partial v(x_j, \{z, X_{-i,j}\})}{\partial z} dz$. Thus $\nu_i - \nu_j = \int_{x_j}^{x_i} \left(\frac{\partial v(z, \{x_j, X_{-i,j}\})}{\partial z} - \frac{\partial v(x_j, \{z, X_{-i,j}\})}{\partial z} \right) dz$ that is positive because for any $j \neq i$, $\partial v(x_i, X_{-i})/\partial x_i \geq \partial v(x_i, X_{-i})/\partial x_j$.

financial asset can neither be created nor destroyed, that is:

$$\sum_i q_i = 0.$$

We denote with $\mathcal{T}$ the set of all trading profiles. A trading profile $T \in \mathcal{T}$ is said to be *self-financed* if

$$\sum_i t_i \leq 0,$$

That is, the sum of the cash spent by traders is not inferior to the sum of cash received by traders. Thus a self-financed trading profile does not require cash inflow from a non-trader, say the market infrastructure, but the market infrastructure can receive some net cash from the traders.

We say that trading $T \in \mathcal{T}$ is *ex-post regret-free given* X, if after knowing all traders' signals X, no trader i would be strictly better-off withdrawing his trade. That is:

$$q_i \; v(x_i, X_{-i}) + t_i \geq 0, \forall i \tag{1}$$

Social Welfare Given this simple setting, if trader i sells to trader j at price p, then the value created by this trade, given X, is equal to

$$\nu_j - \nu_i$$

Note that the trading price p at which the trade occurs is irrelevant to determine the value created by a deal, as the transaction price p is a mere transfer from the buyer to the seller.

It is straightforward to see that trading profile that maximize ex-post social welfare must be ex-post allocation efficient as defined below.

Definition 2. *A trading profile is ex-post allocation efficient given* X*, if each trader with* $\nu_i > \nu^M$ *buys one unit of the asset from one trader with* $\nu_j < \nu^M$*, where* ν^M *is the median valuation. Because* $\sum_i q_i = 0$*, and* N *is an odd number, the trader with the median valuation* ν^M *does not trade.*

Thus, for a given realization of the private signals X, the ex-post maximum social welfare is equal to

$$W^{\star}(X) := \sum_{\nu_i(X)>\nu^M(X)} \nu_i(X) - \sum_{\nu_j(X)<\nu^M(X)} \nu_j(X) \geq 0 \tag{2}$$

Direct mechanism We define a direct mechanism as a mapping from the set of possible private signals profiles $\mathcal{X}$ into the set of trading profiles $\mathcal{T}$. Namely a direct mechanism associates to any given signal profile $X \in \mathcal{X}$ and each trader i a transaction $q_i(X) \in \{-1, 0, 1\}$ such that $\sum_i q_i(X) = 0$, and a monetary transfer $t_i(X)$. As standard since Myerson (1979), the mechanism function as follows: each trader privately observes his private signal, and is asked to report a signal to the direct mechanism. The mechanism then realizes a transaction profile based on what traders reported. An incentive compatible direct mechanism is such that each trader has an incentive to report his true signal. Furthermore the direct mechanism is individually rational if no trader expects to lose from the resulting trading profile. We focus on *ex-post mechanisms*, that guarantee that no matter what a trader believes about the other trader's signals, each trader i) reports his true signal and ii) he is sure not to lose from the resulting transaction.

Formally, we say that trading mechanism is ex-post incentive compatible (ex-post IC, henceforth) if for each trader i, any given signal profile of the other traders X_{-i},

$$q_i(x_i, X_{-i})v(x_i, X_{-i}) + t_i(x_i, X_{-i}) \geq q_i(x_i', X_{-i})v(x_i, X_{-i}) + t_i(x_i', X_{-i}), \text{ and } \forall x_i, x_i' \in [\underline{x}, \overline{x}] \tag{3}$$

This condition guarantees that each trader wants to truthfully report his signal to the mechanism no matter what he believes the realization of the other traders' signals to be.

We say that the trading mechanism is ex-post individually rational (ex-post IR, henceforth) if for each trader i

$$q_i(X)v(x_i, X_{-i}) + t_i(X) \geq 0, \text{ and } \forall X \in \mathcal{X} \tag{4}$$

That is, for any given realization of $X \in \mathcal{X}$, the trading profile induced by the trading mechanism is ex-post regret-free, in the sense of (1), given X.

2.1 The direct mechanism benchmark

In this section we propose an ex-post mechanism that always leads to ex-post efficient allocations. Fix any trader i and a signal profile $X_{-i} \in \mathcal{X}_{-i}$ for all other traders. In order to ensure the ex-post IC constraints (3) for trader i, the transfer to trade i must satisfy

$$t_i(x_i, X_{-i}) = t_i(x_i', X_{-i}),\ \forall x_i, x_i' \in \mathcal{X}_i, \tag{5}$$

In words, a trading mechanism is ex-post incentive compatible only if a trader's monetary transfer related to a specific trading position does not depend on what he reports to the mechanism, although it can depend on what the other traders report, i.e., on X_{-i}. Thus, rewriting the transfer to trader i as only function of what the other traders report, the ex-post IC condition (3) becomes:

$$q_i(x_i, X_{-i})\nu_i + t_i(X_{-i}) \geq q_i(x_i', X_{-i})\nu_i + t_i(X_{-i}) \tag{6}$$

To achieve an ex-post efficient allocation it must be that for any given realization X, all $(N-1)/2$ traders whose valuation $\nu_i > \nu^M$ purchase from the $(N-1)/2$ whose valuation $\nu_i < \nu^M$. Hence we have that a trading direct mechanism leads to an ex-post efficient allocation if and only if

$$q_i(X) = \begin{cases} -1 \text{ if } \nu_i < \nu^M \\ 0 \text{ if } \nu_i = \nu^M \\ 1 \text{ if } \nu_i > \nu^M \end{cases} \tag{7}$$

Now consider any ex-post mechanism leading to an ex-post efficient allocation. If for signals profile realization X, trader i buys, it must have valuation $\nu_i > \nu^M$. Because $\nu_i = v(x_i, X_{-i})$ and the function $v(\cdot)$ satisfies the private signal ordering, the trader i also has a private signal x_i strictly larger than the median private signal x^M. Thus for a given signal profile X with median x^M, each trader i whose $x_i > x^M$ has $q_i(X) = 1$. Taking into account ex-post IR condition (5), one has

$$v(x_i, X_{-i}) + t_i(X_{-i}) \geq 0\ \forall x_i \geq x^M$$

As $v(\cdot)$ satisfies the private signal ordering, necessary and sufficient condition for the above IR constraint to be satisfied for all $x_i \geq x^M$, is

$$t_i(X) \geq -v(x^M, X_{-i}) \tag{8}$$

By symmetry, it follows that an ex-post mechanism that leads to an ex-post efficient allocation must offer to a trader i whose $x_i < x^M$ to sell the asset for a price $t_i(X) \geq v(x^M, X_{-i})$.

Thus we have

Proposition 1. *A direct mechanism where q_i satisfies*

$$(q_i(X), t_i(X)) = \begin{cases} (-1, v(x^M, X_{-i})) \text{ if } x_i < x^M \\ (0, 0) \text{ if } x_i = x^M \\ (1, -v(x^M, X_{-i})) \text{ if } x_i > x^M \end{cases} \tag{9}$$

is an ex-post mechanism leading to an ex-post efficient allocation.

The proof of the proposition is in Appendix A.1. The proposition should be read as an allocation characterization: ex post efficiency holds if and only if agents with valuations above (below) the median are assigned to buy (sell), with the median agent not trading. Intuitively, this median-cutoff rule maximizes realized surplus by assigning trade to the agents with the highest gains from trade subject to feasibility. Moreover, ex post incentive compatibility requires transfers to be independent of an agent's own report; this condition is necessary but not sufficient on its own, since agents could still gain by misreporting to cross the median cutoff and change roles. Incentive compatibility is achieved by designing payments so that the median type is the marginal (indifferent) trader, i.e., no type benefits from deviating across the cutoff.

The proposed direct mechanism, however, has two important limitations that undermine its direct implementability. First, it is assumed that the mechanism knows the evaluation function $v : \mathcal{X} \to \mathbb{R}$ of market participants: traders submit their own signals x_i and the mechanism defines a specific price $t_i(X') = v(x'^M, X_{-i})$ for each trader. Second, the proposed mechanism is not necessarily self-financed. To see this consider the following example with $N = 3$ and suppose

X is such that $x_1 < x_2 < x_3$. Then $x^M = x_2$, trader 1 sells for a price of $t_1(X) = v(x_2, (x_2, x_3))$ whereas trader 3 buys for $t_3(X) = -v(x_2, (x_2, x_1))$. Now, it is natural to assume that valuation function $v(\cdot)$ is non-decreasing in all its arguments, then $t_1(X) \geq -t_3(X)$ implying that the seller sells for a price that is larger or equal than what the buyer pays. This result is reminiscent of the Myerson and Satterthwaite (1983) theorem which shows that, in a bilateral framework, it cannot exist an ex-post efficient budget-compatible mechanism.[5]

Proposition 1 isolates the basic difficulty of benchmark price formation in financial markets. Efficiency requires the venue to trade against the marginal valuation, but manipulation-proofness requires that this marginal valuation not be chosen by a trader who benefits from moving it. The direct mechanism achieves the right allocation, but at the cost of informational demands and possible budget imbalance that make it implausible as a real trading protocol. The remaining sections ask whether actual market rules can replicate these same design principles without requiring the venue to know investors' payoffs or inject capital into the market. We first consider the baseline private value framework and show a simple implementation of the direct mechanism that we call the median price trading rule and is ex-post IC, ex-post IR and self-financed. Then, we move to the case of interdependent valuations where we study the tradeoff between ex-post regret and trading volume.

3 Median Price Trading Protocol with Private Values

Consider a simple setting where each trader perfectly knows his own valuation for the asset. Namely we assume $v(x_i, X_{-i}) = x_i$.

The first best social welfare in (2) boils down to

$$W^\star = \sum_{x_i > x^M} x_i - \sum_{x_i < x^M} x_i \geq 0 \tag{10}$$

[5]By contrast, if traders' valuations are weakly decreasing functions of the other trader's signal, the proposed ex-post mechanism is also self-financed. However, assuming valuations decrease in other traders' signals is difficult to reconcile with the assumption that valuations are strictly increasing in own signals.

Below we describe a simple trading mechanism that achieves $W^\star$.

Median Price Trading (MPT) protocol Trading occurs as follows.

First, each trader i submits a price $p_i \in \mathbb{R}$ to the trading platform. Prices are submitted simultaneously, that is, when trader i submits his price, he does not observe the other traders' prices.

Second, the trading platform orders prices from the highest to the lowest and determines the median price p^M so that there are $(N-1)/2$ traders who submitted a price above p^M and $(N-1)/2$ traders who submitted a price below p^M.

Third, any given trader i buys for p^M if he submitted $p_i > p^M$, sells for p^M if he submitted $p_i < p^M$, and does not trade if he submitted $p_i = p^M$. Therefore, trader i's realized payoff is

$$\begin{cases} p^M - x_i & \text{if } p_i < p^M \\ 0 & \text{if } p_i = p^M \\ x_i - p^M & \text{if } p_i > p^M \end{cases} \tag{11}$$

Proposition 2. *Within the private value setting, in the Median Price Trading mechanism each trader i submitting $p_i = x_i$ is an equilibrium in a weakly dominant strategy. Transaction price is $p^M = x^M$. This trading mechanism is* ex-post efficient, ex-post incentive compatible, ex-post individually rational *and* ex-post budget balance compatible. *Moreover, the mechanism attains the first best social optimum and hence maximizes trading volume.*

The proof of the proposition is in Appendix B.1. As truthful reporting is an ex-post equilibrium in the MPT implies that each agent's belief about other agents' valuations plays no role in determining the equilibrium play and outcome. For this reason, we did not make any assumption on traders private valuations distributions.[6] Most importantly, in this pure private valuation setting the mechanism leads to a regret-free outcome. That is, because each trader already knows his own valuation for the asset (x_i), there is no room for learning from the bids

[6]We only assume that valuations are drawn from a continuous distribution. This reflects the idea that the probability of two traders having the exact same valuation for the asset is nil, and making thus tie-breaking rules irrelevant.

of the other traders; therefore, as long as he does not pay more than x_i for the asset or does not sell the asset for less than x_i, he has no regret when trading. As we show in Section 4, this property is not satisfied when traders' valuations are interdependent.

For double auctions, where buyers' and sellers' roles are predefined, Rustichini, Satterthwaite, and Williams (1994) show that the Myerson and Satterthwaite (1983) theorem does not hold if the number of agents tends to infinity - the price impact of each trader is close to zero. Proposition 1 shows that when market participants do not have predefined roles as buyers or seller in the mechanism, it is possible to design a mechanism that is efficient, incentive compatible, individually rational, and budget balance compatible by having three or more traders active in the market. We thus show that absent predefined roles for market participants, the Myerson and Satterthwaite (1983) theorem does not hold even with just three traders.

It is important to compare the MPT mechanism and its properties with other trading mechanism. The MPT mechanism and standard Walrasian Double auction mechanism lead to the same outcome if i) there is an infinite number of traders, and ii) the population is perfectly divided between buyers and sellers (Reny and Perry (2006)). With atomistic traders, in virtue of the strictly increasing truthfully reporting bidding strategy $p_i = x_i$, the median signal is the price that clears the Walrasian auction when the mass of buyers equals the mass of sellers.

However both conditions i) and ii) are necessary for the equivalence between the outcomes of the two trading mechanism. Walrasian auction caps potential trading in case of inbalance on one side (buyers or sellers) of the market. By contrast, the median mechanism systematically maximizes aggregate gain from trade as it is more agnostic on the role of potential traders, capturing the idea that many market participants decide their trading direction according to market conditions. Also, having atomistic traders in the Walrasian Double auction is necessary to ensure that no trader can impact the equilibrium price through his order. With finitely many traders, each trader has incentive to misreport in the hope of tilting the price in his favor (buyers underbid and sellers overbid respectively), the worst-case cost of misreporting is to not be executed, leading to a nil payoff. In the MPT system there is no incentive to misreport or to manipulate the trading price. By changing his own price an individual can change the sign of

his trade, and potentially realize a negative payoff, but not the price at which he can buy or sell the asset. The median trader can change the trading price but, as he does not trade, he has no incentive to manipulate it. As in a second price auction, the trading price is not determined by the agents who eventually trade but by an agent that is eventually excluded from the transaction. This feature makes the MPT price manipulation free.[7]

A potential concern is that, unlike standard trading platforms where participants submit buy or sell orders, the MPT mechanism requires agents to submit only a price without specifying direction. However, in many financial markets participants are not exogenously committed to being buyers or sellers; rather, their trading direction is endogenous to their information and to prevailing prices. The MPT mechanism captures this feature directly: direction is determined by the ranking of submitted prices, so agents with sufficiently high valuations become buyers and those with sufficiently low valuations become sellers. Agents whose valuations are close to the clearing price may trade in either direction, but ex-post individual rationality guarantees that any realized trade yields non-negative surplus. Thus, the absence of exogenously specified trading roles is not a restriction but a deliberate modeling choice that reflects the endogenous direction of trade in many financial environments. Finally, Appendix D.2 develops a multi-quantity extension of the mechanism that allows traders to submit demand for multiple units. The extension improves the institutional realism of the framework while leaving the main economic conclusions unchanged in qualitative terms.

4 Interdependent Values

In the previous section we assumed that each trader's valuation coincides with his private signal. As a result, trader i knows his own valuation exactly, and any information about other traders' valuations that he may infer from the trading outcome is payoff-irrelevant. In this section, we

[7] McAfee (1992) proposes a direct mechanism in a private-values setting where truthful reporting is a dominant strategy, assuming other agents also report truthfully. Unlike our approach, his mechanism relies on (i) an exogenously given market structure with distinct buyers and sellers, and (ii) a crossing condition, requiring that buyers' and sellers' valuations intersect for trading to occur. As a result, allocative efficiency crucially depends on the distribution of agent types, its expected value is difficult to characterize analytically, and the realized number of trades may substantially deviate from the efficient benchmark.

turn to the more demanding case, and the one more closely aligned with financial markets, in which traders update their valuations from prices and order flow. This gives rise to the central tension of the paper. A well-designed venue should enable prices to aggregate information, but it should not enable traders to profit by distorting the information that prices convey. In this environment, order flow serves a dual purpose: it reallocates the asset across investors and simultaneously communicates information. For this reason, venue design is central. A market that reveals too little information may leave traders vulnerable to ex post regret, while a market that reveals too much information too early may encourage the manipulation of beliefs and, ultimately, prices.

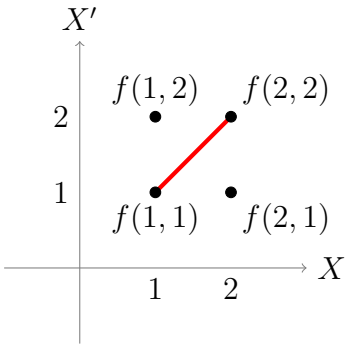
We model this tension using the standard interdependent-values environment with affiliated signals (Milgrom and Weber, 1982, 2000). The assumption is natural in financial markets: investors process different information sets, but their signals are correlated because they concern the same asset and the same aggregate shocks. Specifically, we assume that each trader's evaluation is *interdependent*: not only does it depend on the trader's own signal but also on the realization of other traders' signals. Within our setting this implies restricting assumption A.2 to the case

$$\frac{\partial v(x_i, X_{-i})}{\partial x_i} \geq \frac{\partial v(x_i, X_{-i})}{\partial x_j} > 0, \forall j \neq i.$$

Meaning that traders have an incentive to learn other traders' signals to update their own valuation. In line with the auction theory literature, we assume that trader signals X are *positively affiliated*, so that the realization of a trader's signal may provide information about the realization of the other traders' signals.[8] More specifically, in this section, we adopt the *interdependent values* setting introduced by Milgrom and Weber (1982), which consists of two ingredients: valuation functions mapping multiple signals into the asset evaluations of each trader, and affiliation of signals. This provides a powerful and rather general setting for characterizing how learning impacts trading strategies and, hence, the mapping between traders' strategies and equilibrium prices in financial markets. Consequently, this setting is well suited for studying adverse selection—i.e., asymmetric information across traders about the asset's fundamental value (Du and

[8]For a non-exhaustive list of papers using affiliation see Persico (2000), Bergin (2001), Reny and Perry (2006) and Mezzetti, Pekeč, and Tsetlin (2008).

Figure 1: **Positive Affiliation**



This figure reports an example with 2-dimensional signal vectors X and X' in the $[1,2]^2$ domain. Affiliation implies that the product of the weights associated with the points on the line $(1,1)$ and $(2,2)$ (where signals are similar, e.g., either high or low) exceeds the product of the weights associated with the off-diagonal points, e.g., $(1,2)$ and $(2,1)$ where one value is high and the other is low. Essentially, there is a higher probability to extract points along the red diagonal segment than off-diagonal.

Zhu, 2017).

Signals affiliation We assume that signals X are *positively affiliated* drawn from an unspecified continuous probability distribution f bounded in the interval $[\underline{x},\ \overline{x}]^N$. Formally, given two different vectors of signals, X and X', let $X \vee X'$ denote the component-wise maximum and $X \wedge X'$ denote the component-wise minimum. Then, X and X' are positively affiliated if:

$$f(X \vee X')f(X \wedge X') \geq f(X)f(X') \tag{12}$$

The condition expressed in equation (12) implies that the random variables are positively correlated. Intuitively, affiliation implies that it is more likely to observe either all signals low or all signals high than dispersed signals. Figure 1 reports an example with a two-dimension setting where signals have the same support. Equation (12) implies that, in this example, it is more likely that the realization of the two signals lies close to the positively sloped diagonal - between $f(1,1)$ and $f(2,2)$) - rather than far from it.

In terms of learning, this environment is particularly interesting when equation (12) is satisfied with inequality: it formalizes the assumption that a trader with signal x_i thinks other traders will have similar private values.[9] Economically, the *positive affiliation* structure captures the idea

[9]In the limit, when the condition is satisfied with equality, it encompasses independent signals so that all the

that market participants tend to have a similar view on a stock even though each comes to the market with different private information. This framework captures environments with market participants relying on different techniques to extract signals (Dugast and Foucault (2018)): in the case of a positive news, all traders come to the market with a different opinion, although all traders are likely to agree on the positive nature of the shock.

An essential property of affiliation is that, as $v(\cdot)$ is an increasing function, one has that for any integer $L \leq N$, the function $G(Y)$ defined as $G(Y) := E[v(\tilde{X})|\{X_i\}_{i=1}^{L} = Y]$, is an increasing function in its arguments.[10] This implies that public information preserves evaluation ranks. For example take any set of traders who privately know their own private signals. If these traders updates their belief based on some public additional information, then the trader with highest private signal x_i has the highest expected valuation conditional on the new public information. Hence, conditionally on the same information set, the valuation function and the expectation on the evaluation are order-preserving.

Compared to the private value setting of Section 3, the interdependent valuation case presents two complications for the design of an ex-post efficient and incentive compatible trading protocol: First, the possibility of ex-post regret if traders signals are disclosed once trading is over.[11] Second, the possibility of manipulation, when the protocol involves traders moving sequentially: a trader's action can affect the beliefs, and thus the actions, of those trading afterwards.[12]

In what follows, we study first the Median Price Trading mechanism which is robust to potential market manipulation; however, in this mechanism no information is released and therefore the potential ex-post regret is maximum. We subsequently propose a sequential trading mechanism balancing information release and trading volume.

results established in this section have general validity as they hold for the benchmark *i.i.d.* setup.

[10] See Theorem 5 in Milgrom and Weber (1982).

[11] Ex-post regret can emerge even in a simultaneous move setting if after knowing the trading outcome, an agent might regret the trade he executed (or the one he did not). A classical example is the well-known Winner's Curse in auctions.

[12] See for example Calcagno and Lovo (2006) for belief manipulation, and Avery and Zemsky (1998) for rational herding behavior

4.1 Median Price Trading Mechanism

We can now describe the equilibrium in the median price mechanism when traders have interdependent valuations. It is important to note that, in this framework, market participants face two layers of uncertainty. The first, which is also present in the private value setting, concerns the endogenous equilibrium price at which they will eventually trade. The second, unique to the interdependent value framework, relates to the uncertainty in their valuations: without knowledge of the other $N-1$ signals, each market participant can only conjecture the asset's value.

Within this setting, participants in the Median Price Trading protocol face maximum uncertainty because each participant forms his expectation based on only one of the N dispersed signals. In the previous private value setup, each market participant was indifferent to the other $N-1$ signals, as their valuation was assumed to be equal to their private signal. For this reason, the equilibrium strategy was belief-free in the sense that it did not depend on what a trader believed about the other traders' signals.

However, this is no longer the case when values are interdependent. Each trader now has an equilibrium bidding function that incorporates both their own signal (the known element) and their belief about the signals received by the other traders. Formally, let's define the function $g : \mathbb{R}^2 \to \mathbb{R}$, as follows:

$$g(x, y) = E[v(x_i, \tilde{X}_{-i})|x_i = x, x^M = y] \tag{13}$$

This function $g(x, y)$ is the expected value of the asset given trader i private signal $x_i = x$, conditional on the median signal being $x^M = y$. Note that because $v(\cdot)$ is increasing in all its arguments, $g(x, y)$ is increasing in both arguments. Then we have

Proposition 3. *Within the interdependent value setting, in the Median Price Trading protocol each trader i submitting $p_i = g(x_i, x_i)$ is an equilibrium that brings to the first best social optimum. Transaction price is $p^M = g(x^M, x^M)$, and the equilibrium trading profile is self-financed.*

Proposition 3 shows that MPT is robust on the allocation margin even when valuations are interdependent (the proof of the proposition is in Appendix C.1). Intuitively, in equilibrium each trader bids a price, $p_i = g(x_i, x_i)$, that is a function only of his own signal, and is equal to

the expected value of the asset conditional on his own signal x_i and on the median signal x^M being equal to x_i. This trading strategy ensures that market participants expected gain from trade is positive. This outcome is consistent with equilibrium bidding in auction mechanisms in which the price is set by market participants that are excluded from trading, Milgrom and Stokey (1982).

We now investigate the effect of pre-trade transparency of this mechanism and consider a slightly modified framework in which the trading platform announces the median execution price, p^M, after all traders have submitted their prices. The trading platform then allows market participants to choose either to execute at the median price or to refrain from trading. This new mechanism is reminiscent of the quasi-dark periodic auction trading system introduced in Europe in 2015, where market participants only observe the indicative auction price, Johann, Putniņš, Sagade, and Westheide (2019).[13] We show that enhancing pre-trade transparency by announcing the execution price does not alter the equilibrium auction outcome:

Corollary 1. *Announcing the median price before or after the trade does not alter the trading decisions of market participants.*

The corollary directly follows from the proof of Proposition 3, where bidding decisions were taken conjecturing traders knew the median price. According to the current market design of periodic auctions, each auction closes according to the standard walrasian pricing rule. Corollary 1 suggests that substituting the current pricing rule with the MPT mechanism would enhance the robustness of the periodic auctions mechanisms with regards to manipulation.

Although pre-trade transparency enhances the realism of our proposed mechanism and increases the information set available to market participants, it does not eliminate potential ex-post regret. Intuitively, if all private signals were disclosed after the trade, some traders might regret their trading decision. For example, the announced median price p^M could result from having signals above the median $x_i > x^M$ close to the upper bound $\overline{x}$ and signals below the

[13] As part of MiFID II, European trading venues created call auction phases. During the two-minute call phase, traders submit orders that are collected but not immediately executed and potentially deleted. Differently from lit sessions, traders have only aggregated information on the current order situation: they know the indicative price/volume at which they would trade if the auction ended immediately.

median $x_j < x^M$ close to x^M. Conversely, it could result from signals above the median $x_i > x^M$ being close to x^M and signals below the median $x_j < x^M$ being close to the lower bound $\underline{x}$. Clearly, if traders' signals were publicly disclosed after the trade, all traders' valuations would be higher in the former case than in the latter case. In the former case, some sellers might regret their trade, whereas in the latter case, some buyers might regret their decision.

Overall, this section shows that the paper's central design principle – a trader should not affect the price at which he trades – remains powerful in the presence of learning. But the proposition also clarifies the limit of one-shot benchmark formation. A single clearing price can implement the efficient crossing of orders without fully aggregating the information dispersed across traders. In finance terms, an auction can print a benchmark that is manipulation-proof and self-financing and yet still be informationally incomplete. This is the sense in which the close, dark pools, and other batch venues face not only an allocation problem but an information-design problem. Corollary 1 reinforces this interpretation. Announcing the indicative price before execution does not overturn the equilibrium. That is relevant for actual closing and periodic auctions, where venues disseminate indicative prices and imbalances prior to uncrossing. The result suggests that the key issue is not transparency per se, but the structure of transparency: which statistics are released, when they are released, and whether traders can still exploit that information before their own executions are locked in. In the next section, we analyze mechanisms that permit the venue to release information gradually while maintaining the no-manipulation property. The setting is particularly relevant for financial markets, where price discovery and execution are intrinsically linked, yet full pre-trade transparency would invite strategic trading.

4.2 Multi Round Market - Sequential Trading

In this section, we propose a sequential trading protocol with two primary objectives. First, to ensure that each investor has no incentive to manipulate other investors' beliefs. Second, conditional on achieving the first objective, to minimize ex-post regret, there should be trading transparency along the whole process.

The protocol consists of $\frac{N-1}{2}$ rounds. In the first round, traders are divided into two groups:

one group of $\frac{N-1}{2}$ buyers and another group of $\frac{N-1}{2}$ sellers. In each of the subsequent rounds, a buyer purchases one unit from a seller and then both the buyer and the seller exit the market. Thus by the end of the final round, $N-3$ traders have exchanged $\frac{N-3}{2}$ units of the asset.

The no-manipulation property is ensured by guaranteeing that the price at which a trader trades the asset does not depend on his own actions throughout the entire trading process. Given this constraint, transparency property is achieved by disclosing, at the beginning of each round, all relevant actions taken by those who have already traded and hence exited the market.

Sequential Trading Mechanism (STM) protocol

- In the starting round, $t = 0$, each trader i simultaneously submits a price p_i to the trading platform. Let $P \in \{\mathbb{R}^+\}^N$ denote the initial price profile. The $(N-1)/2$ traders who submitted a price p_i higher than the median price p^M are potential buyers, and the $(N-1)/2$ traders who submitted a price p_i lower than the median price are potential sellers. The trader who submits the median price exits the game, and the trading platform reveals his price.

- In each following round, $t \geq 1$:

  - Each potential buyer who is still active at t submits a bid $b_{i,t}$, simultaneously, each potential seller who is still active at t submits an ask $a_{i,t}$. Bids and asks are only observed by the trading platform.

  - The potential buyer (seller) who submitted the highest (resp. lowest) bid, is delivered (resp. delivers) one unit of the asset from (resp. to) the trading platform and pays (resp. is paid) the trading platform an amount equal to the highest (resp. lowest) among his competing buyers' (resp. seller's) bids, that we denote $\hat{b}_t$ (resp. $\hat{a}_t$).

  - Round t's buyer and seller exit the game and their bids are disclosed to all traders.

- The trading mechanism stops when there are only one buyer and one seller left, or at the first t such that

$$\sum_{\tau \leq t} (\hat{b}_\tau - \hat{a}_\tau) < 0 \tag{14}$$

That is, trading stops if it would require the trading platform to inject its own cash on top of what it has accumulated from previous trades.

In this protocol in the very first round the median trader is excluded. Then there are at most $T = (N-1)/2 - 1$ trading rounds, and in each trading round $t \in \{1, \ldots, T\}$ there is an equal number $N_t := T - t + 2$ of active potential buyers and active potential sellers. Denoting with $B_t \in \mathbb{R}^{N_t}$ and $A_t \in \mathbb{R}^{N_t}$ the bid and ask vectors of time t active buyers and sellers, one has that time t buyer buys for price $\hat{b}_t = b_t^{N_t-1}$ and the seller sells for $\hat{a}_t = a_t^2$, the rest of the traders are informed of the best bid and ask quotes $(b_t^{N_t}, a_t^1)$. Thus the publicly available information at the beginning of each period $t = \{1, ...T\}$ is $H_t = \left(p^M, \{(b_\tau^{N_\tau}, a_\tau^1)\}_{\tau<t}\right)$.

The sequential trading mechanism proposed here can be decomposed into a pre-trading phase in which market participants market side is determined by their initial bid that is based only on their private signal (as for the Median Price Trading Protocol) and on a sequence of second-price auctions in which market participants define their bidding with information gradually released from realized trades. Notably, the announced prices (from which market participants extract the private signals) are the winning bids; hence, they provide information about the signals of traders who just traded and exited the game. Transaction prices, $\hat{a}_t$ and $\hat{b}_t$, are not revealed because - as noted in Milgrom and Weber (1982) - otherwise the mechanism would release signals of market participants still active in the market thus incentivizing potential manipulation behavior. Moreover, the trading mechanism stops when either the budget balance compatibility condition is not met, or when, with only one trader on each side of the market, the second-price auction would no longer be feasible. The second price auction in each room ensures that (i) market participants bid competitively, (ii) the market prices are endogenously determined without relying on external price setters or trading venues knowing the evaluation function. Consequently, after the median signal is released, a market spread naturally rises as buyers and sellers trade iteratively at their different marginal valuations. Therefore, the market venue cashes in the spread as compensation for facilitating trades. It is important to note that this trading mechanism requires a third active party that crosses the spread, providing an economic foundation of intermediaries in financial markets. In contrast to standard market microstructure

theory that considers the bid-ask spread as remuneration for liquidity provision in virtue of asymmetric and inventory costs, in our framework, the market maker gains the spread to ensure an efficient price to market participants. In other words, we theoretically motivate the bid-ask spread as compensation for guaranteeing truthful information disclosure. Finally, in order to preserve competitive bidding, we show in Appendix C.4 that the sequential mechanism cannot be refined to give the last two traders a chance to trade.

We focus on symmetric equilibria: For any given trader i with signal $x_i = x$ we denote with $p_0(x)$ the price he submits in the starting round, and with $b_t(x, H_t)$ (and $a_t(x, H_t)$) the equilibrium bid (ask) made by trader i if he is an active buyer i (active seller) in trading round t. For each trader i, denote the ordered set of the remaining $(N-1)$ signals as Y_{-i} where y_{-i}^{N-1} is the largest and y_{-i}^{1} the smallest signal.

Proposition 4. *The above mechanism has the following symmetric equilibrium bidding strategies of buyers and sellers respectively :*

- *At the last round T, each active buyer (sellers) i bids the expectation of ν_i conditional on the information disclosed in previous rounds, his own signal x_i and the highest (lowest) of the other still active buyers' (sellers') signal being equal to x_i. Formally*

$$b_T(x, H_T) = E[\nu_i | H_T, x_i = y_{-i}^{N-T} = x] \tag{15}$$

$$a_T(x, H_T) = E[\nu_i | H_T, x_i = y_{-i}^{T} = x] \tag{16}$$

- *For any round $t < T$ each active buyer (active seller) bids the expectation of the buying (selling) price he would pay (be paid) in next period conditional H_t and on the highest (lowest) of his competing buyers' (competing sellers') signal being equal to his own signal.*

$$b_t(x, H_t) = E[\hat{b}_{t+1} | H_t, x_i = x, y_{-i}^{N-t} = x] \tag{17}$$

$$a_t(x, H_t) = E[\hat{a}_{t+1} | H_t, x_i = x, y_{-i}^{t} = x] \tag{18}$$

- *At round 0, any monotone function of private signals is an equilibrium strategy, hence*

without loss of generality, traders submit $p_0 = x_i$

The proof of the proposition is in Appendix C.2. In the last trade, each buyer (seller) uses the available public information and bids as if his own signal were just high (low) enough to win - hence by conjecturing that his own signal were just equal to the other remain extreme buyer's (seller's) signal. For any round $t < T$, traders anticipate other future opportunities to trade in case they lose the current bidding round: they bid the price they expect to pay (receive) in the next round. Conditional on their own signal being just high (low) enough to win the current trading round, this price makes them indifferent between executing in the current or in the next round. Finally, at time 0, as there is no trade but only positioning, traders send their private signal to the market venue. In general, traders misreport when this allows them to trade at more advantageous price. Note that because no trader action has any impact on the price at which he trades, traders do not misreport neither in the first round, when there is only buyer-seller positioning, nor in the following trading rounds.

Corollary 2. *The trading dynamic in equilibrium is the following:*

- *The expected price paid (received) by buyers (sellers) increases (decreases) over trading rounds, hence the expected revenue made by the trading venue,* $\hat{b}_t - \hat{a}_t$*, also increases over rounds*
- *Traders with more extreme private signals trade first, those with less extreme signals gather information and trade later*

The proof of the Corollary is in Appendix C.3, where we show that the expected buy (sell) price at $t+1$ conditional on winning the auction at t is greater (smaller) than the price actually paid (sold) at t. The reason is that if a buyer (seller) wins the auction at time t, his signal is larger (smaller) than any other signal in the market. Therefore, upon execution, his revealed signal will cause an upward (downward) update of the expectation of the buyer (seller) holding the largest (smallest) signal among the remaining buyers (sellers). As this is the trader who will set the execution price tomorrow, the expected buy (sell) price will increase (decrease).[14] Intuitively, the

[14]This result is consistent with Milgrom and Weber (2000) and Mezzetti et al. (2008) who show that in a

more information the price reveals, the larger the expected difference between the price at which traders buy and the one at which they sell, which is a measure of price impact. Because this expected spread is increasing over trading rounds, traders do not have an incentive to postpone their trading over time.

The equilibrium has a clear market-design interpretation. Traders with the strongest signals in absolute terms (more distant from the median signal) trade first because they are least dependent on learning from others; traders closer to the margin wait and trade after more information has been revealed. The mechanism thus endogenously matches execution urgency to informational needs. In Figure 2, we graphically report a potential trading path of the sequential mechanism with 21 initial traders. In the next corollary, we qualify the profits of the trading venue:

Corollary 3. *In each trading round, the profit of the trading venue is positive. The mechanism stops when a second price auction among buyers and sellers is no longer feasible.*

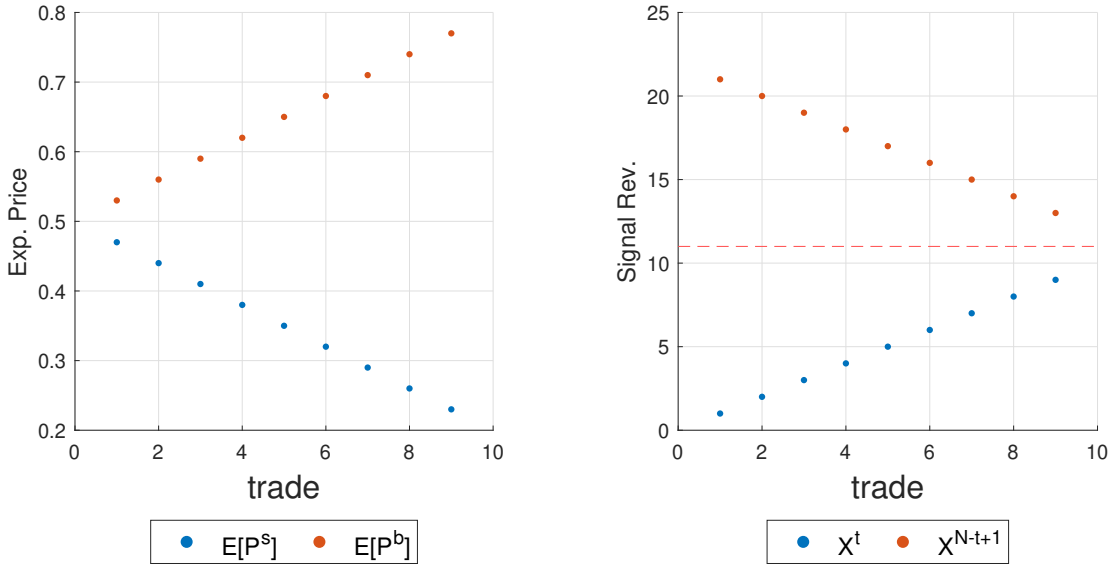
In each trading round t, the market venue sells at the second highest bid ($\hat{b}_t = b_t(x^{N-t}, H_t)$) among buyers and buys at the second lowest ask among sellers ($\hat{a}_t = a_t(x^{t+1}, H_t)$). Due to the increasing nature of the valuation function on its own signal ($y_{N-t} > y_{t+1}$) and considering that all buyers (sellers) have private signals above (below) the median, it follows that $B_t > A_t$. Therefore, the trading venue has an incentive to market making in each round until there is only one trader left on each side of the market.

It is therefore important to note that the profits earned by the trading venues do not derive from compensation for inventory or adverse selection risks. It is the endogenous payment required to sustain truthful revelation when valuations depend on information dispersed across market participants. The more information market participants need to have before trading, the more they need to compensate through the execution schedule.

From the two above properties, it follows that a sequential mechanism optimizes the trade-off

standard sequential auction with multiple buyers with interdependent valuations and a single seller, prices are subject to an upward drift. The key property that buy (selling) prices follow a sub (super) martingale is due to the common component of traders' valuation. In the special case of private values, the bid and ask prices define martingale processes. In case traders are subject to a convex execution cost, the upward drift is partially mitigated if not reversed depending on valuation functions and execution costs (Mezzetti, 2011).

Figure 2: **Dynamic of the Sequential Trading Game**
These plots report possible dynamics of the sequential trading game for $N = 21$. In the left plot, in virtue of Corollary 2, we report the expected buy and sell price over trading rounds. In the right plot, we report the signals revealed after each trading round. The red dashed line reports the median signal $X^{\frac{N+1}{2}} = 11$. The only signals not revealed are the closest to the median, $X^{\frac{N+1}{2}-1} = 10$ and $X^{\frac{N+1}{2}+1} = 12$.



between high trading volume and low potential ex-post regret. Compared to the Median Price Mechanism, volume decreases by one unit, but potential ex-post regret significantly decreases. In both mechanisms, the median trader does not trade as his price is instrumentally used to separate buyers and sellers. In the sequential mechanism, also the traders adjacent to the median do not trade to guarantee competitive prices iteratively, hence one extra unit of volume is lost. In terms of welfare, however, the loss is minimal as the traders excluded from trading are the closest to the median and therefore those generating the lowest gains from trade $(v(x^{M+1}) - v(x^{M-1}))$.

In terms of ex-post regret, the sequential mechanism substantially improves compared to the MPT. Traders with a signal far from the median trade first with less information: these traders, holding stronger signal, prefer to trade with less available information in order to execute at more competitive prices. Given their own strong signal and the competitive price at which they trade, they are naturally less exposed to regret their trades ex-post. Moreover, as the trading game progresses, the traders closer to the median signal (the ones more exposed to ex-post regret

in the MPT mechanism) have less uncertainty as public information increases (at the cost of less competitive execution prices), thus naturally reducing potential ex-post regret. Overall, the ex-post regret reduces as the sequential mechanism mediates the needs and structural differences between impatient traders (by offering competitive execution prices) and patient traders (by providing more information).

The trade-off between the expected spread and pre-trade information crucially depends on the number of shares traded in each round. In the baseline specification, one share is traded per round. The next corollary characterizes how equilibrium outcomes change when $n \in 1, N - 1/2$ units of the asset is exchanged in each round. That is if n units are traded in every round, each auction is $(n+1)^{th}$ price auction. For example, on the buyers side, each of the $n-$ highest bidders purchase one asset and each pays the $(n+1)^{th}$ highest bid. For $n = N - 1/2$, we are back to the MPT mechanism with a single round of trading occurring at the median bid.

Corollary 4. *By decreasing (increasing) the speed of the STM—that is, by decreasing (increasing) the number of units traded per round—the average number of signals available before trading increases (decreases), while the spread increases (decreases).*

Corollary 4 highlights a design trade-off for the social planner in choosing the number of units traded per period.[15] A *faster market* (higher number of shares traded per round) decreases spreads, while a *slower market* (smaller number of shares traded per round) improves information aggregation and reduces ex-post regret. Intuitively, the fewer traders the market maker needs to incentivize to reveal the signals, the lower the rents he can extract. Following Budish et al. (2015), subsequent work has sought "a richer understanding of the costs of lengthening the time between auctions." Du and Zhu (2017) show that increasing trading frequency allows investors to adjust more quickly to news but reduces per-auction depth because traders anticipate future opportunities; in extensions with heterogeneous speeds, fast (slow) traders prefer short (long) auction durations. We complement this literature by endogenizing information release: market participants release information by trading rather than taking arrivals as exogenous. Consistent with Du and Zhu (2017), we find that slowing the market (intended as more auctions held to

[15] Appendix C.5 provides an analytical discussion of the Corollary.

allocate K shares among N traders) increases the informativeness of trading but can raise per-auction price impact.[16] Finally, our results echo Jagannathan (2022), who shows that lengthening the interval between batch auctions makes markets less prone to crashes.

In Appendix C.6, we propose a simple sequential mechanism that minimizes ex post regret at the expense of trading volume. In the next section, we discuss how our mechanisms connect with existing trading mechanisms.

5 Discussion

How do the Median Price Trading and Sequential Trading mechanisms compare with the trading mechanisms used in practice? Our new mechanisms isolate which institutional features invite strategic distortion and which features prevent it. In other words, the paper speaks not to one venue alone but to a class of markets in which the execution price is economically valuable and market activity is informative.

Closing auctions. Closing auctions are the clearest application because they combine one-shot execution with benchmark formation. The closing price is not merely the price at which the final trade occurs; it is a reference price used for valuation, performance measurement, index replication, and settlement. That is precisely the environment in which finite-trader price-taking becomes fragile. MPT addresses that fragility directly. It internalizes the benchmark inside the venue and assigns price-setting power to the trader who does not trade. The mechanism is therefore best interpreted as a benchmark-forming auction that preserves allocative efficiency without rewarding traders for moving the benchmark they face.

MPT provides a mechanism-design response to this finite-trader benchmark problem. The efficient allocation requires an endogenous cutoff—the marginal trader—and MPT implements it through the median separation of buyers and sellers. At the same time, the no-manipulation

[16]The results are consistent with Du and Zhu (2017), as they define Trading frequency as the number of double auctions per unit of clock time. Hence, keeping constant clock time, increasing the number of auctions implies faster markets. In an environment close to Du and Zhu (2017), Paul, Thibaut, and Mathieu (2021) estimate optimal auction durations on the order of minutes and find that continuous limit-order books are mildly sub-optimal in terms of price discovery.

requirement is satisfied because price-setting power is assigned to the trader who does not trade. This is the sense in which MPT addresses a foundational closing-auction concern: it yields an internally formed benchmark while preventing participants from profiting by moving the benchmark price they face.

In information-rich closing environments, however, allocative efficiency is only part of the objective. The closing print is often treated as a public signal, so informational efficiency matters as well. Our interdependent-values analysis clarifies why one-shot benchmark formation can be informationally incomplete even when it is allocatively efficient and manipulation-free: revealing only a single clearing statistic can leave substantial private information unaggregated. This is where STM becomes relevant as a design for controlled disclosure. Its central lesson is not that real-world closing auctions should be replaced one-for-one by our sequential protocol, but that *the timing of transparency* is itself a mechanism-design choice. Information that is released while a trader can still affect his own terms of trade invites strategic signaling; information that is released only after a trader is committed can improve price discovery without rewarding manipulation. This principle maps naturally into the design of indicative-price dissemination, imbalance publication, and late-order flexibility in auction windows: the question is not "transparency versus opacity," but "which statistics are released when, and with what commitment from the traders who generate them."

Dark pools and crossing venues. Dark pools have been created so traders could potentially trade large quantities with a minimum price impact. Most of dark pools in Europe and in the U.S. follow a derivative pricing rule: the execution price generally relies on pricing information sourced from lit venues.[17] Given their low transparency, the difficulty of monitoring, and the reliance on outside trading prices, the literature has established three channels through which dark pools can reduce market participants' welfare. First, the execution price may not be set consistently with the reference price in the lit market.[18] Second, given the proliferation of co-

[17]Traditional dark pools match customer orders using either the midpoint of the national (European) best bid and offer (NBBO or EBBO), or the volume-weighted average price (VWAP) of the primary market, Zhu (2014).

[18]In June 2014, the Financial Industry Regulatory Authority (FINRA), which oversees securities firms in the U.S., accused Goldman Sachs' SIGMA X dark pool of executing nearly 400,000 trades between July 29, 2011, and August 9, 2011, at prices lower than those available on public exchanges. Although Goldman Sachs did not

located and sophisticated traders, large "slow" institutional traders are exposed to stale quote costs (Securities, Commission, et al. (2015) and Aquilina et al. (2024)): the execution price in the dark pool may rely on an old, outdated, reference price derived from the lit venues, creating latency arbitrage opportunities. This is a concern especially for high priced-liquid stocks, when all stale trading involves high frequency traders.[19] Third, traders can potentially manipulate the reference price in the dark pool by submitting orders in the lit market. This may be a concern especially for illiquid stocks and during periods of market turmoil.

MPT directly isolates and removes these vulnerabilities by internalizing benchmark formation. In MPT the execution price is a function of the submitted prices inside the venue—the median—rather than an externally supplied benchmark. This eliminates the cross-venue manipulation channel by construction: trading in a lit market cannot move the dark execution price because the mechanism does not import that price. It also neutralizes latency-arbitrage logic because execution does not hinge on being faster than the reference-price updating process in a different venue. Most importantly, it makes manipulation of the execution price privately unprofitable even when the set of participants is finite. The reason is the price-taking principle embedded in the design: the only trader who can affect the median is the median agent, and he does not trade, so he does not benefit from moving the price. The point of the comparison is the design principle: if a venue relies on an external benchmark, then the mechanism inherits the external venue's informational fragility and becomes vulnerable to cross-venue strategic behavior. A mechanism that determines the benchmark internally and that decouples price-setting power from trading incentives can mitigate these vulnerabilities at their source. Nor is the endogenous determination of trade direction in MPT obviously at odds with dark-pool participation. Empirical evidence suggests that a meaningful share of dark-pool activity, especially by some high-frequency traders, is opportunistic rather than tied to a fixed ex ante need to buy or sell (Aquilina et al., 2024). Read in that light, the mechanism's endogenous sorting of buyers and sellers is not an artificial feature, but one that matches an important component of trading

admit wrongdoing, they agreed to settle FINRA's allegations by paying an $\$800,000$ penalty.

[19] Aquilina et al. (2024) observe both a positive correlation between stale trading and market capitalization, and a weak correlation between stale trading and nominal stock prices. Hence, the stale trading social cost is likely to be higher for highly liquid, high-volume stocks, as HFTs can rapidly and almost freely balance in lit markets any position they assume in a dark pool.

observed in these venues.

Continuous limit order books. Continuous limit order books remain the dominant trading protocol for many instruments, largely because they provide immediacy and continuous adjustment. Yet the same features that make continuous markets attractive also create well-understood fragilities. When orders can be posted and cancelled at high frequency, the order book becomes not only a trading device but also a communication device: strategic traders can submit orders to move beliefs and then cancel them, a logic formalized by Williams and Skrzypacz (2020) and empirically shown in Brogaard et al. (2022). At the same time, continuous time priority and rapid updating create incentives for speed races that can dissipate rents and shape market quality (Budish et al., 2015; Foucault, Hombert, and Roşu, 2016). These concerns are fundamental because they arise from equilibrium incentives created by the protocol, not from rare deviations.

Our sequential mechanism clarifies a design route to mitigate these equilibrium incentives: ensure that no trader can affect the price at which he trades *and* ensure that information becomes public only after the trader who generated it can no longer exploit it. STM implements this via second-price formats, exit after trade, and post-trade disclosure of bids. The mechanism therefore targets belief manipulation by severing the link between "being informative" and "having future strategic leverage." It also eliminates the speed advantage inherent in continuous time priority because trading occurs in discrete rounds with well-defined participation and without order-cancellation races.

At the same time, the comparison with continuous limit order books must be made with caution. Our baseline environment abstracts from several features that are central to the economic role of continuous trading, including inventory management, repeated position adjustment, and the provision of immediacy through standing limit orders. These considerations help explain why continuous markets exist and why observed bid-ask spreads reflect not only adverse selection and strategic trading incentives, but also compensation for inventory risk and immediacy provision. For that reason, our results should not be read as a general case against limit order books.

The more appropriate implication is comparative. The relative advantage of a continuous limit order book is greatest when immediacy is especially valuable and when the informational

and manipulative vulnerabilities of continuous trading are limited. By contrast, the case for the Sequential Trading Mechanism is stronger in environments in which immediacy is less important, liquidity is thin, or uncertainty is high—precisely the settings in which continuous trading is more exposed to strategic order placement, cancellation-based distortions, and price manipulation. In such markets, an auction-based mechanism may offer larger marginal benefits because it internalizes price formation, disciplines the release of information, and reduces the scope for traders to exploit the trading process itself.

6 Conclusion

This paper studies a basic problem in financial market design: how to allocate assets and form prices when traders may profit both from trading on information and from influencing the mechanism that aggregates it. The paper's central point is that these distortions are often not external frictions layered on top of a neutral institution. They are equilibrium outcomes generated by the rules of price formation and information release.

Our analysis isolates two core design requirements. Efficient trade requires an endogenous marginal trader who separates buyers from sellers. Manipulation-resistance requires individual price-taking: no trader should be able to affect the price at which he trades. In the private-values environment, the Median Price Trading mechanism translates those principles into a simple, self-financing market rule. In the interdependent-values environment, the paper shows why efficient one-shot crossing is not enough. A venue may remain allocatively efficient and yet leave too much information unaggregated. The Sequential Trading Mechanism addresses that second problem by releasing information only after the trader who generated it has exited and can no longer exploit the beliefs of others.

The broader implication is methodological and economic rather than venue-specific. Mechanism design provides a disciplined way to evaluate and redesign trading rules when the mechanism, not merely the environment, creates opportunities for distortion. Dark pools, closing auctions, and continuous limit order books are three prominent settings in which these issues arise, but the core logic is more general: whenever a venue must both allocate assets and release

information, the structure of price formation and the timing of information release are first-order design choices.

Our paper suggests a different way to think about manipulation. Some forms of strategic behavior are best treated not as pathologies to be policed *ex-post*, but as outcomes to be designed away *ex-ante*. The mechanisms in this paper do not solve every market-design problem. They do, however, show how to start from first principles and build trading rules that align private incentives more closely with allocative efficiency and informative price formation.

References

Aggarwal, Rajesh K, and Guojun Wu, 2006, Stock market manipulations, *The Journal of Business* 79, 1915–1953.

Allen, Franklin, and Gary B Gorton, 1991, Stock price manipulation, market microstructure and asymmetric information.

Andreyanov, Pasha, and Tomasz Sadzik, 2021, Robust mechanism design of exchange, *The Review of Economic Studies* 88, 521–573.

Aquilina, Matteo, Sean Foley, Peter O'Neill, and Thomas Ruf, 2024, Sharks in the dark: Quantifying hft dark pool latency arbitrage, *Journal of Economic Dynamics and Control* 158, 104786.

Avery, Christopher, and Peter Zemsky, 1998, Multidimensional uncertainty and herd behavior in financial markets, *American economic review* 724–748.

Benabou, Roland, and Guy Laroque, 1992, Using privileged information to manipulate markets: Insiders, gurus, and credibility, *The Quarterly Journal of Economics* 107, 921–958.

Bergin, James, 2001, Common knowledge with monotone statistics, *Econometrica* 69, 1315–1332.

Brogaard, Jonathan, Dan Li, and Jeffrey Yang, 2022, Does high frequency market manipulation harm market quality?, *Available at SSRN 4280120* .

Budish, Eric, Peter Cramton, Albert S Kyle, Jeongmin Lee, and David Malec, 2023, Flow trading, Technical report, National Bureau of Economic Research.

Budish, Eric, Peter Cramton, and John Shim, 2015, The high-frequency trading arms race:

Frequent batch auctions as a market design response, *The Quarterly Journal of Economics* 130, 1547–1621.

Calcagno, Riccardo, and Stefano Lovo, 2006, Bid-ask price competition with asymmetric information between market-makers, *The Review of Economic Studies* 73, 329–355.

Chen, Daniel, 2025, Optimal exchange design, *Personal Website* .

Colla, Paolo, Walter Distaso, and Paolo Vitale, 2024, Price manipulation and collusion around the london 4pm fix, *Available at SSRN 4819478* .

Comerton-Forde, Carole, and Tālis J Putniņš, 2011, Measuring closing price manipulation, *Journal of Financial Intermediation* 20, 135–158.

Du, Songzi, and Haoxiang Zhu, 2017, What is the optimal trading frequency in financial markets?, *The Review of Economic Studies* 84, 1606–1651.

Dugast, Jérôme, and Thierry Foucault, 2018, Data abundance and asset price informativeness, *Journal of Financial economics* 130, 367–391.

Ernst, Thomas, Chester S Spatt, and Jian Sun, 2022, Would order-by-order auctions be competitive?, *Available at SSRN 4300505* .

Foley, Sean, Peter O'Neill, and Tālis J Putniņš, 2023, A better market design? applying 'automated market makers' to traditional financial markets, *Applying 'Automated Market Makers' to Traditional Financial Markets (June 15, 2023)* .

Foucault, Thierry, Johan Hombert, and Ioanid Roşu, 2016, News trading and speed, *The Journal of Finance* 71, 335–382.

Glosten, Lawrence R, 1994, Is the electronic open limit order book inevitable?, *The Journal of Finance* 49, 1127–1161.

Glosten, Lawrence R, and Paul R Milgrom, 1985, Bid, ask and transaction prices in a specialist market with heterogeneously informed traders, *Journal of financial economics* 14, 71–100.

Jagannathan, Ravi, 2022, On frequent batch auctions for stocks, *Journal of Financial Econometrics* 20, 1–17.

Jegadeesh, Narasimhan, and Yanbin Wu, 2022, Closing auctions: Nasdaq versus nyse, *Journal of Financial Economics* 143, 1120–1139.

Johann, Thomas, Tālis J Putniņš, Satchit Sagade, and Christian Westheide, 2019, Quasi-dark trading: The effects of banning dark pools in a world of many alternatives .

Krishna, Vijay, 2009, *Auction theory* (Academic press).

Kyle, Albert S, 1985, Continuous auctions and insider trading, *Econometrica: Journal of the Econometric Society* 1315–1335.

Lehar, Alfred, and Christine A Parlour, 2021, Decentralized exchanges, *Available at SSRN* 3905316.

Malinova, Katya, and Andreas Park, 2023, Learning from defi: Would automated market makers improve equity trading?, *Available at SSRN 4531670* .

McAfee, R Preston, 1992, A dominant strategy double auction, *Journal of economic Theory* 56, 434–450.

Mezzetti, Claudio, 2011, Sequential auctions with informational externalities and aversion to price risk: decreasing and increasing price sequences, *The Economic Journal* 121, 990–1016.

Mezzetti, Claudio, Aleksandar Saša Pekeč, and Ilia Tsetlin, 2008, Sequential vs. single-round uniform-price auctions, *Games and Economic Behavior* 62, 591–609.

Milgrom, Paul, and Nancy Stokey, 1982, Information, trade and common knowledge, *Journal of economic theory* 26, 17–27.

Milgrom, Paul R, and Robert J Weber, 1982, A theory of auctions and competitive bidding, *Econometrica: Journal of the Econometric Society* 1089–1122.

Milgrom, Paul R, and Robert J Weber, 2000, A theory of auctions and competitive bidding ii, *The Economic Theory of Auctions Vol. II* 113, 179–194.

Myerson, Robert, 1979, Incentive-compatibility and the bargaining problem, *Econometrica, 47, 61-73* .

Myerson, Roger B, and Mark A Satterthwaite, 1983, Efficient mechanisms for bilateral trading, *Journal of economic theory* 29, 265–281.

O'Hara, Maureen, 2003, Presidential address: Liquidity and price discovery, *The journal of Finance* 58, 1335–1354.

Paul, Jusselin, Mastrolia Thibaut, and Rosenbaum Mathieu, 2021, Optimal auction duration: A price formation viewpoint, *Operations Research* 69, 1734–1745.

Persico, Nicola, 2000, Information acquisition in auctions, *Econometrica* 68, 135–148.

Reny, Philip J, and Motty Perry, 2006, Toward a strategic foundation for rational expectations equilibrium, *Econometrica* 74, 1231–1269.

Rustichini, Aldo, Mark A Satterthwaite, and Steven R Williams, 1994, Convergence to efficiency in a simple market with incomplete information, *Econometrica: Journal of the Econometric Society* 1041–1063.

Securities, Australian, Investments Commission, et al., 2015, Review of high-frequency trading and dark liquidity, *Report* 452, 20.

Tremacoldi-Rossi, Pedro, 2024, Manipulating algorithmic markets, *Available at SSRN 5025393* .

Van Bommel, Jos, 2003, Rumors, *The journal of Finance* 58, 1499–1520.

Vila, Jean-Luc, 1989, Simple games of market manipulation, *Economics Letters* 29, 21–26.

Williams, Basil, and Andrzej Skrzypacz, 2020, Spoofing in equilibrium .

Zhu, Haoxiang, 2014, Do dark pools harm price discovery?, *The Review of Financial Studies* 27, 747–789.

Appendices

A Appendix Section 2

A.1 Proof of Proposition 1

Let first prove that the mechanism is ex-post IC. Take a trader with $x_i > x^M$, if he truthfully reports x_i or any other $x_i' > x^M$ his payoff is $q_i(X)v_i(x_i, X_{-i}) - t_i(X) = v_i(x_i, X_{-i}) - v_i(x^M, X_{-i}) > 0$, where the inequality follows from $x_i > x^M$. If instead, he reports $x_i' = x^M$, he does not trade and gets 0. If he reports $x_i' < x^M$, then he sells for $v_i(x^M, X_{-i})$, realizing payoff of $v_i(x_M, X_{-i}) - v_i(x^i, X_{-i}) < 0$. Hence truthfully reporting his signal is a best response. A symmetric argument applies for a trader whose $x_i < x^M$. Consider now the median signal trader,

i.e. the one with $x_i = x^M$. If he reports any $x_i' \in (x^{M-1}, x^{M+1})$, he remains the median trader and does not trade getting 0. If he reports $x_i' < x^{M-1}$, then the resulting action profile $X' = \{x_i', X_{-i}\}$ has a median signal $x'^M = x^{M-1} < x^M$. Trader i sells for $t_i(X') = v_i(x'^M, X_{-i})$ realizing a payoff of $v_i(x'^M, X_{-i}) - v_i(x_i, X_{-i}) < 0$, where the inequality follows from $x_i = x^M > x^{M-1} = x'^M$. A symmetric argument applies if he reported $x_i' > x^{M+1}$. Thus, the mechanism is ex-post IC for each trader.

Ex-post IR is guaranteed by the fact that each seller (buyer) values the asset $v_i(x_i, X_{-i})$ and sells for (buys for) $t_i(X) = \nu_i(x^M, X_{-i})$ whenever $x_i < x^M$ (resp $x_i > x^M$) that is larger (resp. smaller) than his own valuation.

B Appendix Section 3

B.1 Proof of Proposition 2

We first show that for any given trader i, bidding $p_i = x_i$ is weakly optimal even ex-post i.e., after knowing the other trader's prices. Thus, bidding $p_i = x_i$ is optimal for trader i no matter his beliefs about the other trader's valuations or strategies.

Fix trader i and an arbitrary price profile $P_{-i} \in \mathbb{R}^{N-1}$ for all other traders. Let p^{M+} and p^{M-} be the $(N-1)/2$-th highest, and the $(N-1)/2$-th lowest price among the other traders prices P_{-i}, respectively. Note that $p^{M-} \leq p^{M+}$, and there is no price by i's competitor between p^{M-} and p^{M+}. Depending on p_i, the overall median price for the N traders is:

If trader i sets $p_i > p^{M+}$, then the median price is $p^M = p^{M+}$, and trader i buys for p^{M+} and realizes $x_i - p^{M+}$.

If trader i sets $p_i \in [p^{M-} p^{M+}]$, then the overall median price is $p^M = p_i$, and trader i does not trade and realizes 0.

If trader i sets $p_i < p^{M-}$, then the median price is $p^M = p^{M-}$, and trader i sells for p^{M-} and realizes $p^{M-} - x_i$.

Thus if $x_i > p^{M+}$ ($x_i < p^{M-}$), then all $p_i > p^{M+}$ (resp. $p_i < p^{M-}$) provide the same strictly positive payoff of $x_i - p^{M+}$ (resp. $p^{M-} - x_i$), whereas any price $p_i \leq p^{M+}$ (resp. $p_i \geq p^{M-}$)

provides either 0 (for $p_i \in [p^{M-} p^{M+}]$) or $p^{M-} - x_i < 0$ (resp. $x_i - p^{M+} < 0$) for $p_i < p^{M-}$ (resp. $p_i > p^{M+}$). Thus setting $p_i = x_i$ is a best response to P_{-i}.[20]

The above reasoning also proves that the mechanism is *ex-post incentive compatible* as traders set $p_i = x_i$ no matter their belief about the other trader's types profile X_{-i}.

The mechanism is *ex-post efficient* as the buyers have all an evaluation above the median and sellers all below the median, and the resulting trades lead to the first best social welfare The mechanism is *ex-post individually rational* as buyers know their own valuations, they learn nothing from trade so they have no regret to have traded and their ex-post payoff is non-negative. In case of buyers (sellers) the ex-post gains are strictly positive and in case of the price setter the gain is zero. The mechanism is *ex-post budget balance compatible* as there is one single clearing price that cancels out in aggregate. The mechanism maximizes volume if the probability of extracting two equal signals is nil: in this case the mechanism leads to $\frac{N-1}{2}$ trades.

C Appendix Section 4

C.1 Proof of Proposition 3

Note first that because $v(\cdot)$ is non-decreasing in all its arguments and strictly increasing in at least one argument, the function $g(x, x)$ is invertible. Namely let $\phi(z)$ be the inverse of the bidding function, that is, the x such that $g(x, x) = z$. Note that $\phi(z)$ is increasing in z.

Fix trader i and suppose he expects each other trader j to set price $p_j = g(x_j, x_j)$. Let $p^M = g(x^M, x^M)$ denote the trading price, i.e., the median price, if trader i submits $p_i = g(x_i, x_i)$. There are three scenarios, depending on whether by setting $p_i = g(x_i, x_i)$ trader i is a buyer, a seller, or does not trade.

Consider first the case $p_i > p^M$. That is, by submitting his price, trader i buys for p^M and

[20] For example let $N = 5$ and $P_{-i} = \{4, 7, 10, 20\}$, then $P^{M-} = 7$ and $p^{M+} = 10$. If trader i sets $p_i < 7$, then the trading price is $p^M = 7$, and trader i sells for 7. If $7 \leq p_i \leq 10$, the trading price is $p^M = p_i$, and trader i does no trade. If $p_i > 10$, then trader i buys and pays the trading price of $p^M = 10$.

his payoff is

$$E[v(x_i, X_{-i})|x_i, v^M = p^M] - p^M = g(x_i, \phi(p^M)) - g(\phi(p^M), \phi(p^M)) \tag{19}$$

That is positive because $p_i > p^M$ implies $x_i > \phi(p^M)$, and $g(\cdot)$ is strictly increasing in its first argument. Now suppose that instead trader i deviates to $\tilde{p}_i = g(\tilde{x}, \tilde{x}) \neq g(x_i, x_i)$, and let $\tilde{p}^M$ denote the resulting median price. Then:

- if $\tilde{p}_i > p^M$, then $\tilde{p}^M = p^M$, as the price submitted by trader i does not change the median price as it can only impacts the ordering among the $(N-1)/2$ prices above the median. Trader i keeps buying for p^M and his payoff remains $g(x_i, \phi(p^M)) - g(\phi(p^M), \phi(p^M)) > 0$.

- if $\tilde{p}_i < p^M$, then either trader i becomes the median price trader, hence does not trade and gets 0, or trader i becomes a seller. In the latter case the new median price becomes $\tilde{p}^M < p^M$, and trader i sells at this price with a resulting payoff of

  $$\tilde{p}^M - E[v(x_i, \tilde{X}_{-i})|x_i, v^M = \tilde{p}^M] = g(\phi(\tilde{p}^M), \phi(\tilde{p}^M)) - g(x_i, \phi(\tilde{p}^M)) < 0$$

  where the inequality follows from $g(x_i, x_i) = p_i > p^M > \tilde{p}^M = g(\phi(\tilde{p^M}), \phi(\tilde{p^M}))$. Therefore, it follows that $x_i > \phi(\tilde{p}^M)$, and so $g(\phi(\tilde{p}^M), \phi(\tilde{p}^M)) < g(x_i, \phi(\tilde{p}^M))$. Hence if trader i expects to be a buyer, then setting $p_i = g(x_i, x_i)$ is a best response.

Second, consider the case $p_i < p^M$. That is when playing the prescribed price p_i, the trader i sells for p^M and his expected payoff is

$$p^M - E[v(x_i, X_{-i})|x_i, v^M = p^M] = g(\phi(p^M), \phi(p^M)) - g(x_i, \phi(p^M)) > 0$$

where the inequality follows form $p_i < p^M$ that implies $x_i < \phi(p^M)$. Now suppose that instead trader i deviates to $\tilde{p}_i = g(\tilde{x}, \tilde{x}) \neq g(x_i, x_i)$, and let $\tilde{p}^M$ denote the resulting median price. Then:

- if $\tilde{p}_i < p^M$, there is no change in the median price, trader i keeps selling for p^M and get the same strictly positive payoff as when setting $p_i = g(x_,x_i)$.

- if $\tilde{p}_i > p^M$, then either trader i becomes the median price trader, hence does not trade and gets 0, or trader i becomes a buyer. In the latter case the new median price becomes $\tilde{p}^M > p^M$ and trader i buys at this price with a resulting expected payoff of

$$E[v(x_i, \tilde{X}_{-i})|x_i, v^M = \tilde{p}^M] - \tilde{p}^M = g(x_i, \phi(\tilde{p}^M)) - g(\phi(\tilde{p}^M), \phi(\tilde{p}^M)) < 0$$

  where the inequality follows from $g(x_i, x_i) = p_i < p^M < \tilde{p}^M = g(\phi(\tilde{p}^M), \phi(\tilde{p}^M))$. Therefore, it follows that $x_i < \phi(\tilde{p}^M)$, and so $g(\phi(\tilde{p}^M), \phi(\tilde{p}^M)) > g(x_i, \phi(\tilde{p}^M))$. Hence if trader i expects to be a seller, then setting $p_i = g(x_i, x_i)$ is a best response.

Finally, consider $p_i = p^M$. That is, when playing the prescribed price, trader i is the median price trader, he does not trade, and gets 0. If he deviates to some $\tilde{p}_i \neq g(x_i, x_i)$ but remains the median trader, he does not trade and still gets 0. Suppose instead that the deviation brings a new median price $\tilde{p}^M > p^M$, then it must be that $\tilde{p}_i > \tilde{p}^M$, trader i buys and realizes a payoff of

$$E[v(x_i, \tilde{X}_{-i})|x_i, v^M = \tilde{p}^M] - \tilde{p}^M = g(x_i, \phi(\tilde{p}^M)) - g(\phi(\tilde{p}^M), \phi(\tilde{p}^M)) < 0$$

where the inequality follows from $g(x_i, x_i) = p^M < \tilde{p}^M = g(\phi(\tilde{p}^M), \phi(\tilde{p}^M))$, implying $\phi(\tilde{p}^M) > x_i$ and hence $g(\phi(\tilde{p}^M), \phi(\tilde{p}^M)) < g(x_i, \phi(\tilde{p}^M))$. If instead the deviation brings a new median price $\tilde{p}^M < p^M$, then it must be that $\tilde{p}_i < \tilde{p}^M$, trader i sells and realizes a payoff of

$$\tilde{p}^M - E[v(x_i, \tilde{X}_{-i})|x_i, v^M = \tilde{p}^M] = g(\phi(\tilde{p}^M), \phi(\tilde{p}^M)) - g(x_i, \phi(\tilde{p}^M)) < 0$$

where the inequality follows from the fact that $g(x_i, x_i) = p^M > \tilde{p}^M = g(\phi(\tilde{p}^M), \phi(\tilde{p}^M))$ implying $\phi(\tilde{p}^M) < x_i$ and so $g(\phi(\tilde{p}^M), \phi(\tilde{p}^M)) < g(x_i, \phi(\tilde{p}^M))$. Thus even when the trader expects to be the median price trader, setting $p_i(x_i, x_i)$ is the best response.

Note that the resulting outcome is ex-post allocation efficient as it leads to the first best allocation. Also, the mechanism is *ex-post budget balance compatible* as there is one single clearing price that cancels out in aggregate. The mechanism maximizes volume as the probability of extracting two equal signals is zero, and hence, the probability of having N distinct values is

1. Therefore, the mechanism determines $\frac{N-1}{2}$ trades.

The mechanism is ex-interim individually rational as the expected payoff of all market participants is non-negative. Namely, traders expect to profit from trade ex-ante, and once they know the trading price, they do not expect to lose by bidding $p_i = g(x_i, x_i)$.

C.2 Proof of Proposition 4

The proof of the equilibrium strategies is by backward induction.
The bidding equilibrium strategies for buyers and sellers in each period $t = \{1, ...T\}$ is an extension of Mezzetti et al. (2008): their paper defines the equilibrium bidding strategy of a sequential T round second price auction with bidding announcement. We extend their framework by considering a second price auction on each side of the market with common information set and showing that at time $t = 0$, truthful reporting is weakly dominating.

The proof shows by backward induction that truthful reporting is a weak dominant strategy (i) in the last trading round, (ii) in a generic first round trading, and (iii) in the first "separating" round.

C.2.1 The last trading period T.

In the last period of the trading game $T = (N-1)/2-1$, there are two bidders in each room that compete in a second price auction style. Information H_T is publicly available to all active traders. This includes the bids of all traders who already traded, and the fact that the bidders in the buyers' (sellers') room have signals larger (lower) than the median trader signal and lower (larger) than all other traders who already purchased (sold) the asset in the previous rounds. At this stage the game in the buyers room reduces to a classic second price bid auction with public information available to all active traders, and the equilibrium bidding strategy is standard as it directly follows from Milgrom and Weber (1982), that is, active bidder i bids the expected value of the asset conditional on H_T, and on the highest of his competing bidder's type being equal to x_i. Symmetric result applies to the sellers' room. Thus equations (15) and (16).

C.2.2 The intermediate trading periods $1 \leq t < T$.

Consider the bid auction. In a symmetric equilibrium, at any given round t, a trader's strategy only depends on his private signal x_i and the public information H_t. Namely let $b_t(x, H_t)$ denote the equilibrium bidding strategy of an active buyer i at time t with signal $x_i = x$. Let $V_t(x, H_t)$ denote an active trader expected equilibrium continuation payoff at the beginning of period t. This can be seen as the probability of buying the asset by the end of the game times the expected profit from trade. Namely

$$V_t(x, H_t) := \underbrace{W_t(x, H_t)}_{\text{Probability of buying}} \underbrace{(v_w(x, H_t) - P_w(x, H_t))}_{\text{Profit conditional on buying}} \tag{20}$$

where $W_t(x, H_t)$ is the equilibrium probability of buying the asset between t and the last trading round given $x_i = x$, and H_t. The variables $v_w(x, H_t)$ and $P_w(x, H_t)$ are the equilibrium expected value of the asset and the expected payment, respectively, when these expectations are conditional on $x_i = x$, H_t, and buying the asset between round t and the last trading round.

For any given active trader i at t, let $\tilde{y}_{-i}^{T-t}$ denote the highest private signal among this trader's competitors in round t, and let $f_t(\cdot)$ denote the density of $\tilde{y}_{-i}^{T-t}$ conditional on all information available to trader i at the beginning of round t, that is, $x_i = x$ and H_t.

Fix a time $t > 0$, $t < T$, consider a trader i facing other traders who use strategy $b_\tau(\cdot)$, $\tau \geq 0$. The trader's best response at t given $x_i = x$ and the public history H_t solves:

$$\max_{z \in [\underline{x}, \overline{x}]} \int_{\underline{x}}^{z} (E[\nu_i | x_i = x, H_t, \tilde{y}_{-i}^{T-t} = y] - b_t(y, H_t)) f_t(y) dy + \int_{z}^{\overline{x}} V_{t+1}(x, \{H_t, \tilde{y}_{-i}^{T-t} = y\}) f_t(y) dy \tag{21}$$

where the objective function is trader i's expected payoff in round t from making a bid equal to $b_t(z, H_t)$. Namely, the first integral reflects the trader's expected payoff when buying the asset at t, and the second integral is the continuation expected payoff in case he does not buy the asset at t. Note that at the end of each round, the trading protocol discloses the bid of the winning trader. Because equilibrium bids are monotonic in signals, this also discloses time t buyer's signal. Thus if trader i does not buy at t, the buyer will be his competitor with the

highest signal and thus the public history at $t+1$ is $H_{t+1} = \left\{H_t, \tilde{y}_{-i}^{T-t} = y\right\}$. This explains the argument of $V_{t+1}(\cdot)$ in equation (21).

In equilibrium, the first order condition requires that the derivative of the objective function with respect to z is equal to zero for $z = x$. Simplifying and rearranging, the f.o.c. can be written as follows

$$\begin{aligned}
b_t(x, H_t) &= E[\nu_i | H_t, x_i = x, \tilde{y}_{-i}^{T-t} = x] - V_{t+1}(x, \left\{H_t, \tilde{y}_{-i}^{T-t} = x\right\}) \\
&= E[\nu_i | H_t, x_i = x, \tilde{y}_{-i}^{T-t} = x] \\
&\quad -W_{t+1}(x, \left\{H_t, \tilde{y}_{-i}^{T-t} = x\right\})(v_w(x, \left\{H_t, \tilde{y}_{-i}^{T-t} = x\right\}) - P_w(x, \left\{H_t, \tilde{y}_{-i}^{T-t} = x\right\}) \\
&= E[\nu_i | H_t, x_i = x, \tilde{y}_{-i}^{T-t} = x] - v_w(x, \left\{H_t, \tilde{y}_{-i}^{T-t} = x\right\}) + P_w(x, \left\{H_t, \tilde{y}_{-i}^{T-t} = x\right\}) \\
&= E[\nu_i | H_t, x_i = x, \tilde{y}_{-i}^{T-t} = x] - E[\nu_i | \left\{H_t, x_i = x, \tilde{y}_{-i}^{T-t} = x\right\}] + P_w(x, \left\{H_t, \tilde{y}_{-i}^{T-t} = x\right\}) \\
&= P_w(x, \left\{H_t, \tilde{y}_{-i}^{T-t} = x\right\}).
\end{aligned}$$

where the second equality comes from the definition of $V_t(\cdot)$ in (20). The third equality follows from $W_{t+1}(x, \left\{H_t, \tilde{y}_{-i}^{T-t} = x\right\}) = 1$ because if at time t trader i's competitor with the highest type has $\tilde{y}_{-i}^{T-t} = x = x_i$, then it must be that at time $t+1$ it is trader i who has the highest type $x_i = x$, and hence he is certain to win the auction and buy the asset in round $t+1$.[21] This also implies that conditioning on $\left\{x_i = x, H_t, \tilde{y}_{-i}^{T-t} = x\right\}$, the expected value of the asset upon winning, $v_w(x, \left\{H_t, \tilde{y}_{-i}^{T-t} = x\right\})$, is equal to $E[\nu_i|, x_i = x, H_t, \tilde{y}_{-i}^{T-t} = x]$. Thus, the fourth and fifth equalities.

To guarantee the first order condition identifies a maximum, we now show that the objective function in (21) is quasi-concave. Let $\Pi(x, z|H_t)$ denote the objective function in (21) shall a trader of type $x_i = x$ behave as if his type was z. Namely:

$$\Pi(x, z|H_t) := \int_{\underline{x}}^{z} (E[\nu_i|H_t, x_i = x, \tilde{y}_{-i}^{T-t} = y] - b_t(y, H_t)) f_t(y) dy + \int_{z}^{\overline{x}} V_{t+1}(x, \left\{H_t, \tilde{y}_{-i}^{T-t} = y\right\}) f_t(y) dy \tag{22}$$

It is sufficient to show that $\frac{\partial^2 \Pi(x,z|H_t)}{\partial z \partial x} > 0$.

[21] Ties in types are excluded from the continuous distribution assumption of types.

Note that

$$\frac{\partial \Pi(x,z|H_t)}{\partial z} = f_t(z)\left((E[\nu_i|x_i=x,H_t,\tilde{y}_{-i}^{T-t}=z]-b_t(z,H_t))-V_{t+1}(x,\left\{H_t,\tilde{y}_{-i}^{T-t}=z\right\}\right), \quad (23)$$

and so

$$\frac{\partial^2 \Pi(x,z|H_t)}{\partial z \partial x} = \left(\frac{\partial E[\nu_i|x_i=x,H_t,\tilde{y}_{-i}^{T-t}=z]}{\partial x} - \frac{\partial V_{t+1}\left(x,\left\{H_t,\tilde{y}_{-i}^{T-t}=z\right\}\right)}{\partial x}\right) f_t(z)$$

Define $\hat{V}_{t+1}(x,y|H_{t+1})$ as the payoff a trader of type $x_i = x$ is expected to make if starting from period $t+1$ on, the trader behaves as if his type was y instead of x. This is

$$\hat{V}_{t+1}(x,y|H_{t+1}) := max_y W_{t+1}(y,H_t)(v_w(x,H_{t+1}) - P_w(y,H_{t+1}))$$

Because the trading game has a finite number T of rounds and the last round is a standard second-price auction (for which truthful bidding is an equilibrium), a backward-induction argument implies that in every subgame starting at $t+1$ it is optimal for a trader of type x to report truthfully from $t+1$ onward. In other words, the maximizer $y^*(x,H_{t+1})$ in the expression above satisfies $y^*(x,H_{t+1}) = x$.[22]

In equilibrium it must be optimal for the trader to choose $y^* = x$, that implies the first order condition $\frac{\partial \hat{V}_{t+1}(x,y|H_{t+1})}{\partial y}\Big|_{y^*=x} = 0$ must be satisfied. Therefore the envelope theorem for $V_{t+1}(x,H_{t+1}) = V_{t+1}(x,y^*(x),H_{t+1})$ is equal to:

$$\frac{d}{dx}V_{t+1}(x,H_{t+1}) = \frac{dV_{t+1}(x,y^*(x),H_{t+1})}{dx} + \underbrace{\frac{dV_{t+1}(x,y^*(x),H_{t+1})}{dy^*(x)}}_{=0} + \frac{dy^*(x)}{dx}$$

Only $v_w(x,H_{t+1}) = \mathbb{E}[\nu_i \mid x_i = x, H_{t+1}]$ depends directly on x, so

[22]Start with the two-period version of the game. In the last period, by construction , truthful reporting is an equilibrium bidding strategy. Hence, in the two-period game, when choosing their first-round bid, $y^*(x,H_{T=2}) = x$ holds. Hence also in the first round truthful reporting is an equilibrium strategy. Now consider the three-period game. The subgame starting in period 2 is exactly the two-period environment just analyzed, so traders bid truthfully from period 2 onward. Therefore, when choosing bids in period 1 of the three-period game, traders again know that $y^*(x,H_{t+1}) = x$. The argument continues iteratively for any finite game.

so

$$\frac{\partial \hat{V}_{t+1}(x, y \mid H_{t+1})}{\partial x} = W_{t+1}(y, H_{t+1}) \frac{\partial E[\nu_i \mid x_i = x, H_{t+1}]}{\partial x}.$$

Evaluating at $y = y^*(x, H_{t+1}) = x$ and at $H_{t+1} = \{H_t, \tilde{y}_{-i}^{T-t} = z\}$ yields

$$\frac{\partial V_{t+1}\big(x, \{H_t, \tilde{y}_{-i}^{T-t} = z\}\big)}{\partial x} = W_{t+1}\big(x, \{H_t, \tilde{y}_{-i}^{T-t} = z\}\big) \frac{\partial E[\nu_i \mid x_i = x, H_t, \tilde{y}_{-i}^{T-t} = z]}{\partial x}. \tag{24}$$

Substituting (24) into the expression for the cross-partial, we obtain

$$\frac{\partial^2 \Pi(x, z \mid H_t)}{\partial z \, \partial x} = \frac{\partial E[\nu_i \mid x_i = x, H_t, \tilde{y}_{-i}^{T-t} = z]}{\partial x} \big(1 - W_{t+1}\big(x, \{H_t, \tilde{y}_{-i}^{T-t} = z\}\big)\big) f_t(z).$$

Thus substituting this last equality in equation (23), we obtain:

$$\frac{\partial^2 \Pi(x, z | H_t)}{\partial z \partial x} = \left(\frac{\partial E[\nu_i | x_i = x, H_t, \tilde{y}_{-i}^{T-t} = z]}{\partial x} (1 - W_{t+1}(x, \{H_t, \tilde{y}_{-i}^{T-t} = z\}) \right) f_t(z)$$

For $z > x$, the above expression is strictly positive because $\frac{\partial E[\nu_i | x_i = x, H_t, \tilde{y}_{-i}^{T-t} = z]}{\partial x} > 0$, $f_t(y) > 0$ and $W_{t+1}(x, \{H_t, \tilde{y}_{-i}^{T-t} = z\})$ is less than 1. For $z < x$ we have $W_{t+1}(x, \{H_t, \tilde{y}_{-i}^{T-t} = z\}) = 1$. As $\frac{\partial \Pi(x, z | H_t)}{\partial z}$ is weakly decreasing in z, the profit function is quasi-concave and hence the f.o.c. identifies a global maximum.[23]

So we have shown that it is optimal for a trader that is active in the buyers room at time t to make a bid equal to $P_w(x, H_{t+1})|_{H_{t+1} = \{H_t, \tilde{y}_{-i}^{T-t} = x\}}$, that is what the trader expects to pay if he buys between $t+1$ and the last round given H_t and $\tilde{y}_{-i}^{T-t} = x$. The event $\tilde{y}_{-i}^{T-t} = x$ however implies that at round $t+1$ trader i is the highest bidder with probability 1, implying the trader is certain to win in round $t+1$ and pay the highest of his competitors' bids at t+1. So we have that $P_w(x, , H_t, \tilde{y}_{-i}^{T-t} = x)$ is equal to the expectation of the highest bid from time $t+1$

[23] Formally, fix a type x and define $g(x; z) := \partial \Pi(x, z \mid H_t)/\partial z$. By the first-order condition we have $g(z; z) = 0$ for all z. Moreover, from the expression for the cross-partial above, $\partial^2 \Pi(x, z \mid H_t)/(\partial z \, \partial x) > 0$ whenever $z > x$ and $= 0$ whenever $z < x$. For any $z > x$,

$$g(z; z) - g(x; z) = \int_x^z \frac{\partial^2 \Pi(\xi, z \mid H_t)}{\partial \xi \, \partial z} \, d\xi > 0,$$

so, since $g(z; z) = 0$, it follows that $g(x; z) < 0$. Hence, for a given x, $\Pi(x, z \mid H_t)$ is strictly decreasing in z for all $z > x$, while for $z < x$ the cross-partial is zero and thus $g(x; z) = g(z; z) = 0$, so $\Pi(x, z \mid H_t)$ is flat to the left of x.

competitors of trader i, given that the strongest competitor type is equal to x. Hence we have

$$b_t(x, H_t) = P_w(x, H_t, \tilde{y}_{-i}^{T-t} = x) = E[\hat{b}_{t+1}|H_t, x_i = x, \tilde{y}_{-i}^{T-t} = x].$$

A symmetric argument applies for the ask auction.

C.2.3 The $t = 0$ period.

To see that $b_0(x) = x$ is an equilibrium, we show that if trader i expects all other traders to truthfully report their own signal at round $t = 0$, and then follow the strategies described above at $t > 0$, then it is a best response for trader i to truthfully report his own signal at $t = 0$. Suppose trader i draws the median signal x^M, if he/she reports truthfully, the payoff is 0. Suppose he reports $b_0 > x_i$ and ends up in the buyers room. Then he knows that i) he will be the buyer with the smallest signal in the room, and ii) because the median report disclosed by the mechanism is larger than true median signal, all bidders' valuation , but his own valuation, will be biased upward. This implies that he is certain not to win any of the bid auctions and thus his expected payoff from deviation is nil. A symmetric argument applies if he by misreporting he ends up in the seller's room. Suppose trader i draws a signal x_i larger than the median signal $x_i > X^M$. By truthfully reporting he will enter the buyer's room. Misreporting leads to two possible scenarios. If his report is larger than X^M, he will not change his own payoff as only the median signal is reported at the end of the round. Thus misreporting has no effect on the trading dynamic and on his own payoff. If his report is not greater than X^M, the trader will have a zero payoff as will be either out of the game at the end of t_0 or will be the trader in the sellers' room with the highest signal and hence will not trade. A symmetric result applies considering a bidder whose truthful report would lead him to the sellers room. (by outbidding there is a higher probability of finishing in the buyer room/ being median, by under-bidding no change in his/her payoff).

C.3 Proof of Corollary 2

We prove the statement by showing that, at time t, the expected bid paid at time $t+1$ conditional on the previous bid , $E[\hat{b}_{t+1}|\hat{b}_t]$, is higher than the realized previous price paid, $\hat{b}_t$ (conversely, the expected ask received conditional on the previous ask, $E[\hat{a}_{t+1}|\hat{a}_t]$, is smaller than the realized previous price received, $\hat{a}_t$). Hence by tower property rules follow that: $E[\hat{b}_{t+1}] > E[\hat{b}_t]$ and $E[\hat{a}_{t+1}] < E[\hat{a}_t]$.

Buyer room. The expression $E[\hat{b}_{t+1}|\hat{b}_t]$ is equal to $E[b_{t+1}(X^{N-t-1})|X^{N-t} = x^{N-t}]$ as conditioning on $\hat{b}_t$ implies that active traders can extract the signal from the price and hence know $X^{N-t} = x^{N-t}$. The condition $X^{N-t} = x^{N-t}$ is equivalent to condition on $Y_{-i}^{N-t} > x^{N-t} = Y_{-i}^{N-t-1}$.Therefore, we can write:

$$E[b_{t+1}(X^{N-t-1})|X^{N-t} = x^{N-t}] = E[b_{t+1}(X^{N-t-1})|Y_{-i}^{N-t} > x^{N-t} = Y_{-i}^{N-t-1}]$$

By the equilibrium submission strategy, the realized price at time t is:

$$\hat{b}_t = b_t(x^{N-t}) = E[b_{t+1}(X^{N-t-1})|X_i = x_i^{N-t}, Y_{-i}^{N-t} = X_i]$$

as the bid setter in stage t (x^{N-t}) bids the expected price in the next round, conditional on his signal being just enough to win the current bid stage. Hence by the properties of affiliation $E[\hat{b}_{t+1}|\hat{b}_t] > \hat{b}_t$.[24]

Seller room. Equivalently, the ask price at $t+1$ conditional on the previous ask is $E[\hat{a}_{t+1}|\hat{a}_t] = E[\hat{a}_{t+1}(X^{t+2})|X^{t+1} = x^{t+1}]$. The condition $X^{t+1} = x^{t+1}$ is equivalent to condition on $Y_{-i}^{t+1} = x^{t+1} > Y_{-i}^{t}$. It follows that:

$$E[\hat{a}_{t+1}(X^{t+2})|X^{t+1} = x^{t+1}] = E[\hat{a}_{t+1}(X^{t+2})|Y_{-i}^{t+1} = x^{t+1} > Y_{-i}^{t}]$$

[24]Consider, for example, the first two bids $\hat{b}_1$ and $\hat{b}_2$. Observing $\hat{b}_1$ is equivalent to know $X^{N-1} = x^{N-1}$ (the second largest signal among all traders). This is equivalent to say that the generic trader i, knows that, among his competitors, there is one signal (X^N) higher than x^{N-1}. On the other hand, at time $t = 1$, the trader with the second largest signal (X^{N-1}) bids the expected next round execution price, conditioning his signal to be equal to the largest signal among all his competitors.

The realized price at time t is obtained by bidding: $\hat{a}_t = E[\hat{a}_{t+1}(X^{t+2})|Y^t_{-i} = x = X^{t+1}]$. Hence by the properties of affiliation $E[\hat{a}_{t+1}|\hat{a}_t] < \hat{a}_t$.

As the inequalities are valid for any possible realizations of $\hat{b}_t$ and $\hat{a}_t$, we can write:

$$E[\hat{b}_{t+1}] = \int E[\hat{b}_{t+1}|\hat{b}_t]d\hat{b}_t \geq \int \hat{b}_t d\hat{b}_t = E[\hat{b}_t] \tag{25}$$

and similarly $E[\hat{a}_{t+1}] \leq E[\hat{a}_t]$.

C.4 Properties Sequential Trading Game

Corollary 5.

1. *In our proposed sequential trading game, it cannot be ensured trading in the last round*
2. *In an interdependent environment, sequential trading mechanisms cannot ensure a single last clearing price*

To prove the corollary, we first provide an important result on budget compatibility for sequential games:

Lemma 1. *In a sequential trading game with interdependent environment, ensuring expected budget balance compatibility in the last trading round is a necessary condition to ensure global expected budget balance compatibility*

Proof Suppose a sequential trading game is **not** expected budget balance compatibility in the last trading round ($E[\hat{a}_T] > E[\hat{b}_T]$). This condition automatically implies that the mechanism is also not expected budget balance compatibility in the previous rounds, as for any $t < T$ the following relationship holds:

$$E[\hat{a}_t] \geq E[\hat{a}_T] > E[\hat{b}_T] \geq E[\hat{b}_t]$$

due to the equilibrium price property $E[\hat{b}_{t+1}] \geq E[\hat{b}_t]$ and $E[\hat{a}_{t+1}] \leq E[\hat{a}_t]$. In general, due to the weakly increasing (decreasing) nature over rounds of buying (selling) prices, if a trading round is not expected budget balance compatibility, all the previous rounds are not expected budget balance compatibility.

Proof Corollary 5.1 In an environment with two traders, with predefined roles and partial uncertainty in their opponent evaluation, the Myerson–Satterthwaite theorem applies. Therefore, as the last round trading prices are not budget compatible ($\hat{a}_T > \hat{b}_T$), I Lemma 1 implies global no expected budget balance compatibility.

Proof Corollary 5.2 Consider first a trading mechanism for which buyers and sellers trade at different prices up to $T-1$ but in the last trading round T there is a single clearing price. Consider the last two periods, as we want a mechanism that satisfies both incentive and budget compatibility, we have the following conditions:

$$\begin{cases} E[\hat{b}_{T-1}] \leq E[\hat{b}_T] = E[\hat{a}_T] \leq E[\hat{a}_{T-1}] \\ E[\hat{b}_{T-1}] \geq E[\hat{a}_{T-1}] \end{cases}$$

where the first condition follows from $E[\hat{b}_{t+1}] \geq E[\hat{b}_t]$ and $E[\hat{a}_{t+1}] \leq E[\hat{a}_t]$ and the second condition ensures budget balance compatibility. The unique solution is given by:

$$E[\hat{b}_{T-1}] = E[\hat{a}_{T-1}]$$

This reasoning applies recursively for any previous round $1 \leq t < T-1$, hence to enforce in expectation budget balance and incentive compatibility, it must hold the following condition

$$E[\hat{b}_t] = E[\hat{b}_T] = E[\hat{a}_T] = E[\hat{a}_t]$$

This condition, the expected price does not change over rounds (hence with information release), is possible only in the case of a private value environment (see Milgrom and Weber 1982), leading to an internal contradiction.

Finally, consider a trading mechanism in which in every round there is a single clearing price (think of a repeated MPT game, for example). In this case, the game satisfies budget compatibility by construction as buyers and sellers trade at the same price in each round, also in this

case it must hold that $E[\hat{p}_t] = E[\hat{p}_T]$. Indeed if the condition were such that $E[\hat{p}_t] \geq E[\hat{p}_T]$ ($E[\hat{p}_t] \leq E[\hat{p}_T]$) buyers (sellers) would have an incentive to misreport at time t.

C.5 Speed of the Sequential Mechanism.

Given N traders, the sequential trading mechanism (STM) ensures that $N-3$ traders trade (the traders who do not trade are the median trader and the two adjacent traders). In our baseline model, we have shown results for an STM in which only one unit is traded at a time. However, the proof and equilibrium strategies do not change if multiple units are traded in each round. Formally, let the number of traded stocks be

$$K = \frac{N-1}{2} - 1,$$

so that the STM satisfies $\sum_\tau r_\tau = K$, where r_k is the number of stocks traded in round τ.

A natural question is what happens when we change the number of stocks traded in each round. Suppose that in each round the same number of stocks is traded. Formally, define the number of stocks traded per round as $k \in \mathbb{N}^+$ such that $\frac{K}{k} \in \mathbb{N}^+$, and hence the general number of trading rounds as $\frac{K}{k}$. In the main body of the text, we consider $k = 1$ and hence the number of traded stocks is equal to the number of trading rounds ($K = \frac{N-1}{2} - 1 = T$).

The average number of bids known to traders before trading in round $t \in [1, \frac{K}{k})$ is then

$$\#H \;=\; \frac{\sum_{l=0}^{K/k-1}(1+2kl)}{K/k} \;=\; K-k+1, \tag{26}$$

where $l = 0$ corresponds to the first round of trading, when only the median signal is revealed. As the STM is incentive compatible, traders are able to extract from each bid the signal and decrease price uncertainty.

Equation (26) implies that when $k = 1$ (our baseline case), the average number of signals available before trading is maximized. More generally, $\#H$ is strictly decreasing in k. Hence, by *slowing down the market* (smaller k), information precision increases. [25]

[25] If the number of units traded per round is not constant, it is generally impossible to obtain a closed-form solution for arbitrary N, as the result depends on the exact allocation of trades. However, the STM mechanism

At the opposite extreme, when $k = K$, the STM becomes equivalent to a standard K-double auction, in which traders on both sides bid

$$b = E[V \mid X_i = x_i,\, X^{M+1} = x_i], \qquad s = E[V \mid X_i = x_i,\, X^{M-1} = x_i],$$

respectively. In this case, the spread is minimal because equilibrium prices are

$$\hat{b} = E[V \mid X_i = x^{M+1}, X^{M+1} = x^{M+1}], \qquad p^s = E[V \mid X_i = X^{M-1}, X^{M-1} = x^{M-1}].$$

Generally speaking, as k increases, the STM becomes a k-sequential double-price auction. The spread decreases in k, because the marginal buyer and seller are closer together in signal space.

C.6 Two-Rounds Market - Minimization Ex Post Regret

In this section, we consider a simultaneous game that minimizes ex-post regret but at the cost of substantially reducing trading volume. To contemporaneously maximize the ex-ante information set of market participants and ensure incentive compatibility, in this mechanism only two market participants can trade. We formally define the trading mechanism:

Single Trading (ST) mechanism

- In the starting round, $t = 0$, each trader i simultaneously submits a price p_i. Let $P \in \{\mathbb{R}^+\}^N$ denote the initial *ordered* price profile. The mechanism publicly discloses all traders' prices, except the two highest and the smallest bids. The $N - 4$ traders whose price has been disclosed cannot trade.

- The two traders with the largest (smallest) bid at $t = 0$ run a second price buy (sell) auction. The market sells at $\hat{b}$ and buys at $\hat{a}$, with $\hat{b} > \hat{a}$ and pockets the difference. Trading always occur.

The mechanism minimizes ex-post regret by exacerbating two features of the baseline design. First, only agents with the strongest incentives to trade are active: the lowest-valuation agent

and its equilibrium do not rely on this simplifying assumption.

(seller) and the highest-valuation agent (buyer). By Assumption A.2, these two types are ex post least likely to regret their trading decisions. Second, when deciding whether to trade, each observes the signals of all traders except the signal of its competitor and the top two on the opposite side of the market. Therefore, out of the N fragmented signals, each active trader conditions on $N-3$ signals. This is, ex ante, the largest information set consistent with incentive compatibility and budget balance. Full disclosure of all N signals would induce the highest- and lowest-valuation agents to profitably misreport so as to manipulate both the clearing price and the counterparty's beliefs. Retaining the second price auction type ensures the mechanism can rely on competitive bidding for price formation. Equilibrium bidding then follows the standard logic of second-price auctions with interdependent values (Milgrom and Weber, 1982).

It is interesting to study what happens when instead of releasing $N-4$ signals, the market releases $N-2$ signals: hence in the second round, when actual trading happens, there is no competitive bidding to determine the trading prices. In this case, the market has to rely on excluded traders to provide potential trading prices as we assume the market does not know traders' evaluation function. Therefore, the market has to subsidize at least one excluded trader to provide quotes, leading to our next corollary:

Corollary 6. *In the ST mechanism, the lack of competitive bidding in the last stage implies only expected budget compatibility.*

Intuitively, as the market needs to pay the price maker independently of whether trading happens or not (otherwise, the trader would have an incentive to misreport to increase the likelihood of trading) and trading under this regime is only in probability as active traders' information sets differ from the ones who set the price. Therefore, conditional on trading, the trading venue can incur in some loss.

Proof. In case $N-2$ signals are released at $t=0$, the price maker has to submit a bid, $\hat{b}$, and ask price, $\hat{a}$, such that the bid (ask) uses as private signal the highest (smallest) signal revealed as private value to ensure incentive compatibility at time $t=0$.

The price maker can for example submit a bid (ask) equal to the expected value of the asset of the trader with the highest (smallest) signal revealed:

- $\hat{b} = v(x^{N-1}, \{x^i\}_{i=2}^{N-2}, E[x^1|\{x^i\}_{i=2}^{N-1}], E[x^N|\{x^i\}_{i=2}^{N-1}])$
- $\hat{a} = v(x^2, \{x^i\}_{i=3}^{N-1}, E[x^1|\{x^i\}_{i=2}^{N-1}], E[x^N|\{x^i\}_{i=2}^{N-1}])$

Clearly these bids ensure incentive compatibility in the two stage game and $\hat{b} > \hat{a}$ in virtue of the order preserving property. However, the valuations of the traders still in game are:

- $E[\nu^N|\{x^i\}_{i=2}^{N-1}] = v(x^N, \{x^i\}_{i=2}^{N-1}, E[x^1|\{x^i\}_{i=2}^{N}])$
- $E[\nu^1|\{x^i\}_{i=2}^{N-1}] = v(x^1, \{x^i\}_{i=2}^{N-1}, E[x^N|\{x^i\}_{i=1}^{N-1}])$

The two traders still in actions are using $N-1$ signals to determine the expected missing private value (their own private signal still not revealed and $N-2$ known signals). As the expected value of conditional on public information ($E[x^N|\{x^i\}_{i=1}^{N-1}]$) may be bigger than the actual realized value x^N (known only to the potential buyer), it may happen that the suggested bid is larger than the potential buyer expected value (similar reasoning applies to the sell side). As an example, consider 5 traders: by the argued equilibrium we have that the proposed bid is $\hat{b} = v(x^4, \left\{x^3, x^2, E[x^1|\{x^i\}_{i=2}^{4}], E[x^5|\{x^i\}_{i=2}^{4}]\right\})$, whereas the expected value of the active potential buyer is $v(x^5, \left\{x^4, x^3, x^2, E[x^1|\{x^i\}_{i=2}^{5}]\right\})$: depending on the specific signal distribution and evaluation function it may happen that the potential buyer refrains from trading.

Alternative price schedule that ensure interim incentive compatibility have similar problems. As firstly noticed in Milgrom and Weber (1982), when the information set across traders differ, characterize equilibria that ensure both incentive compatibility and truthful reporting becomes challenging.

D Extensions

D.1 Number of Traders is Even

In this Appendix, we relax the assumption that the number of traders (N) is necessarily odd. When N is even, let x^{M^-} and x^{M^+} denote the lower and upper median values, respectively—i.e., the two central order statistics in the sorted profile.

Lemma 2. *In a MPT mechanism with private values and an even number of traders, buyers buy at x^{M^+} and sellers at x^{M^-}.*

Proof Consider the buyers' side. Suppose the trader with signal x^{M^+} deviates and bids $\tilde{x}$ such that $\tilde{x} > x^{M^+ + 1}$. By changing his position in the order, the trader incurs a loss equal to $x^{M^+ + 1} - x^{M^+} > 0$. The argument on the sellers' side follows by symmetry. The presence of a bid-ask spread, $x^{M^+} > x^{M^-}$, when the number of traders is even, is consistent with the logic that x^{M^+} and x^{M^-} represent, respectively, the lower- and upper-bound marginal valuations for buyers and sellers. This immediately follows from the fact that no single value divides the two sides of the market evenly when N is even.

Suppose instead that trading venue defines a unique equilibrium price

$$p_0^M = \alpha x^{M^+} + (1 - \alpha) x^{M^-}, \quad \text{with } \alpha \in (0, 1).$$

In this case, for certain values of α, the trader with valuation x^{M^+} may have an incentive to participate as a buyer, since the resulting payoff

$$x^{M^+} - \tilde{p}_0^M > 0$$

can be strictly positive. This occurs in particular when $\alpha \to 0$, as $\tilde{p}_0^M \to x^{M^-}$. Conversely, as $\alpha \to 1$, $\tilde{p}_0^M \to x^{M^+ + 1}$, and the payoff becomes negative, removing the incentive to deviate.

A symmetric argument applies on the sellers' side. Consider a trader with valuation x^{M^-} who may have an incentive to misreport in order to sell at a price

$$\tilde{p}_0^M = \alpha x^{M^+} + (1 - \alpha) x^{M^- - 1}.$$

The incentive to misreport increases as $\alpha \to 1$, since $\tilde{p}_0^M \to x^{M^+}$, which may exceed the trader's true valuation. As these conditions have to be simultaneously satisfied, it is impossible to define a single clearing price.

Lemma 3. *In a MPT mechanism with interdependent values and an even number of traders, buyers buy at* $g(x^{M^+}, x^{M^+})$ *and sellers at* $g(x^{M^-}, x^{M^-})$.

Proof The proof follows immediately from the private value case.

Lemma 4. *In the sequential game with an even number of traders, both* x^{M^+} *and* x^{M^-} *are announced at* $t = 0$.

Proof Consider the evolution of the game backwards. After the signals x^{M^+} and x^{M^-} are announced, the protocol sorts traders into two rooms, each with an equal number of agents, ensuring that buyers and sellers can be matched. Within each room, buyers and sellers compete in a second-price auction. Therefore, a trader's payoff depends only on the number of rounds in which they are eligible to participate, not on x^{M^+} and x^{M^-}.

At time $t = 0$, traders with signals different from x^{M^+} and x^{M^-} have no incentive to misreport. Any deviation results in a zero payoff—either because they are excluded at $t = 0$ or because they are reassigned to the opposite room, where they do not trade in equilibrium. In contrast, truthful reporting guarantees a weakly positive payoff.

A trader with signal exactly equal to x^{M^+} or x^{M^-} obtains a payoff of zero regardless of their report at $t = 0$: either they are excluded from trading at $t = 0$, or - if they remain in the game - they are always outbid (underbid) in the trading rounds.

Lemma 5. *The Two-Rounds Market is not affected by the number of traders*

Proof Since the buy (sell) price corresponds to a second price auction on both sides, the number of revealed signals—whether odd or even—does not alter the strategic structure of the game.

D.2 Multi–Quantity Median Price Trading (MQ–MPT)

In this Appendix, we relax the unit-trade restriction and allow traders to submit *multiunit demand* with decreasing marginal values. The mechanism extends the baseline Median Price

Trading logic by combining (i) a median-based split into buyer and seller rooms and (ii) a multiunit allocation and pricing rule within each room.

General setting. Consistent with the baseline trading environment, we assume there are N (odd) traders and each trader i receives a private signal $x_i \in [\underline{x}, \overline{x}]$ drawn from a continuous distribution. Differently from the baseline model, trader i's valuation for q units is

$$v(x_i, q) = x_i \sum_{\ell=1}^{q} f(\ell),$$

where x_i is the private signal and $f(\cdot)$ is a non-negative and non-increasing function. We standardize $f(1) = 1$ so that bidding for one quantity is equal to the private signal itself. The function captures the concepts that players have a decreasing marginal utility by increasing the number of stocks they own.

It follows that the valuation function is bounded by:

$$x_i \leq v(x_i, q) = x_i \sum_{l=1}^{q} f(l) \leq x_i q$$

the lower-bound follows from the non-negative assumption ($\forall l \geq 2, f(l) = 0$), whereas the upper-bound follows from the non-increasing assumption ($\forall l \geq 2, f(l) = 1$). Both bounds can be handled by the MPT mechanism proposed in the main body of the paper. The lower-bound is our classic environment as any marginal quantity after the first one has zero evaluation. The upper bound is a risk neutral environment where instead of trading 1 unit, players submit orders of q units.

The Multi–Quantity Median Price Trading (MQ–MPT) protocol. We propose a two-stage mechanism in which traders (i) first are split among buyers and sellers, (ii) each trader can submit orders of up to $q = \frac{N-1}{2} - 1$ units:

- In the starting round, $t = 0$, each trader i simultaneously submits a price p_i to the trading platform. Traders reporting above the median are assigned to the buyer room; those below

are assigned to the seller room. Each room therefore contains $q+1$ traders. The trader who submits the median price exits the game.

- In the second round, $t=1$, within each room, each trader reports a valuation vector for q units. The trading venue allocates the q units to the traders with the highest (smallest) marginal q evaluations. Winning bidders pay the relevant displaced (competing) marginal bids (equivalently, the externality pricing rule).

- Let P^B denote total payments collected from buyers and P^S total transfers paid to sellers. Trade is executed whenever total payments collected from buyers exceed: total transfers to sellers is $P^B \geq P^S$.

Reminiscent of the proposed mechanisms for interdependent values, the market venue cashes in the difference between buyers' and sellers' total payments. Economically, this profit is the wedge between the demand–side and supply–side marginal externality prices. Because scarcity and strategic displacement are priced separately in the two rooms, these shadow prices need not coincide. So consistent with our analysis on the main text, when the marginal valuation that separates buyers and sellers is not unique, a natural spread arises and the market maker steps in to ensure incentive compatibility. The multiunit extension requires an additional restriction: tradable quantity must be bounded relative to market size. This condition is novel and necessary to ensure a truthful equilibrium.

Proposition 5. *Consider a private-values environment with N traders and let the order size available to traders be $q=\frac{N-1}{2}-1$. In the Multi-Quantity Median Price Trading mechanism (MQ-MPT), truthful reporting of the (marginal) valuation vector at stage $t=1$,*

$$\overline{p}_i = \big(x_i, x_i f(2), \ldots, x_i f(q)\big)$$

is a weakly dominant strategy for each trader i. Moreover, trade always occurs and the resulting allocation is ex post efficient. The mechanism is ex post incentive compatible, ex post individually rational, and ex post budget balanced.

Proof. The proof is by backward induction.

- At $t = 1$. In each trading room, the Vickrey multi object auction has a truthful reporting equilibrium with an efficient allocation (Krishna, 2009).

- At $t = 0$. Traders anticipate that truthful reporting is an equilibrium outcome in the next round. Hence, any trader with a signal equal or below the median knows that, by overbidding, may be assigned to the buyer room. By inflating the price submitted at $t = 0$, however, each trader enters a room formed by $q = \frac{N-1}{2} - 1$ traders with strictly higher initial signals. Since truthful reporting remains an equilibrium strategy in stage 1, exactly q units are allocated to the q highest bids. Therefore, overbidding at stage 0 yields zero trading probability: by construction there are at least q bids above the deviator's highest bid, so the mis-reporting trader cannot win any unit. An analogous argument applies to deviations into the seller room.

  - As the marginal function is common across agents, i.e. $f_i(\ell) \equiv f(\ell)$ for all i and ℓ, with $f(\ell) > 0$ and non–increasing and every trader assigned to the buyer room has a strictly larger signal than every trader assigned to the seller room, it follows that

$$\min_{i \in B} x_i = \underline{x}^B \; > \; \max_{j \in S} x_j = \overline{x}^S. \tag{Sep}$$

  Since bids in stage 1 are truthful and marginal values take the form $x_i f(\ell)$, this implies the point-wise dominance condition

$$x_i f(\ell) \; > \; x_j f(\ell), \qquad \forall\, x_i > x_j, \; \forall\, \ell \leq q.$$

  Now recall that each room contains $q + 1$ traders and exactly q units are allocated. In a q–unit VCG (multi–unit Vickrey) auction with $q+1$ participants, it follows that the player with the smallest (highest) signal in the buyer (seller) room never trades in equilibrium. Since payments in the VCG auction are computed from *competing* marginal bids (replacement bids in the relevant counterfactual allocations), the marginal bids of the excluded trader in each room can, at best, enter as pricing bids; if they do not enter, this implies

that the pricing bids used are even larger on the buyer side (and even smaller on the seller side). Therefore, the following condition holds

$$P^B \geq \sum_{l=1}^{q} \underline{x}^B f(l) > \sum_{l=1}^{q} \overline{x}^S f(l) \geq P^S,$$

The venue's clearing condition is always satisfied and trade occurs with certainty. Therefore, once players are within each trading room, there is no incentive to misreport their valuation vector to influence trading.[26]

■

As probability of trading is endogenous, differently from the standard environment, where N sellers can buy q quantity from a fixed and inelastic supplier, we here need to impose two conditions. First, analogous to the mechanisms proposed in the main text, the mechanism separates buyers from sellers as otherwise players would have the incentive to influence the other side of the market. Second, the mechanism imposes a bound on the tradable quantity in each room, limiting it to strictly fewer units than the number of traders in each room. This condition is crucial for incentive compatibility along two dimensions: *room selection* at stage $t = 0$, and *quantity reporting* at stage $t = 1$.

The condition eliminates profitable manipulations of the initial split at $t = 0$. If a trader misreports at $t = 0$ to change role, then—under truthful stage–$t = 1$ reporting—there are q traders with higher (smaller) relevant marginal values who exhaust the q available trades, so the deviator does not trade.[27] Moreover, trading q units with $q + 1$ traders in each room facilitates

[26]The mechanism allows to consider heterogeneous function across traders, namely f_i, conditional that the decay on marginal utilities for player with higher signals is not faster than for players with smaller signals ($x_i f_i(\ell) > x_j f_j(\ell), \quad \forall x_i > x_j$). This condition is necessary and sufficient to ensure point-wise dominance. Therefore, if either a buyer or the median player were trying to misreport to enter the seller room, the player would always face q traders with $x_j f_j(q)$ marginal values lower than him. Therefore, the mis-reporting trader would never trade, making his deviation not profitable. In the buyer room, the condition is always satisfied due to the non increasing assumption of f.

[27]Suppose there are 7 traders and the marginal evaluation function f is such that we have the following evaluation matrix:

$$\begin{pmatrix} 7 & 3.5 & 1.75 \\ 6 & 3 & 1.5 \\ 5 & 2.5 & 1.25 \\ 4 & 2 & 1 \\ .... & ... & .... \end{pmatrix} \tag{27}$$

clearing. It implies that total payments collected from buyers exceed total transfers to sellers ($P^B > P^S$). Therefore, trade occurs with certainty and in each room market participants do not have incentive to strategic bid shading to satisfy the clearing condition.

The analysis highlights a simple thickness–manipulation tradeoff. When the market is *thick* (many participants), relatively large quantities can be accommodated without creating profitable incentives to distort the market, because competition on each side is sufficient to absorb the order while preserving truthful behavior. By contrast, in *thin* markets (few participants), large orders relative to the number of traders amplify market power and create scope for manipulation (e.g. via misreporting to affect the room composition or clearing condition). In such environments, imposing tighter quantity caps—or, more generally, restricting order size relative to market depth—can be necessary to preserve incentive compatibility and robust clearing.

The thickness–manipulation tradeoff highlighted by the mechanism is closely related to standard insights in market microstructure. In liquid markets with substantial depth, large orders can be absorbed by a sufficiently competitive set of counter-parties, so that no single trader can profitably distort prices or allocations through strategic misreporting. By contrast, in thin markets where participation is limited, large orders relative to market depth generate substantial price impact and create scope for manipulation, as traders may benefit from repositioning themselves across sides of the market or from influencing the clearing condition. This provides a mechanism-design rationale for the order-size limits, and liquidity safeguards commonly observed in practice, especially in closing auctions and other batch trading environments.

Now consider the median trader. If three units were traded, he would have an incentive to misreport $\tilde{x}_4 > 5$ at t_0 (thereby making $x_5 = 5$ the median trader). In the second round, even though he would then bid truthfully $x_4 = 4$ (by incentive compatibility of the Vickrey mechanism), he would obtain strictly positive profit. Indeed, the highest bid vector is $b = \{7, 6, 4\}$. Therefore, excluding trader 4, the highest competing bid vector is $pb_3^{-4} = \{7, 6, 3.5\}$. Since $4 > 3.5$ (while $2 < 6$), trader 4 can buy one unit and pays 3.5, yielding a strictly positive payoff. By contrast, if only two units are traded, the highest bid vector is $b = \{7, 6\}$, which eliminates any incentive to misreport at t_0. A similar argument applies to the seller room