Algorithmic Trading Architecture and Price Informativeness

Snehal Banerjee and Martin Szydlowski*

August 2026

Abstract

We develop a model of algorithmic trading in which a better-informed algorithm competes with human traders who choose whether to acquire costly information and participate in the market. We compare a *model-free* algorithm, which learns how to trade based on realized profits alone, with a *model-based* algorithm that learns from a model of its price impact. We show that model-based learning induces more aggressive trading, raises the algorithm's profits, and can reduce price informativeness by crowding out human information acquisition. This implies that mandates for "interpretable" algorithms can lead to worse market outcomes. Restrictions on algorithmic trading intensity or access to information can improve price informativeness when the resulting increase in human participation outweighs the loss of information from algorithmic trading.

JEL Classification: D82, D83, G12, G14, G18

Keywords: Algorithmic trading, Reinforcement learning, Human-AI interaction, Interpretable algorithms, Price informativeness

*Banerjee (snehalb@umich.edu) is at the University of Michigan - Ann Arbor and Szydlowski (mszydlowski@ust.hk) is at the Hong Kong University of Science and Technology.

1 Introduction

Algorithmic trading allows investors to acquire and act on information more cheaply and at greater scale, and now accounts for the majority of trading volume across markets.[1] The impact of such trading on market outcomes, specifically price informativeness, is not well understood. Common wisdom suggests that the increased presence of better-informed algorithmic traders should improve price informativeness (e.g., Grossman and Stiglitz, 1980). However, the empirical evidence appears mixed. On the one hand, Weller (2018) shows that algorithmic trading deters costly information acquisition and reduces the information content of prices. On the other hand, Bonelli and Foucault (2023) argue that the use of alternative data sources (e.g., satellite data on retail parking lots) improves price informativeness, while Katona, Painter, Patatoukas, and Zeng (2025) argue that access to such data does not immediately improve price discovery.

Understanding the market impact of algorithmic trading is also important for regulatory policy. A common feature of recent regulatory proposals is a preference for "interpretable" algorithms, whose behavior can be understood and audited, over "black-box" systems (e.g., Rudin, 2019). MiFID II requires algorithmic trading strategies to be documented, tested, and auditable; the SEC has proposed rules governing the use of predictive data analytics by broker-dealers; and the EU Artificial Intelligence Act imposes documentation and transparency obligations on providers of AI systems.[2] Evaluating the impact of such policies requires a framework for understanding how algorithmic trading affects market outcomes and how these effects depend on the algorithm's architecture.

We model a repeated, strategic trading environment in which an algorithmic trader with superior information trades against human traders who decide whether to acquire costly, noisy information. We compare trading by a **model-free** algorithm that learns only from its realized profits to that by a **model-based** algorithm that learns from an explicit, but misspecified, model of price impact. We show that the model-based algorithm trades more aggressively, earns higher profits, is more likely to crowd out human traders, and consequently can make prices *less* informative. Since the model-based algorithm is more interpretable, our analysis implies that regulations that mandate "interpretable" algorithms can lead to worse outcomes by harming human trading profits and decreasing price informativeness.

[1] See, for example, Hendershott, Jones, and Menkveld (2011), Chaboud, Chiquoine, Hjalmarsson, and Vega (2014), Brogaard, Hendershott, and Riordan (2014), Abis (2020), and Boehmer, Fong, and Wu (2021).

[2] See Article 17 of MiFID II (Directive 2014/65/EU) and the accompanying technical standards (RTS 6, Commission Delegated Regulation (EU) 2017/589); the SEC's proposed rule on Conflicts of Interest Associated with the Use of Predictive Data Analytics (Exchange Act Release No. 34-97990, 2023); and the EU Artificial Intelligence Act (Regulation (EU) 2024/1689).

Model. Our framework is a repeated version of the single-period Kyle (1985) model. In each period, a number of risk-neutral human traders decide whether to pay a cost to acquire noisy private information about the asset's payoff and participate in the market. The algorithmic trader has superior information about the payoff and submits trades according to a policy it learns over time. A risk-neutral, competitive market maker observes the aggregate order flow from human, algorithmic, and noise traders, and sets the price as the conditional expectation of the asset payoff.

The assumption that the algorithm is better informed than human investors reflects two complementary features of algorithmic trading at proprietary trading desks and hedge funds. First, algorithms can act not only on traditional financial market data, but also on alternative data that human analysts may not be able to process at scale, including transaction records, web traffic, geolocation data, satellite imagery, and text. Second, they can bring far more computing power to bear on that data, and so extract informative, tradable signals. Market participants and regulators alike are concerned about how the lack of strategic sophistication embedded in such strategies can lead to overcrowded trades and excessive volatility.[3]

We consider two benchmark architectures for how the algorithm learns to trade. Under the **model-free** architecture, the algorithm uses an actor-critic framework and learns the optimal trading intensity from realized profits alone.[4] Specifically, it perturbs its trading intensity to identify which values increase realized profits. We show that the steady-state trading intensity corresponds to the best response to the market maker and human traders. In essence, the model-free algorithm learns how to play the Bayes-Nash equilibrium, even though it does not have a formal model of other market participants.

Under the **model-based** architecture, the algorithm conjectures a parametric model of how its trades move prices and updates its estimate of price impact, or λ, from realized outcomes. In this case, we show that in the steady state, the algorithm correctly estimates λ. However, the algorithm does not account for the presence of human traders that trade on correlated information, and so trades more aggressively than under the model-free archi-

[3]For example, see "Hedge funds and high-frequency traders are converging," *Financial Times* (September 29, 2025), which describes how algorithmic traders underestimated the overlap in their strategies. The resulting losses in the summer of 2025 echo the "quant quake" of August 2007. Also, see "Trading algorithms bring benefits but fears of accidents grow," *Financial Times* (June 1, 2016) for an earlier discussion of similar concerns.

[4]Actor-critic is one of the foundational algorithms in reinforcement learning (e.g., Sutton and Barto (2018), Ch. 13.5). It is suited for continuous action spaces, unlike Q-learning which requires a discrete action space (e.g., Sutton and Barto (2018), Ch. 6.5). The algorithm consists of an "actor" who learns the policy (e.g., via a gradient method), and a "critic" who maintains an estimate of a value function (e.g., via a temporal difference method). It is model-free since it requires no knowledge of the underlying environment. This means the algorithm does not know the distribution of asset values, how its own trades affect its profits, how prices are formed, or that it faces a market maker and other traders.

tecture.[5] Notably, the misspecified model under the model-based architecture gives rise to an endogenous form of commitment to more aggressive trading. We show that this *commitment effect* is profitable: the model-based algorithm earns strictly higher profits than the model-free algorithm and leaves human traders strictly worse off.[6]

Market Impact of Algorithmic Trading. The impact of algorithmic trading on market outcomes, and price informativeness in particular, depends on two opposing forces. First, consider an increase in the algorithm's trading intensity, holding fixed the number of human traders. Human traders trade less aggressively and earn lower average profits. Moreover, since the algorithm trades on more precise information, we show that price informativeness unambiguously increases. We refer to this channel as the *intensity effect*.

With endogenous human participation, a sufficiently large increase in the algorithm's trading intensity leads to crowding out: when their expected profits fall below the cost of acquiring information, human traders do not participate in the market. We refer to this channel as the *participation effect*. Each exit leads to a discrete drop in price informativeness, but a discrete jump in expected profits for the algorithm and the remaining human traders.

As a result, price informativeness exhibits a saw-tooth pattern in algorithmic trading intensity: because human traders exit at discrete thresholds, price informativeness increases between thresholds (intensity effect) but discretely drops at the thresholds (participation effect). Importantly, the relative impact of these forces depends on the algorithm's architecture. Because the model-based algorithm trades more aggressively, it tends to crowd out more human traders. We provide sufficient conditions under which (i) the presence of an algorithm leads to less informative prices relative to an economy with only human traders, and (ii) the model-based architecture leads to lower price informativeness than the model-free architecture.

Policy Implications. Our framework provides a natural benchmark to evaluate a number of policy proposals.[7]

[5] As we discuss further in Section 4.2, even if the algorithm's model specification allowed for the possibility of trading by other participants, experimenting with its own trades would not reveal it.

[6] While Bernhardt and Boulatov (2023) and Yan, Yang, Zhang, and Zhou (2024) show that commitment is not valuable in the single-trader Kyle model with normally distributed payoffs, it can be valuable in settings with multiple traders and non-normal payoffs (e.g., Banerjee, Malikov, and Prokopev (2024b)). However, as we discuss in Section 5.1, it is unlikely that traders can credibly commit to trading intensities in practice since trading is unobservable and anonymous. Our results suggest, however, that the use of model-based algorithms might be a practical approach to doing so.

[7] Our analysis focuses primarily on price informativeness as the policy target. Welfare, which we can measure as aggregate profits across all market participants net of information costs, is driven completely by aggregate information acquisition costs. This is because all traders and the market maker are risk-neutral, and so trading is zero-sum, i.e., aggregate expected trading profits are zero. In turn, because the model-based architecture tends to reduce participation by humans, it would always be preferred by a policymaker who only wishes to maximize welfare.

First, interpretability mandates can lead to worse market outcomes. In our setting, the model-based algorithm is the interpretable one, since its estimate of price impact has a transparent meaning that can be communicated to, and audited by, a regulator. In contrast, the model-free algorithm is a black box. A mandate that requires algorithms to be interpretable therefore *selects* the model-based architecture. This leads to higher profits for the algorithmic trader, lower aggregate profits and more crowding out for human traders, and can lower price informativeness, even though it makes the algorithm easier to understand.

Second, policies that aim to level the informational playing field across human and algorithmic traders can have very different consequences. For instance, policies that subsidize access to information for human traders (e.g., the SEC's Market Data Infrastructure rules) lower their cost of participation and tend to unambiguously lead to higher price informativeness. In contrast, restricting the information available to algorithmic traders (e.g., privacy regulation to restrict the collection of alternative data) operates through both the intensity and participation effects. It reduces price informativeness by reducing the intensity and informativeness of algorithmic trading, but may increase price informativeness by inducing more participation by human traders. For a regulator whose objective is price informativeness, this asymmetry favors subsidizing human traders over restricting the algorithm.

Third, caps on trading intensity (e.g., through limits on position sizes for algorithmic traders) can have nuanced implications.[8] While standard intuition would suggest that such caps should reduce price informativeness by restricting trading by the (better informed) algorithm, they may have the opposite effect by inducing increased participation by human traders. Our analysis also suggests that, all else equal, trading caps or position limits might be preferable to transaction taxes since the efficacy of the latter depends on whether a model-based architecture accounts for the tax when specifying its model of trading profits.

Overview. The rest of the paper is organized as follows. Section 2 discusses the related literature. Section 3 presents the model and discusses important assumptions, and presents preliminary analysis that characterizes the trading equilibrium for an arbitrary algorithmic trading intensity. Section 4 describes the two learning architectures and establishes steady state outcomes. Section 5 analyzes the crowding out of human traders and price informativeness, and compares the two architectures. Section 6 discusses the policy implications, and Section 7 concludes. All proofs are in Appendix A, Appendix B considers extensions to allow for (i) noisy information for the algorithm and (ii) transaction taxes, and Appendix C contains simulation results establishing convergence to steady state strategies.

[8] In Appendix B.3, we show that transaction taxes that target algorithmic traders have similar implications, while a blanket transaction tax unambiguously lowers price informativeness.

2 Related literature

Our paper contributes to a number of strands of literature.

Learning algorithms in financial markets. The most closely related papers study how algorithmic traders interact with other market participants. Banerjee and Szydlowski (2025) analyze a continuous-time economy in which Q-learners trade against rational, risk-averse investors and show how feedback between realized returns and the Q-learners' future demand can endogenously generate stochastic volatility, return predictability, and rents from risk sharing. In contrast, our paper focuses on how trading by a better-informed algorithm affects the behavior of strategic human traders. The two papers provide complementary analyses of human-algorithm interaction: the former focuses on asset pricing and welfare implications in a competitive (price-taking) setting, while the current paper focuses on implications for price informativeness and regulatory policy in a setting in which traders have price impact.

In related work, Dou, Goldstein, and Ji (2025a) study algorithms with explicit planning modules that coordinate trading over multiple periods, create and exploit price bubbles, and interact with positive-feedback traders. Dou, Goldstein, Li, and Yang (2025b) study competition between better-informed but boundedly rational humans and AI investors that learn from prices. They show that sophisticated humans can outperform AI because the latter learn from market data generated by average human behavior. Our analysis is complementary. In our setting the algorithm is better informed but strategically naive, while human traders are rational but must pay to acquire information. This reflects the informational advantage of proprietary trading desks and hedge funds, which combine market data with alternative sources at a scale and speed human analysts cannot match.

An earlier set of papers uses simulations to study interactions among algorithmic market participants. Colliard, Foucault, and Lovo (2022) show that algorithmic market makers in a setting based on Glosten and Milgrom (1985) adapt to adverse selection but charge markups above competitive price. Moreover, they find that markups increase and prices become less competitive as adverse selection decreases. In a similar setting, Guarino, Jehiel, and Symons-Hicks (2025) show that algorithmic quotes range from loss-free to more competitive prices depending on the design and calibration of the market maker's Q-learning algorithm. Gufler, Sangiorgi, and Tarantino (2025) show that interactions among multiple deep-learning traders create a negative learning externality through exploration noise. We complement this work by providing an analytic characterization of trading between algorithmic and rational traders, who have different degrees of information.

More closely related is Dou, Goldstein, and Ji (2023), who show that multiple informed Q-learning speculators can autonomously sustain collusive outcomes. Their algorithms earn supra-competitive profits by trading *less* aggressively on their information, either through price-trigger strategies or homogenized learning biases. Our analysis complements this work by focusing on the interaction between algorithmic traders and strategic human traders, who can exit the market. The (model-based) algorithm in our setting earns higher expected profits by trading *more* aggressively on its information and crowding out human traders.[9]

Algorithm design. Our distinction between model-free and model-based learning is related to Barberis and Jin (2023), who combine the two learning systems to explain investor behavior, and to Asker, Fershtman, and Pakes (2022), who show that pricing outcomes depend on whether algorithms evaluate counterfactual actions. In earlier work, Routledge (1999) establishes conditions under which adaptive learning converges to rational-expectations demands in a financial market. Consistent with this benchmark, our model-free algorithm learns to play the Bayes-Nash equilibrium in steady state. However, we also show that a model-based architecture can lead to a steady state strategy that is more aggressive.

Misspecified learning and commitment. The steady state of our model-based algorithm is related to learning under misspecified models (e.g., Esponda and Pouzo, 2016; Heidhues, Kőszegi, and Strack, 2018; Fudenberg, Lanzani, and Strack, 2021). It is also related to models of strategic delegation, in which committing to an objective other than profit maximization can be advantageous (e.g., Vickers, 1985; Fershtman and Judd, 1987). In our setting, the algorithm does not intentionally choose a commitment strategy. Instead, commitment to aggressive trading arises endogenously from the interaction between its learning rule and its representation of the pricing rule.

Strategic trading and participation. More generally, our model builds on Kyle (1985) and the literature on strategic trading with multiple informed traders (e.g., Holden and Subrahmanyam (1992), Foster and Viswanathan (1993), Tighe and Michener, 1994, Foster and Viswanathan (1996)). The endogenous participation by human traders follows the tradition of costly information acquisition (e.g., Grossman and Stiglitz (1980), Verrecchia (1982)) and, in strategic settings, Banerjee and Breon-Drish (2020); the discrete entry structure and associated multiplicity are standard from the analysis of entry with fixed costs (Bresnahan and Reiss, 1990).

[9]In our setting, trading less aggressively would invite more entry by human traders and lower profits (e.g., see Section 6.3).

Figure 1: Timeline

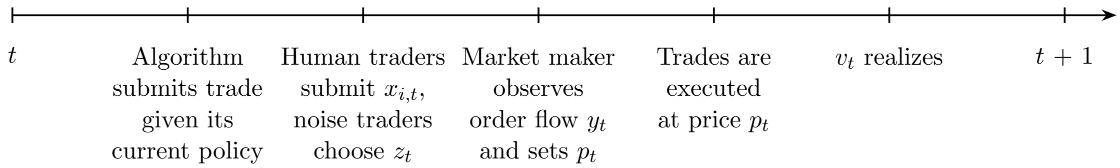

3 Model setup

Our framework is based on a repeated version of the single-period Kyle (1985) model. Time is discrete and indexed by $t = \{0, 1, 2, \dots\}$. The market participants trade a sequence of one-period assets. At the beginning of period t, the asset trades at a price p_t; at the end of period t, the asset pays off $v_t \sim \mathcal{N}(0, \sigma_v^2)$, where v_t is independently distributed over time. The sequence of events within period t is summarized in Figure 1.

There are $\bar{K}$ risk-neutral human traders indexed $i = 1, \dots, \bar{K}$, one algorithmic trader, and noise traders. At the beginning of each period t, each trader i decides whether to pay cost $c > 0$ to participate in the market and acquire a signal $s_{i,t}$, where

$$s_{i,t} = v_t + \varepsilon_{i,t} \quad \text{and} \quad \varepsilon_{i,t} \sim \mathcal{N}(0, \sigma_\varepsilon^2).$$

We refer to c as the (per-period) cost of information acquisition, and let K_t denote the number of active human traders at date t. For notational convenience, we define the signal-to-noise ratio of these signals as

$$\tau_H \equiv \frac{\sigma_v^2}{\sigma_v^2 + \sigma_\varepsilon^2}.$$

Conditional on acquiring the signal, trader i submits a trade $x_{i,t}$ to maximize expected profits in the current period net of costs:

$$\mathbb{E}\left[x_{i,t}(v_t - p_t)\right] - c.$$

The algorithmic trader observes v_t perfectly at the beginning of period t, and submits trade $x_{A,t}$ to maximize the expected present value of profits

$$\mathbb{E}\left[\sum_{t=0}^{\infty} \gamma^t x_{A,t}(v_t - p_t)\right],$$

where $\gamma \in (0, 1)$ is the discount factor. The noise traders submit an aggregate trade $z_t \sim \mathcal{N}(0, \sigma_z^2)$. We assume that $\{v_t, \varepsilon_{i,t}, z_t\}$ are independent across each other and over time.

There is a risk-neutral market maker who observes aggregate order flow $y_t = x_{A,t} + \sum_{i=1}^{K_t} x_{i,t} + z_t$ and sets the price of the date-t risky asset equal to her conditional expectation:

$$p_t = \mathbb{E}_t \left[v_t \right].$$

We focus on linear equilibria in which (i) the equilibrium trade by investor i is given by $x_{i,t} = b_t \, s_{i,t}$, (ii) the equilibrium trade by the algorithmic trader is given by $x_{A,t} = \beta_t v_t$, and (iii) the market maker's pricing rule is linear in the order flow i.e.,

$$p_t = \lambda_t y_t,$$

where λ_t corresponds to the equilibrium price impact, or Kyle's lambda, and (iv) K_t human traders participate in the market.

3.1 The algorithmic trader's architecture

The algorithmic trader uses a reinforcement-learning algorithm to update its trading strategy from past experience. We consider two possible **architectures**, which we denote by $a \in \{A, M\}$:

- Under the *model-free* architecture ($a = A$), the algorithm learns from realized profits alone, using an actor-critic scheme. It maintains no representation of the market it trades in.
- Under the *model-based* architecture ($a = M$), the algorithm conjectures a parametric model of how its own trades move prices, learns the model's parameters from realized prices, and chooses its policy to maximize profits under the conjectured model.

Section 4 describes both architectures formally and characterizes their steady states. We show that these architectures lead to different steady state strategies and, consequently, have different implications for expected profits, price informativeness and welfare.

3.2 Discussion of assumptions

Repeated static structure. Payoffs, signals, and noise trades are independent across periods, so the trading game is effectively static: the repetition matters only because it gives the algorithm a stationary environment in which to learn. This assumption is made for tractability because it allows us to characterize the algorithm's behavior in steady state and facilitates the welfare analysis. One can interpret this as an approximation of a more complex environment (e.g., with additional information flows, serial dependence in fundamentals and noise trading) which may still yield stationary equilibria, after appropriately conditioning on the relevant state variables.

Short-term traders. We assume that human traders are myopic for tractability, which abstracts from repeated-games type complications (see Dou et al. (2025b) for a similar assumption).[10] For instance, long-term human traders could engage in tacit collusion, as they interact with each other repeatedly over time, or trade strategically to manipulate the evolution of the algorithm. While such dynamics are potentially interesting, a complete analysis is beyond the scope of the current paper and is therefore left for future work.

Sophistication of market participants. Our model features an algorithmic trader who is better informed but must trade against rational human traders who understand the strategic environment better. Given the ability of trading algorithms to collect, process and trade on vast amounts of information in a timely manner, we believe this is an empirically relevant setting to study. Moreover, the assumption justifies the use of algorithmic trading in equilibrium.[11]

The informational advantage of algorithms is also consistent with recent empirical evidence. Zhu (2019) shows how the increased use of consumer-transaction and satellite data raises price informativeness and predicts fundamentals ahead of disclosure. Katona et al. (2025) document that satellite coverage of retail parking lots predicts earnings surprises, and that the gains accrue to the sophisticated investors who buy the data. Abis (2020) finds that quantitative funds process far broader information sets than discretionary funds, and Bonelli and Foucault (2023) argue that big data reduces the value of traditional human expertise. Coleman, Merkley, and Pacelli (2022) show that automated Robo-analysts revise recommendations more frequently, incorporate information that human analysts overlook, and make recommendations that outperform human analysts. Grennan and Michaely (2021) document the rapid entry of equity market intelligence financial technology firms supplying machine-generated analysis that aggregates traditional and non-traditional data sources and uses artificial intelligence to make investment recommendations.[12]

In our main analysis, the algorithmic trader observes v_t perfectly, which gives it an unambiguous informational advantage over human traders. This makes the result that algorithmic

[10] The resulting setting is equivalent to having a sequence of short-lived human traders who arrive in each period, decide whether or not to enter and trade, and then leave.

[11] This is in contrast to existing work on human-AI interaction (e.g., Banerjee and Szydlowski (2025), Dou et al. (2025a), Dou et al. (2025b)) in which algorithms only generate trading profits at the expense of less sophisticated or constrained human traders.

[12] Relatedly, Cao, Jiang, Yang, and Zhang (2023) show that firms adjust disclosure once machines are the primary readers, which itself presumes an audience with a processing advantage.

trading can reduce price informativeness more stark. Appendix B.2 relaxes this assumption and considers an algorithm that observes a noisy signal of the payoff. We show that the qualitative results are unchanged provided the algorithm's signal is more precise than the human traders' signals.

3.3 Preliminary analysis

The focus of our analysis will be on the steady state equilibrium for different architectures of the algorithm. To facilitate this analysis in Section 4.3, in this subsection, we present a characterization of the stationary equilibrium as a function of an *arbitrary strategy* β chosen by the algorithmic trader, in terms of stationary functions $\{b(\beta), \lambda(\beta), K^*(\beta)\}$.

Since information and payoffs are independent over time and v_t is mean zero, we have that

$$p_t = \mathbb{E}_t\left[v_t\right] = \mathbb{E}[v_t|y_t].$$

Moreover, since payoffs and signals are jointly normal and mean zero in each period, this implies that the price is given by $p_t = \lambda y_t$, where

$$\lambda \equiv \frac{\mathbb{C}(v_t, y_t)}{\mathbb{V}(y_t)} = \frac{\Gamma\sigma_v^2}{\Gamma^2\sigma_v^2 + Kb^2\sigma_\varepsilon^2 + \sigma_z^2}, \qquad \text{and} \qquad \Gamma \equiv Kb + \beta. \tag{1}$$

Conditional on paying c to acquire the signal, human trader i chooses $x_{i,t}$ to maximize:

$$\pi_{H,t} \equiv \mathbb{E}[(v_t - p_t)x_{i,t}|s_{i,t}] = x_{i,t}\big(\mathbb{E}[v_t|s_{i,t}] - \lambda\,\mathbb{E}[y_t|s_{i,t}]\big).$$

Since $s_{i,t} = v_t + \varepsilon_{i,t}$ and all noise terms are independent, $\mathbb{E}[v_t|s_{i,t}] = \mathbb{E}[s_{j,t}|s_{i,t}] = \tau_H s_{i,t}$. This implies:

$$\pi_{H,t} = x_{i,t}\left(\tau_H s_{i,t} - \lambda\left(\left(\Gamma - b\right)\tau_H s_{i,t} + x_{i,t}\right)\right),$$

since $\Gamma - b = (K-1)b + \beta$. Rearranging the first-order condition and matching terms with the conjectured strategy $x_{i,t} = bs_{i,t}$ implies that b satisfies

$$b = \frac{1 - \lambda\left(\Gamma - b\right)}{2\lambda}\tau_H. \tag{2}$$

Solving the equations (1) and (2) yields the following characterization for the trading equilibrium.

Proposition 1. *Given architecture $a \in \{A, M\}$, suppose the algorithmic trader submits a trade $x_{A,t} = \beta v_t$ and the number of participating human traders is fixed at K. Then, there exists a linear equilibrium, which is characterized by human trading intensity $b(\beta)$ and price impact $\lambda(\beta)$, given by:*

$$b(\beta) = \frac{-\left(2-\tau_H\right)\beta\sigma_v^2 + \sqrt{4K\sigma_z^2\tau_H\sigma_v^2 + \left(2-\tau_H\right)^2\beta^2\sigma_v^4}}{2K\sigma_v^2}, \tag{3}$$

and

$$\lambda(\beta) = \frac{\tau_H}{b(\beta)\left(2+(K-1)\tau_H\right)+\tau_H\beta}. \tag{4}$$

The expected (per-period) profit for human traders is given by:

$$\Pi_H(\beta; K) \equiv \mathbb{E}[\pi_{H,t}] = \frac{b(\beta)^2\sigma_v^2}{b(\beta)\left(2+\left(K-1\right)\tau_H\right)+\beta\tau_H}, \tag{5}$$

and the expected (per-period) profit for the algorithmic trader with architecture $a \in \{A, M\}$ is given by:

$$\Pi(\beta; K) \equiv \mathbb{E}[\pi_{a,t}] = \frac{\beta\, b(\beta)\left(2-\tau_H\right)\sigma_v^2}{b(\beta)\left(2+(K-1)\tau_H\right)+\tau_H\beta}$$

Moreover, b and Π_H are decreasing in β and K.

The nature of the equilibrium is intuitive. More aggressive trading by the algorithm (higher β) leads to less aggressive trading by human traders and lower expected profits. Similarly, higher competition among human traders (i.e., higher K) leads each investor to trade less aggressively. However, aggregate trading intensity by humans increases (i.e., Kb increases) and, as a result, expected profits decrease.

The above result implies that expected profits for an individual human trader are maximized for $K = 1$ and when there is no algorithmic trader in the market. In this case, one can express the expected profit as:

$$\Pi_H(0; K=1) = \frac{\sqrt{\tau_H}\,\sigma_v\sigma_z}{2}. \tag{6}$$

In line with Kyle (1985), the expected profit for the human trader increases with the signal-to-noise ratio τ_H, the prior uncertainty about payoffs σ_v and the volatility of noise trading σ_z. To ensure participation by human traders, we make the following assumption for the remaining analysis.

Assumption 1. *The cost of information acquisition is bounded above as follows:*

$$c < \frac{\sqrt{\tau_H}\,\sigma_v \sigma_z}{2}.$$

Finally, to rule out boundary cases, we assume that $\bar{K}$ is sufficiently large so that $\Pi_H(0, \bar{K}) < c$. That is, the market cannot bear all $\bar{K}$ human traders.

4 Algorithmic trading and learning

This section describes how the algorithmic trader learns to trade. We first present the model-free architecture and then the model-based architecture. We then characterize the steady states of both, and establish that the model-free algorithm replicates a rational trader while the model-based algorithm trades strictly more aggressively.

4.1 The model-free architecture

Under the model-free architecture ($a = A$), the algorithmic trader uses an actor-critic algorithm,[13] which is a foundational reinforcement learning algorithm. The algorithm is split between the "actor," who learns the policy, and the "critic," who learns a performance criterion. In each iteration, the critic's performance criterion is used to guide the actor's learning. In our model specifically, the actor learns the trading intensity $\beta_{A,t}$ and the critic learns the expected value from trading.

The actor-critic algorithm has no knowledge of the environment. It does not know the distribution of asset values, the relation between its own trading strategy and profits, its own price impact, or the existence of the market maker and other traders. As we will see in Section 4.3, the actor-critic algorithm still learns to play a best response to the market maker's and human traders' strategies despite this lack of knowledge.

We now explain the specifics of the algorithm.

Perturbation Noise. The algorithm perturbs its trading intensity around $\beta_{A,t}$. Specifically, the algorithmic trader's demand is given by $x_{A,t} = \tilde{\beta}_{A,t} v_t$, where

$$\tilde{\beta}_{A,t} = \beta_{A,t} + \varepsilon_{A,t}.$$

[13]See e.g., Sutton and Barto (2018), Ch. 13.5.

The perturbation noise $\varepsilon_{A,t} \sim \mathcal{N}(0, \sigma_A^2)$ is i.i.d. and independent of $(v_t, \varepsilon_{i,t}, z_t)$.[14] This noise is necessary for the algorithm to learn the marginal value of changing the trading intensity.[15]

Critic. The critic maintains an estimate of the value function, and updates it using a temporal difference method.[16] The estimate is given by the parametric approximation

$$Q_t(\tilde{\beta}_{A,t}) = (\tilde{\beta}_{A,t} - \beta_{A,t})w_t + q_t, \tag{7}$$

with parameters w_t and q_t, which approximate the gradient and expected value, respectively. Intuitively, by varying the policy around $\beta_{A,t}$, the algorithm attempts to form a local estimate of the gradient. The functional form in Equation (7) is a *compatible* function approximation and ensures that the algorithm learns an unbiased estimate of the gradient.[17] The critic updates the Q-values taking into account the profit surprise

$$\delta_t = \pi_{A,t} + \gamma Q_t\left(\beta_{A,t}\right) - Q_t(\tilde{\beta}_{A,t}), \tag{8}$$

where $\pi_{A,t} = x_{A,t}(v_t - p_t)$ is the algorithmic trader's realized profit and where $\gamma \in (0, 1)$ is a discount factor. Specifically,

$$w_{t+1} - w_t := \alpha \delta_t \frac{\partial Q_t(\tilde{\beta}_{A,t})}{\partial w_t} = \alpha \delta_t (\tilde{\beta}_{A,t} - \beta_{A,t}) \tag{9}$$

and

$$q_{t+1} - q_t := \alpha \delta_t \frac{\partial Q_t(\tilde{\beta}_{A,t})}{\partial q_t} = \alpha \delta_t. \tag{10}$$

Here, $\alpha > 0$ is the learning rate, which governs how much the algorithm updates its parameters. Intuitively, δ_t is the value updating term for the Q-values[18] and it directly enters the updating equation for the value estimate q_t. The gradient estimate w_t is updated in Equation (9) based on how the profit surprise δ_t interacts with the perturbation in trading intensity. When profits are higher than expected ($\delta_t > 0$) and simultaneously the policy deviates upwards ($\tilde{\beta}_{A,t} > \beta_{A,t}$), the gradient estimate increases ($w_{t+1} > w_t$).

[14] See Lillicrap, Hunt, Pritzel, Heess, Erez, Tassa, Silver, and Wierstra (2015) for a similar procedure in the context of deep reinforcement learning. We will call $\beta_{A,t}$ the "trading intensity" throughout for simplicity. In the following analysis, we characterize the small noise limit where $\sigma_A \to 0$, i.e., the perturbations vanish.

[15] See Silver, Lever, Heess, Degris, Wierstra, and Riedmiller (2014), Sec. 4.2.

[16] See Sutton and Barto (2018), Ch. 6.5.

[17] See Silver et al. (2014), Th. 3. The choice of function approximation may bias the algorithm's learning, so that the algorithm cannot locally identify the gradient. The parametric form in Equation (7) ensures this is not the case. Note also that since the trading intensity β_A is continuous, the Q-matrix in a standard Q-learning algorithm would be infinite-dimensional. Using a parametric function approximation avoids dealing with infinite-dimensional objects.

[18] See specifically Sutton and Barto (2018), Ch. 6.1, Eq. 6.5.

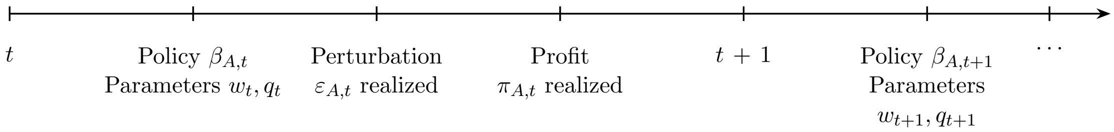

Figure 2: Actor-Critic Update

Actor. The actor uses a deterministic policy gradient method to learn the policy.[19] Given the approximation Q_t, the actor updates the trading intensity $\beta_{A,t}$ as follows:

$$\beta_{A,t+1} - \beta_{A,t} := \alpha \frac{\partial Q_t(\tilde{\beta}_{A,t})}{\partial \tilde{\beta}_{A,t}} = \alpha w_t. \tag{11}$$

Specifically, whenever the gradient estimate w_t is positive, the actor adjusts the trading intensity upwards.

Overall, an actor-critic algorithm consists of a parameter vector $\Psi_t = (\beta_{A,t}, w_t, q_t)$, updating Equations (9)–(11), which describe the evolution of the algorithm for any history of prices $\{p_s\}_{s=1}^t$, asset payoffs $\{v_s\}_{s=1}^t$, and perturbations $\{\varepsilon_{A,s}\}_{s=1}^t$, and hyperparameters $\alpha > 0$, $\gamma \in (0,1)$, and $\sigma_A > 0$, which denote the learning rate, the discount factor, and the perturbation noise. Figure 2 illustrates the updating across time.

Definition 1. A *steady state* of the algorithm is given by $\Psi^* = (\beta_A^*, w^*, q^*)$, such that

$$\mathbb{E}[\Psi_{t+1} - \Psi_t \mid \Psi_t = \Psi^*] = 0.$$

Conditional on having reached the steady state (β_A^*, w^*, q^*), the algorithm on average does not update the policy $\beta_{A,t}$ or the parameters (w_t, q_t). Equations (9)–(11) imply that in a steady state

$$w^* = 0, \quad \mathbb{E}[\delta_t(\tilde{\beta}_{A,t} - \beta_A^*)|\Psi_t = \Psi^*] = 0, \text{ and } q^* = \frac{\mathbb{E}[\pi_{A,t}|\Psi_t = \Psi^*]}{1-\gamma}. \tag{12}$$

In a steady state, the trading intensity is characterized by a first-order condition as the slope estimate is zero ($w^* = 0$), and q^* correctly estimates the discounted profit. The term $\mathbb{E}[\delta_t(\tilde{\beta}_{A,t} - \beta_A^*)|\Psi_t = \Psi^*] = 0$ implies that in the steady state, the derivative estimate w^* is correct on average, since the profit surprise is uncorrelated with perturbations in trading intensity.

While the perturbations $\varepsilon_{A,t}$ allow the algorithm to learn the marginal value of $\beta_{A,t}$, they also introduce noise into the trading intensity. As a result, the algorithmic trader behaves in part like a noise trader. In the following analysis, we focus on the *small noise limit* when

[19] See Silver et al. (2014).

the perturbations vanish, i.e., when $\sigma_A \to 0$.[20] This eliminates the noise component in the demand, but preserves the algorithm's learning.

4.2 The model-based architecture

A model-based algorithm ($a = M$) instead incorporates basic knowledge of the environment. In each period, the algorithmic trader knows its own demand $x_{M,t}$, and observes both the realized value v_t and the price p_t. Thus, the algorithmic trader can infer that the profit is given by

$$\pi_{M,t} = x_{M,t}(v_t - p_t), \tag{13}$$

and does not use an algorithm to learn how prices, asset values, and demand affect profits.[21] However, without detailed understanding of the environment, the algorithmic trader may not know how its own $x_{M,t}$ impacts the price p_t and, more generally, the strategic environment (e.g., the existence or number of other traders and the market maker).

We therefore assume that the algorithmic trader uses its own trading and price data to learn about its price impact. The algorithm is model-based, in that it conjectures a parametric form for the price:[22]

$$\widehat{p}_t = p_{0,t} + \widehat{\lambda}_t x_{M,t}. \tag{14}$$

Here, $p_{0,t}$ represents the trader's estimate of the average price and $\widehat{\lambda}_t$ represents its estimate of the price impact. Beyond the profit function (13) and the price conjecture (14), the algorithmic trader has no knowledge of the environment.[23] The algorithm learns the parameters $(p_{0,t}, \widehat{\lambda}_t)$ and the policy $\beta_{M,t}$ directly by observing prices p_t and asset values v_t.

Similar to Section 4.1, the algorithm uses a perturbation around the trading intensity $\beta_{M,t}$, so that $\tilde{\beta}_{M,t} = \beta_{M,t} + \varepsilon_{M,t}$, where $\varepsilon_{M,t} \sim \mathcal{N}(0, \sigma_M^2)$ is i.i.d. and independent of all $(v_t, \varepsilon_{i,t}, z_t)$. Given the parametric form in Equation (14), the profit estimate $Q_t(\tilde{\beta}_{M,t})$ becomes

$$Q_t(\tilde{\beta}_{M,t}) = \tilde{\beta}_{M,t} v_t (v_t - p_{0,t} - \widehat{\lambda}_t \tilde{\beta}_{M,t} v_t).$$

Unlike in Section 4.1, the algorithm no longer uses the profit surprise δ_t in Equation (8) to update the parameters of the Q-values. Instead, it uses its prior knowledge about the profit

[20]See e.g., Silver et al. (2014), Sec. 3.3 for a similar approach in reinforcement learning.

[21]By comparison, in the model-free algorithm, the algorithmic trader does not exploit this knowledge and instead lets the algorithm learn the profit via Equation (7).

[22]Similar approaches are found in Weretka (2011), where traders have a linear model of their own price impact, and Kyle, Obizhaeva, and Wang (2018), where traders conjecture that *other* traders use a linear trading rule, which then leads each trader to have a particular linear model of price impact.

[23]That is, as in Section 4.1, the algorithm does not know the distribution of asset values or the existence of the market maker or other traders.

(Equation (13)) and its price conjecture (Equation (14)).

The algorithm learns the parameters $(p_{0,t}, \widehat{\lambda}_t)$ via updating equations similar to Equations (9)–(10), i.e.

$$p_{0,t+1} - p_{0,t} = \alpha \left(p_t - \widehat{p}_t \right), \tag{15}$$

and

$$\widehat{\lambda}_{t+1} - \widehat{\lambda}_t = \alpha \left(p_t - \widehat{p}_t \right) (\tilde{\beta}_{M,t} - \beta_{M,t}) v_t. \tag{16}$$

Finally, the algorithm updates $\beta_{M,t}$ using a gradient method similar to Equation (11),

$$\beta_{M,t+1} - \beta_{M,t} = \alpha \frac{\partial Q_t(\tilde{\beta}_{M,t})}{\partial \tilde{\beta}_{M,t}} = \alpha v_t (v_t - p_{0,t} - 2\widehat{\lambda}_t \tilde{\beta}_{M,t} v_t). \tag{17}$$

Overall, the model-based algorithm consists of a parameter vector $\Psi_{M,t} = (\beta_{M,t}, p_{0,t}, \widehat{\lambda}_t)$, parametric conjectures for the profit and price in Equations (13) and (14), updating Equations (15), (16), and (17), and hyperparameters $\alpha > 0$ and $\sigma_M > 0$. A steady state $\Psi_M^* = (\beta_M^*, p_0^*, \widehat{\lambda}^*)$ of the algorithm satisfies

$$\mathbb{E}[\Psi_{M,t+1} - \Psi_{M,t} \mid \Psi_{M,t} = \Psi_M^*] = 0. \tag{18}$$

4.3 Equilibrium characterization across architectures

This subsection defines our equilibrium concept and characterizes the equilibrium under each architecture. Note that, under either architecture, the algorithm must perturb its policy to learn, which introduces noise into its demand. In what follows, we characterize the limit as these perturbations vanish and focus on the resulting steady state. This allows us to tractably characterize the equilibrium since, in the limit, we are able to leverage the assumption that aggregate order flow and fundamentals are jointly normal in equilibrium. In Appendix C, we use simulations to establish that the algorithm's trading strategy converges to the steady state limit that we characterize analytically.

Definition 2. Given an architecture $a \in \{A, M\}$, the strategies $(\beta_a^*, b_a^*, \lambda_a^*, K_a^*)$ are an **algorithmic trading** equilibrium if (1) β_a^* is a steady state trading intensity of architecture a in the small noise limit, (2) each human trader's strategy b_a^* is a best response to λ_a^*, β_a^*, K_a^*, and other human traders' strategies, (3) λ_a^* solves Equation (1) given b_a^* and β_a^*, and (4) K_a^* satisfies

$$\Pi_H(\beta_a^*, K_a^*) \geq c > \Pi_H(\beta_a^*(K_a^*+1), K_a^*+1), \tag{19}$$

where $\beta_a^*(K)$ is the algorithm's equilibrium trading intensity conditional on K human traders participating.

In equilibrium, the algorithm is in a steady state (i.e., it on average does not update the policy β_a^*), and this steady state is consistent with the market maker's and human traders' strategies. Given the algorithm's trading intensity β_a^*, human traders and the market maker play a standard Bayes-Nash equilibrium.

The entry condition (19) implies that human traders will drop out if expected trading profits do not exceed the fixed cost c, until (1) the profit of the remaining traders is positive and (2) entry by any trader currently not participating in the market would be unprofitable. Condition (19) assumes that a trader contemplating entry can anticipate how the equilibrium will change following that trader's entry.[24]

We next characterize the equilibrium under the model-free architecture.

Proposition 2. *In an algorithmic trading equilibrium with the model-free architecture, β_A^* is given by*

$$\beta_A^* = \frac{1 - b_A^* \lambda_A^* K_A^*}{2\lambda_A^*}, \tag{20}$$

b_A^* and λ_A^* are given by Equations (3) and (4), and K_A^* satisfies Equation (19), with $\beta = \beta_A^*$.

Equation (20) is the algorithmic trader's best response to b_A^* and λ_A^*. Thus, the model-free algorithm learns the trading rate that maximizes its profit given the market maker's and human traders' strategies — even though, being model-free, it does not know that it faces a market maker or other traders, or what their strategies are.[25] Specifically, the algorithmic trader scales back its trading intensity to account for the aggregate price impact of human traders in the economy, as is reflected in the $b_A^* \lambda_A^* K_A^*$ adjustment in the steady-state trading intensity β_A^*.

The next result characterizes how the steady state changes under the model-based architecture.

Proposition 3. *In an algorithmic trading equilibrium with the model-based architecture, β_M^* is given by*

$$\beta_M^* = \frac{1}{2\lambda_M^*}, \tag{21}$$

b_M^* and λ_M^* are given by Equations (3) and (4), and K_M^* satisfies Equation (19), with $\beta = \beta_M^*$.

[24] From the literature on entry in oligopoly (e.g., Bresnahan and Reiss (1990)), it is well known that fixed costs lead to multiple equilibria. In particular, the equilibrium is indeterminate in whether firm i or firm i' stays in the market if the market cannot accommodate both. However, such indeterminacy does not affect the profits of the remaining traders, the algorithmic trader's profit, price informativeness, or traders' incentives to enter. We therefore focus the following analysis on the quantities which are determinate.

[25] Note that since the algorithmic and human traders play best responses and the market maker is rational, the algorithmic trading equilibrium is identical to the Bayes-Nash equilibrium of the game in which the algorithmic trader is replaced by a rational strategic trader with the same information.

The above result highlights that the choice of architecture can have important implications for the equilibrium outcomes. As described above, in the steady state, the model-free algorithm recognizes that human traders act on correlated information and scales back its trading accordingly. In contrast, the model-based algorithm's conjectured price model attributes all systematic price variation to its own trades, so its learning rule never makes this adjustment. As a result, its steady state intensity $\beta_M^* = \frac{1}{2\lambda_M^*}$ exceeds the best response in Equation (20).

Importantly, note that even if the model-based algorithm's specification allowed for trading by other participants, experimenting with its own trades (as in Section 4.2) would not reveal it. The price variation generated by correlated trading is orthogonal to the algorithm's perturbations and is absorbed into the residual it attributes to noise. Separating it out requires the prior hypothesis that other agents trade on the same fundamental, not more data. Moreover, no specification would let the algorithm learn how others' strategies respond to its own, i.e., the nature of $b(\beta)$: order flow is anonymous and its perturbations are transitory, so human traders never react to them. We extend the model-based algorithm in Appendix B.1, where we explicitly take into account uncertainty about other traders' strategic behavior. There, we show that generically the model-based algorithm trades more aggressively than the model-free algorithm.

5 Market impact of algorithmic trading

In this section, we characterize the impact of algorithmic trading on market outcomes. We begin by characterizing how human participation and expected profits differ across architectures. We then characterize sufficient conditions under which an algorithmic trader with an arbitrary trading intensity β can crowd out human traders, and what the implications are for price informativeness.

5.1 Trading profits and market participation

We begin by characterizing how trading profits vary across the two architectures.

Proposition 4. *(1) Suppose the number of active human traders is exogenously fixed at $K \geq 1$. Then we have $\beta_M^* > \beta_A^*$, $\Pi(\beta_M^*; K) > \Pi(\beta_A^*; K)$, and $\Pi_H(\beta_A^*; K) > \Pi_H(\beta_M^*; K)$.*

(2) Now suppose K_a^* denotes the equilibrium number of human traders under architecture $a \in \{A, M\}$. Then we have $K_M^* \leq K_A^*$, and $\Pi(\beta_M^*; K_M^*) \geq \Pi(\beta_A^*; K_A^*)$ with strict inequality in expected profits whenever $K_A^* \geq 1$.

Figure 3: Profit rankings across algorithmic architectures

The figure fixes the number of active human traders, K, and compares equilibrium expected profits under the model-free and model-based architectures. Panel (a) shows that the model-based algorithm earns strictly higher expected profits, $\Pi(\beta_M^*; K) > \Pi(\beta_A^*; K)$, for every $K \geq 1$. Panel (b) shows that each human trader earns strictly higher expected profits against the model-free algorithm, $\Pi_H(\beta_A^*; K) > \Pi_H(\beta_M^*; K)$. Parameters are $\tau_H = 0.5$ and $\sigma_v = \sigma_z = 1$.

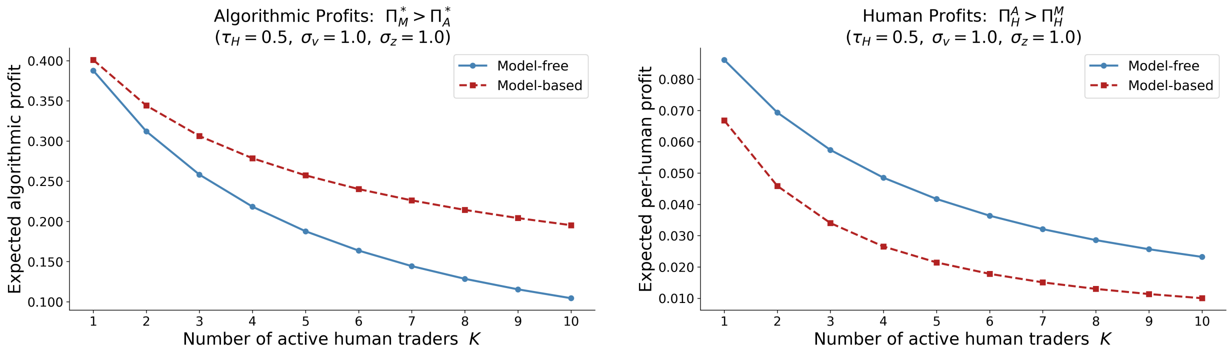

(a) Algorithmic profits

(b) Per-human profits

Part (1) of the above result characterizes how market outcomes vary with architecture for an exogenously fixed K, and Figure 3 provides an illustration. Intuitively, the model-based algorithm's misspecification functions as a commitment device which leads it to trade more aggressively (i.e., $\beta_M^* > \beta_A^*$). In equilibrium, the human traders best respond to the algorithm's trading intensity, and so trade less aggressively against the model-based algorithm, which yields higher profits for the algorithm (i.e., $\Pi(\beta_M^*; K) > \Pi(\beta_A^*; K)$), and lower profits for human traders (i.e., $\Pi_H(\beta_M^*; K) < \Pi_H(\beta_A^*; K)$). We refer to this as the *commitment effect*. The logic parallels strategic delegation in oligopoly, where committing to an objective other than profit maximization confers an advantage (e.g., Vickers, 1985, Fershtman and Judd, 1987). Here, the commitment arises endogenously from the algorithm's architecture.

Since the model-based architecture yields strictly higher profits for any fixed number of human traders, it is natural to ask what the ranking is once human participation is endogenous. Part (2) of the above result establishes that this result survives: the model-based architecture induces weakly more exit, and exit further raises the algorithmic trader's profit. Intuitively, the *commitment effect* benefits the algorithmic trader twice. Holding participation fixed, the model-based algorithm extracts higher profits from the same set of rivals. In addition, the lower profits it leaves to human traders push more of them below their participation threshold, and each exit removes a competitor and raises the algorithm's expected profit further.

Note that in practice, it is unlikely that market participants can credibly commit to a trading intensity because orders are anonymous and trading is simultaneous. However, our results above suggest that an algorithmic trader *may* be able to effectively commit to trading more aggressively by adopting a model-based architecture. The aggressive trading arises not from an explicit ability to commit, but endogenously from learning by the algorithm.

5.2 Crowding out by algorithmic trading

In this subsection, we derive sufficient conditions under which the algorithmic trader crowds out all human traders in equilibrium. The analysis is in terms of an arbitrary trading intensity β, which nests the model-free steady state (β_A^*) and the model-based steady state (β_M^*). When all human traders are crowded out, the two architectures coincide (both reduce to the Kyle monopolist), so the sufficient conditions for full crowding out below apply to either architecture.

First, consider an equilibrium in which only the algorithmic trader is in the market and chooses the steady state trading intensity β_0^*. The following result characterizes the expected profit in this equilibrium.

Lemma 1. *When there are no human traders, the steady-state equilibrium under either architecture is characterized by*

$$\beta_0^* = \frac{\sigma_z}{\sigma_v}, \quad \lambda_0^* = \frac{1}{2}\frac{\sigma_v}{\sigma_z}, \quad \textit{and} \quad \Pi(\beta_0^*) = \frac{1}{2}\sigma_v\sigma_z.$$

Note that this follows immediately from recognizing that $\beta_M^* = \beta_A^* = \frac{1}{2\lambda^*}$ when $K = 0$ and that the resulting steady state equilibrium corresponds to the single-player Kyle (1985) equilibrium.

Next, note that Proposition 1 establishes that for an arbitrary algorithmic trading intensity β, the expected profits for human traders are decreasing in the number of humans K and in the intensity of algorithmic trading β. This implies the following result.

Lemma 2. *There exists a unique threshold $\underline{\beta}$ such that human traders participate iff $\beta \leq \underline{\beta}$, where $\underline{\beta}$ satisfies:*

$$\Pi_H(\underline{\beta}; K = 1) = c$$

Moreover, $\underline{\beta}$ is decreasing in c and increasing in σ_z.

Finally, we can ensure that entry by any human trader is not profitable if the expected profit $\Pi_H(\beta_0^*; K = 1) < c$. The following result characterizes a sufficient condition for this to arise.

Proposition 5. *There are no human traders in the market if $c > c_A$ under the model-free architecture and $c > c_M$ under the model-based architecture, where*

$$c_A \equiv \frac{\tau_H \sigma_v \sigma_z}{(4-\tau_H)\sqrt{\tau_H + (2-\tau_H)^2}} \quad \textit{and} \quad c_M \equiv \frac{\tau_H \sigma_v \sigma_z}{4\sqrt{4-\tau_H}}, \tag{22}$$

and $c_M < c_A$.

Intuitively, when the trading intensity of the algorithmic trader is sufficiently large, the algorithmic trader crowds out human traders from the market. The thresholds in (22) above ensure that the expected profit from acquiring the signal and trading optimally against the algorithmic trader with intensity $\beta_a^*(K = 1)$ under architecture $a \in \{A, M\}$ is lower than the cost of information c.

Figure 4 provides an illustration of how expected profits and human trader participation vary with algorithmic trading intensity β. Panel (a) shows that algorithmic profits are initially increasing in β, maximized at β_0^*, and then decreasing. Panel (b) shows that total human profits (dashed line) decrease with β. However, average human profits (solid line) are non-monotonic: Π_H decreases in β when the number of human traders K is fixed, but spikes up when a human trader exits the market. These exits correspond to discrete drops in total human profits ($K\Pi_H$) and discrete jumps in algorithmic profits.

Figure 4: Expected profits versus β and full crowding out

Panel (a) plots the algorithmic trading profit Π as a function of trading intensity β, while panel (b) plots the total human trading profit $K\Pi_H$ (dashed) and the average human trading profit Π_H (solid) as a function of β. The vertical dotted line in both panels corresponds to the optimal trading intensity with zero human traders, β_0^*, and the horizontal dotted line in panel (b) corresponds to the cost of information acquisition for a human trader, c.

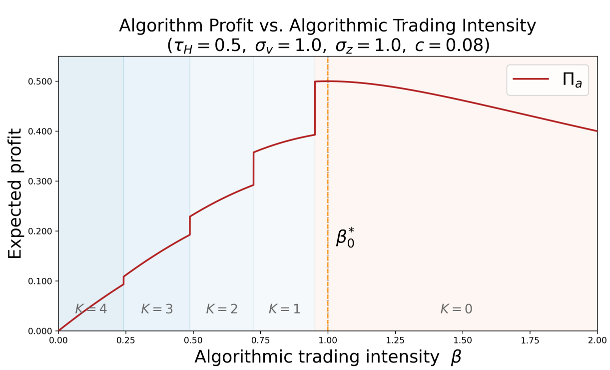

(a) $\Pi(\beta)$ versus β

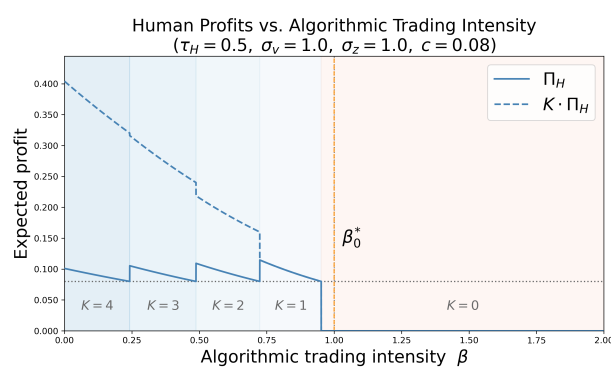

(b) $\Pi_H(\beta)$ versus β

5.3 Price informativeness

Since the equilibrium price p_t is proportional to the aggregate order flow y_t, a natural measure of price informativeness is the percentage reduction in uncertainty about v_t, i.e.,

$$PI = \frac{\mathbb{V}(v_t) - \mathbb{V}(v_t|p_t)}{\mathbb{V}(v_t)} = \frac{\mathbb{C}(v_t, y_t)^2}{\mathbb{V}(y_t)\mathbb{V}(v_t)} = \Gamma\lambda = \frac{\Gamma^2\sigma_v^2}{\Gamma^2\sigma_v^2 + Kb^2\sigma_\varepsilon^2 + \sigma_z^2}, \tag{23}$$

where $\Gamma \equiv Kb + \beta$.[26] The following result characterizes how price informativeness changes with algorithmic trading intensity when the number of human participants is fixed.

Lemma 3. *For a fixed K, price informativeness is increasing in the trading intensity of the algorithmic trader, i.e., PI is increasing in β.*

Intuitively, an increase in algorithmic trading intensity makes the order flow more informative through two channels. First, more aggressive trading by the algorithm directly increases the informativeness of the order flow by increasing its sensitivity to the algorithm's private information, v_t. Second, an increase in β leads to a decrease in the trading intensity of human traders, which leads to a reduction in noise that their private signals introduce to the order flow (i.e., a decrease in the weight on σ_ε^2). This implies that when the number of human traders is fixed, increased trading by the algorithmic trader always increases price informativeness.

However, as the analysis in the previous section highlights, more aggressive algorithmic trading can crowd out participation by human traders. This can have the opposite effect on price informativeness, as the following result establishes.

Proposition 6. *Suppose the conditions in Proposition 5 hold and*

$$K\sigma_v^2 > \sigma_v^2 + 2\sigma_\varepsilon^2. \tag{24}$$

Then price informativeness is lower in the presence of the algorithmic trader than in the economy with only K human traders.

Intuitively, Proposition 6 ensures that the presence of the algorithmic trader crowds out all human traders. In an economy with K active human traders and no algorithmic traders, we can show that price informativeness is given by

$$PI_K = \frac{K\tau_H}{(K-1)\tau_H + 2}, \tag{25}$$

[26] Note that *price efficiency* is always maximized in our model, since the price is a conditional expectation of fundamentals given the order flow, and so the price is semi-strong-form efficient by construction.

Figure 5: Price informativeness versus algorithmic trading intensity

The solid line corresponds to the equilibrium price informativeness with endogenous participation. The dot-dashed, horizontal line corresponds to price informativeness when only human traders are present, while the dashed horizontal line corresponds to the equilibrium price informativeness when only the algorithmic trader is present and is trading at β_0^*.

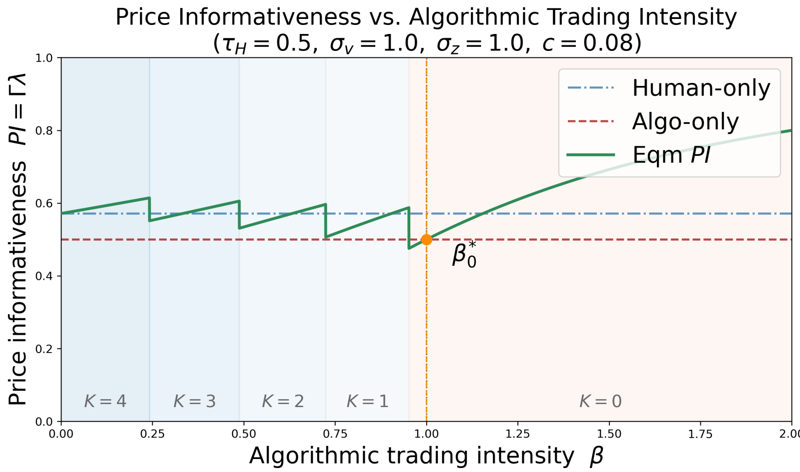

while in an economy with only the algorithmic trader, price informativeness is given by

$$PI_0 = \lambda_0^* \beta_0^* = \frac{1}{2}.$$

Condition (24) ensures that $PI_K > PI_0$. This implies that, together with the condition in Proposition 5, condition (24) is sufficient to ensure that price informativeness decreases in the presence of the algorithmic trader.

Figure 5 provides an illustration of these results. The endogenous exit decision by human traders gives rise to a saw-tooth pattern in price informativeness as β increases. Specifically, for each region where the number K of human traders is fixed, price informativeness increases in β. But, at threshold levels of β, human traders exit and this leads to a discrete drop in price informativeness.

5.4 Model-free versus model-based crowding out

In general, the algorithmic trader does not fully crowd out human traders. Instead, the two architectures induce different trading intensities and, by Proposition 4, different levels of human participation. We now compare price informativeness across architectures with partial crowding out. Let $PI_a(K)$ denote price informativeness in the algorithmic trading equilibrium with architecture $a \in \{A, M\}$ and K active human traders, so that equilibrium price informativeness under architecture a is $PI_a(K_a^*)$.

Proposition 7. *Price informativeness in the algorithmic trading equilibrium with $K \geq 1$ active human traders is given by*

$$PI_A(K) = \frac{2 + (K-1)\tau_H}{4 + (K-2)\tau_H} \quad \textit{and} \quad PI_M(K) = \frac{2 + (2K-1)\tau_H}{2\left(2 + (K-1)\tau_H\right)}.$$

1. *For fixed $K \geq 1$, $PI_M(K) > PI_A(K)$, and both are strictly increasing in K, with $PI_A(0) = PI_M(0) = PI_0 = \frac{1}{2}$.*

2. *With endogenous participation, $PI_M(K_M^*) < PI_A(K_A^*)$ if and only if $K_A^* > 2K_M^*$. A sufficient condition is*

$$\frac{\tau_H \sigma_v \sigma_z}{4\sqrt{4-\tau_H}} < c \leq \frac{\tau_H \sigma_v \sigma_z}{(4-\tau_H)\sqrt{\tau_H + (2-\tau_H)^2}}. \tag{26}$$

3. *Price informativeness is lower than in the economy with only K human traders whenever*

   (a) $K_A^* < K + 1 - \frac{2}{\tau_H}$ *under the model-free architecture, and*

   (b) $2K_M^* < K + 1 - \frac{2}{\tau_H}$ *under the model-based architecture.*

Intuitively, the comparison across architectures depends on the *intensity effect* and the *participation effect*. Part 1 is the intensity effect: for the same set of participants, the model-based algorithm's more aggressive trading makes the order flow more informative, so if participation were held fixed, mandating the model-based architecture would improve price informativeness.

Part 2 is the participation effect: because human traders earn strictly lower profits against the model-based algorithm, it induces weakly more exit (Proposition 4), and each exit lowers price informativeness discretely. When the additional crowding out is strong enough, equilibrium price informativeness is strictly lower under the model-based architecture. A special case is when the model-based algorithm empties the market of human traders that the model-free algorithm would retain — this is the condition in Equation (26).

Part 3 is another implication of the participation effect. Since the algorithmic trader is more aggressive under the model-based architecture, ensuring lower price informativeness in the presence of algorithmic trading requires more crowding out of human traders under the model-based architecture.

Figure 6 illustrates the comparison across architectures. Panels (a) and (b) mark the steady-state intensities β_A^* and β_M^* on the expected profit curves for the algorithm and human traders, respectively, and panel (c) marks them on the price informativeness curve

Figure 6: Model-free versus model-based algorithmic trading

Panels (a) and (b) reproduce the profit curves of Figure 4 and panel (c) reproduces the price informativeness curve of Figure 5, marking the model-free and model-based steady-state intensities β_A^* and β_M^*. Equilibrium price informativeness is $PI = 0.57$ under the model-free architecture and $PI = 0.50$ under the model-based architecture.

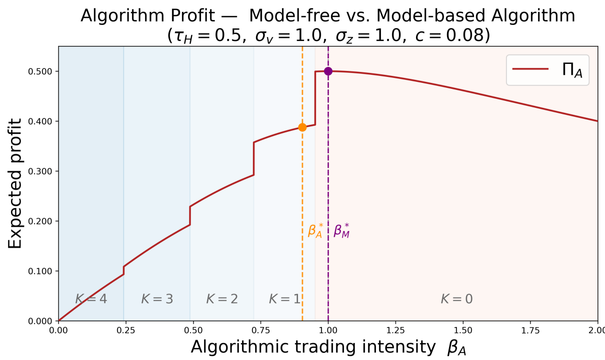

(a) Π versus β

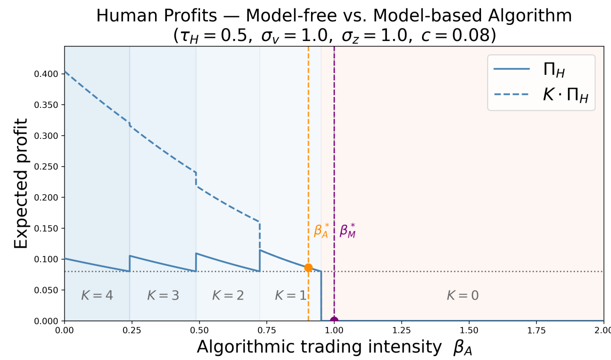

(b) Π_H versus β

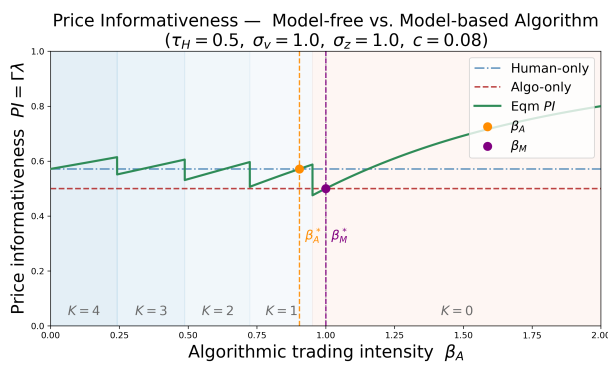

(c) Price Informativeness vs. β

from Figure 5. In the example, the model-based algorithm's additional aggressiveness pushes one more human trader past the exit threshold: consistent with Propositions 4 and 7, the model-based algorithm earns a strictly higher profit, while equilibrium price informativeness is strictly lower ($PI = 0.50$ versus $PI = 0.57$). The privately optimal architecture is the one under which prices reveal less.

6 Regulation

Algorithmic trading has attracted a wide range of regulatory proposals. In this section, we use our framework to evaluate three prominent classes of policies: (i) mandates that trading algorithms be interpretable and auditable, (ii) policies that alter the relative quality of information available to human and algorithmic traders, and (iii) caps on algorithmic trading. Our analysis suggests that the impact of policy depends on the architecture of algorithmic trading and on the resulting endogenous participation by human traders.

6.1 Interpretability mandates

Regulators increasingly require that trading algorithms be documented, tested, and interpretable. In the European Union, MiFID II requires firms engaged in algorithmic trading to maintain descriptions of their algorithms and subjects these algorithms to testing and audit requirements.[27] In the United States, the SEC has proposed rules governing the use of predictive data analytics by broker-dealers and investment advisers.[28] More broadly, the EU Artificial Intelligence Act imposes transparency and documentation obligations on providers of AI systems (see Regulation (EU) 2024/1689). A common thread in these proposals is a preference for "interpretable" algorithms, whose behavior can be understood and audited, over "black-box" systems.

In our setting, the model-based architecture corresponds to an interpretable algorithm (see Rudin (2019)): its behavior is derived from an explicit conjecture about price formation (Equation (14)), and its parameter (i.e., the price impact estimate $\widehat{\lambda}_t$) has a transparent interpretation that can be communicated to, and audited by, a regulator. By contrast, the model-free architecture is a black-box: it learns from realized profits alone and maintains no interpretable representation of the market in which it trades. In the language of our model, an interpretability mandate requires the algorithmic trader to be model-based.

Our analysis suggests that such mandates may harm market quality. Mandating the model-based architecture leads to more aggressive trading and higher expected profits for the algorithm in the steady state. As we show in Section 5, this leads to lower profits and possible crowding out of human traders and can decrease price informativeness. As such, these mandates may be counterproductive.

6.2 Subsidizing human data versus handicapping the algorithm

A second class of policies aims to level the informational playing field between algorithmic and human traders. These policies come in two varieties. The first improves the information available to all traders: for instance, the SEC's Market Data Infrastructure rules

[27] Article 17 of the Markets in Financial Instruments Directive (Directive 2014/65/EU, "MiFID II"); the accompanying technical standards (Commission Delegated Regulation (EU) 2017/589, "RTS 6") specify documentation, testing, and audit requirements.

[28] Conflicts of Interest Associated with the Use of Predictive Data Analytics by Broker-Dealers and Investment Advisers, Exchange Act Release No. 34-97990 (July 26, 2023). The CFTC's Regulation AT, proposed in 2015, would have given the regulator access to firms' source code, but was later withdrawn in favor of principles-based rules.

expand the content of cheap, consolidated market data with the explicit goal of narrowing the gap between public data feeds and the proprietary products purchased predominantly by algorithmic traders.[29] The second variety instead restricts the information available to algorithmic traders: privacy regulation limits the collection and resale of the alternative data (e.g., geolocation and transaction records) that feed algorithmic traders' signals,[30] and proposals to curb co-location and proprietary exchange data feeds recur in policy debates (see Budish, Cramton, and Shim (2015) for a related critique of the latency arms race).

In our model, these two approaches map into different primitives. Subsidizing public data lowers the cost of information acquisition c by human traders (equivalently, it makes Assumption 1 easier to satisfy), while handicapping the algorithm raises the noise in its signal, σ_η^2, in the extension in Appendix B.2.

In our setting, these approaches can have starkly different implications. Lowering c operates only through the participation margin, and only in one direction: given Equation (19), the equilibrium number of human traders K^* weakly increases as c falls, and price informativeness weakly increases with it (Proposition 10). Raising σ_η^2, by contrast, operates through both margins, and they push in opposite directions: for a fixed number of human traders, a less informed algorithm trades less aggressively and prices are less informative, but human traders' profits rise, so that weakly more of them participate (Proposition 10). The net effect on price informativeness is ambiguous — handicapping the algorithm destroys information directly, and only indirectly induces the entry that might restore it.

For a regulator whose objective is price informativeness, this asymmetry favors subsidies over handicaps: improving human traders' access to information weakly raises informativeness, while restricting the algorithm's information may lower it. The logic mirrors Section 5.3: the harm from algorithmic trading operates through the participation of human traders, and the robust remedy targets participation directly rather than the algorithm's informational advantage.

6.3 Caps on algorithmic trading intensity

Regulators also limit the scale of electronic trading through pre-trade risk controls. In the United States, Rule 15c3-5 requires broker-dealers with market access to reject orders that exceed appropriate size parameters. In the European Union, MiFID II requires algorithmic traders to be subject to trading thresholds and limits, while RTS 6 requires controls that

[29] Market Data Infrastructure, Exchange Act Release No. 34-90610 (December 9, 2020).

[30] See, e.g., the EU General Data Protection Regulation (Regulation (EU) 2016/679), which restricts the processing and resale of personal data commonly used in alternative-data products.

block orders exceeding maximum values or volumes.[31] These rules are primarily designed to prevent erroneous or disorderly trading rather than to regulate information aggregation. Nevertheless, they provide a natural analogue of a cap on the aggressiveness of algorithmic trading.

We therefore suppose that the regulator constrains the algorithmic trader's policy to lie in the interval $[0, \bar{\beta}]$. Specifically, for a fixed number K of human traders, let $\beta_a^u(K)$ denote the steady state intensity under architecture $a \in \{A, M\}$, and suppose the constrained intensity is given by:

$$\beta_a^c(K; \bar{\beta}) = \min\left\{\beta_a^u(K), \bar{\beta}\right\}. \tag{27}$$

Human trading intensity, price impact, and participation continue to be characterized by Equations (3), (4), and (19), respectively, evaluated at β_a^c.[32]

The cap has two opposing effects on price informativeness. Holding participation fixed, it removes informative algorithmic trading. At the same time, it raises human traders' profits and can induce entry. These are, once again, the *intensity effect* and the *participation effect.*

Proposition 8. *Fix $K \geq 1$ and suppose that the cap binds, i.e., $0 < \bar{\beta} < \beta_a^u(K)$. Then b^c and Π_H^c are strictly decreasing in $\bar{\beta}$, whereas Π_a^c and price informativeness PI^c are strictly increasing in $\bar{\beta}$.*

With endogenous participation, over any range in which the cap binds, the equilibrium number of human traders is weakly decreasing in $\bar{\beta}$. Hence, as the cap is tightened, price informativeness falls continuously between entry thresholds but jumps upward whenever an additional human trader enters. Consequently, the overall effect of a tighter cap on price informativeness is ambiguous.

Intuitively, a binding cap prevents the algorithm from exploiting some of its superior information. This directly reduces the informativeness of order flow, which is the intensity effect. However, less aggressive algorithmic trading raises human traders' profits and induces entry. Each additional trader contributes information to prices, which is the participation effect. Thus, a cap that is only modestly tightened may reduce price informativeness, while a sufficiently tight cap can improve it by reversing crowding out.

In Appendix B.3, we consider the impact of transaction taxes. We compare the impact of targeted versus untargeted taxes on price informativeness under the model-free architec-

[31] See ESMA Supervisory briefing on algorithmic trading in the EU (2026) and SEC Rule 15c3-5.

[32] Actual regulations generally impose limits on the size or value of individual orders, whereas Equation (27) constrains the sensitivity of the algorithm's demand to its information. We use the latter formulation because a hard cap on realized orders, $|x_{A,t}| \leq \bar{x}$, makes demand nonlinear in v_t and breaks the linear equilibrium characterization. The intensity cap is thus a tractable reduced form for such pre-trade limits.

ture.[33] With a transaction tax that targets algorithmic trading, the same non-monotonicity in price informativeness arises: holding participation fixed, higher taxes harm price informativeness (by reducing algorithmic trading intensity), while with endogenous entry, higher taxes can improve price informativeness (by increasing human participation). An untargeted transaction tax, however, uniformly reduces informed trading and thus leads to lower price informativeness.

6.4 Welfare

The analysis so far characterizes how the algorithm's architecture affects price informativeness. A natural question is how it affects welfare, and who bears the cost. A standard measure of welfare is the sum of expected profits across the algorithmic trader, the human traders, and the noise traders, net of the participation costs the human traders pay to acquire information. However, because the market maker is competitive and trading is zero-sum, this is completely driven by aggregate information costs. Specifically, in the algorithmic trading equilibrium with architecture $a \in \{A, M\}$, welfare is given by:

$$W_a = \Pi + K_a^* \Pi_H + \Pi_{z,a} - \, K_a^* \, c = - \, K_a^* \, c,$$

where $\Pi_{z,a}$ is the expected profit of noise traders, i.e.,

$$\Pi_{z,a} = \mathbb{E}[z_t(v_t - p_t)] = -\mathbb{E}[z_t p_t] = -\lambda(\beta)\sigma_z^2.$$

Moreover, this implies that

$$W_M \geq W_A,$$

with strict inequality whenever the model-based architecture strictly crowds out human traders ($K_M^* < K_A^*$).

The intuition is standard. In a pure-exchange market private information has a social cost but no social value, since trading is zero-sum (e.g., Hirshleifer (1971)). This is why our regulatory analysis above focuses on price informativeness, rather than trading surplus, as the objective. We abstract from explicitly modeling how the information aggregated by the price affects real or allocative decisions (e.g., see Bond, Edmans, and Goldstein (2012)). However, to the extent that price informativeness improves real efficiency, it is a natural objective for regulation, and the crowding out of human traders we analyze is socially costly. In the absence of a role for price informativeness, the model-based architecture is not only

[33] The impact of taxes under the model-based architecture depends on whether the algorithm accounts for the tax when specifying its expected profits.

privately optimal for the algorithmic trader, but also socially optimal since it minimizes information acquisition costs.

7 Conclusion

We develop a model in which an informed algorithmic trader learns to trade strategically against human traders, who choose whether to acquire costly information and participate in the market. We compare two architectures. Under the model-free architecture, the algorithm learns its trading intensity from realized profits alone. Under the model-based architecture, the algorithm uses a misspecified model of price formation to estimate the price impact of its trades. We show that the model-based algorithm leads to more aggressive trading, earns higher expected profits, but can crowd out more human traders and lead to lower price informativeness.

Our model provides a natural benchmark for evaluating regulatory policies. A mandate for more "interpretable" algorithms, which selects the more aggressive model-based algorithm in our setting, can lead to worse market outcomes. Restrictions on algorithmic information or trading intensity can improve price informativeness by inducing higher human participation, while policies that directly lower the cost of human information acquisition avoid the accompanying loss of algorithmic information.

Our analysis suggests several directions for future work. Extending the analysis to allow for multiple, competing algorithms would allow for the possibility of a trading intensity "arms race," whereby each algorithm trades even more aggressively to deter participation by others. Such a setting would also allow the number and architecture of algorithms to be determined endogenously. A central question is whether competition among algorithms disciplines their behavior or instead amplifies the crowding out of human information. It would also be interesting to study how algorithmic trading interacts with real decisions in the economy (e.g., firm investment, capital allocation, or managerial learning from prices), as this would allow for richer welfare analysis. Finally, our analysis suggests that an algorithm's architecture may be inferred from how its trading responds to changes in market conditions. Developing empirical tests that distinguish more aggressive trading due to model misspecification from better informed trading would be useful in evaluating the impact of policy proposals, since these can vary significantly depending on the algorithm's architecture.

References

Simona Abis. Man versus machine: Quantitative and discretionary equity management. Technical report, Working paper, Columbia Business School, 2020.

John Asker, Chaim Fershtman, and Ariel Pakes. Artificial intelligence, algorithm design, and pricing. In *AEA Papers and Proceedings*, volume 112, pages 452–456, 2022.

Snehal Banerjee. Learning from prices and the dispersion in beliefs. *The Review of Financial Studies*, 24(9):3025–3068, 2011.

Snehal Banerjee and Bradyn Breon-Drish. Strategic trading and unobservable information acquisition. *Journal of Financial Economics*, 138(2):458–482, 2020.

Snehal Banerjee and Martin Szydlowski. Trading against algorithms: Price dynamics and risk-sharing in a market with q-learners. *Available at SSRN 5380152*, 2025.

Snehal Banerjee, Jesse Davis, and Naveen Gondhi. Choosing to disagree: Endogenous dismissiveness and overconfidence in financial markets. *The Journal of Finance*, 79(2):1635–1695, 2024a.

Snehal Banerjee, George Malikov, and Filipp Prokopev. Competition and collusion among strategic traders who face uncertainty. *Available at SSRN 4980072*, 2024b.

Nicholas C Barberis and Lawrence J Jin. Model-free and model-based learning as joint drivers of investor behavior. Technical report, National Bureau of Economic Research, 2023.

Dan Bernhardt and Alex Boulatov. Strategic commitment by an informed speculator. Working Paper, University of Warwick, 2023.

Ekkehart Boehmer, Kingsley Fong, and Juan Julie Wu. Algorithmic trading and market quality: International evidence. *Journal of Financial and Quantitative Analysis*, 56(8): 2659–2688, 2021.

Philip Bond, Alex Edmans, and Itay Goldstein. The real effects of financial markets. *Annu. Rev. Financ. Econ.*, 4(1):339–360, 2012.

Maxime Bonelli and Thierry Foucault. Does big data devalue traditional expertise? evidence from active funds. 2023.

Timothy F Bresnahan and Peter C Reiss. Entry in monopoly market. *The Review of Economic Studies*, 57(4):531–553, 1990.

Jonathan Brogaard, Terrence Hendershott, and Ryan Riordan. High-frequency trading and price discovery. *The Review of Financial Studies*, 27(8):2267–2306, 2014.

Eric Budish, Peter Cramton, and John Shim. The high-frequency trading arms race: Frequent batch auctions as a market design response. *Quarterly Journal of Economics*, 130 (4):1547–1621, 2015.

Sean Cao, Wei Jiang, Baozhong Yang, and Alan L. Zhang. How to talk when a machine is listening: Corporate disclosure in the age of AI. *The Review of Financial Studies*, 36(9): 3603–3642, 2023. Publisher: Oxford University Press.

Giuseppe Cappelletti, Giovanni Guazzarotti, and Pietro Tommasino. The stock market effects of a securities transaction tax: Quasi-experimental evidence from Italy. *Journal of Financial Stability*, 31:81–92, 2017.

Alain P Chaboud, Benjamin Chiquoine, Erik Hjalmarsson, and Clara Vega. Rise of the machines: Algorithmic trading in the foreign exchange market. *The Journal of Finance*, 69(5):2045–2084, 2014.

Braiden Coleman, Kenneth Merkley, and Joseph Pacelli. Human versus machine: A comparison of robo-analyst and traditional research analyst investment recommendations. *The Accounting Review*, 97(5):221–244, 2022.

Jean-Edouard Colliard and Peter Hoffmann. Financial transaction taxes, market composition, and liquidity. *Journal of Finance*, 72(6), 2017.

Jean-Edouard Colliard, Thierry Foucault, and Stefano Lovo. Algorithmic pricing and liquidity in securities markets. *HEC Paris Research Paper No. FIN-2022-1459*, 2022.

Eduardo Dávila. Optimal financial transaction taxes. *Journal of Finance*, 78(1):5–61, 2023.

Winston Wei Dou, Itay Goldstein, and Yan Ji. Ai-powered trading, algorithmic collusion, and price efficiency. *Available at SSRN 4452704*, 2023.

Winston Wei Dou, Itay Goldstein, and Yan Ji. Financial market fragility in the era of ai planning. *Available at SSRN 5763222*, 2025a.

Winston Wei Dou, Itay Goldstein, Zigang Li, and Liyan Yang. Human edge, machine limits: Ai-human competition in financial markets. *Machine Limits: AI-Human Competition in Financial Markets (November 18, 2025)*, 2025b.

Ignacio Esponda and Demian Pouzo. Berk–nash equilibrium: A framework for modeling agents with misspecified models. *Econometrica*, 84(3):1093–1130, 2016.

George W Evans and Seppo Honkapohja. *Learning and expectations in macroeconomics.* Princeton University Press, 2001.

Erik Eyster, Matthew Rabin, and Dimitri Vayanos. Financial markets where traders neglect the informational content of prices. *The Journal of Finance*, 74(1):371–399, 2019.

Chaim Fershtman and Kenneth L Judd. Equilibrium incentives in oligopoly. *The American Economic Review*, pages 927–940, 1987.

F Douglas Foster and S Viswanathan. Strategic trading when agents forecast the forecasts of others. *The Journal of Finance*, 51(4):1437–1478, 1996.

F Douglas Foster and Subramanian Viswanathan. The effect of public information and competition on trading volume and price volatility. *The Review of Financial Studies*, 6 (1):23–56, 1993.

Drew Fudenberg, Giacomo Lanzani, and Philipp Strack. Limit points of endogenous misspecified learning. *Econometrica*, 89(3):1065–1098, 2021.

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

Jillian Grennan and Roni Michaely. Fintechs and the market for financial analysis. *Journal of financial and quantitative analysis*, 56(6):1877–1907, 2021.

Sanford J Grossman and Joseph E Stiglitz. On the impossibility of informationally efficient markets. *American Economic Review*, 70(3):393–408, 1980.

Antonio Guarino, Philippe Jehiel, and James Symons-Hicks. Q-learning and algorithmic market making: Loss-free, collusive, or competitive prices? *CEPR Discussion Paper*, 2025.

Ivan Gufler, Francesco Sangiorgi, and Emanuele Tarantino. (deep) learning to trade: An experimental analysis of ai trading and market outcomes. *Available at SSRN 5375160*, 2025.

Paul Heidhues, Botond Kőszegi, and Philipp Strack. Unrealistic expectations and misguided learning. *Econometrica*, 86(4):1159–1214, 2018.

Terrence Hendershott, Charles M Jones, and Albert J Menkveld. Does algorithmic trading improve liquidity? *The Journal of finance*, 66(1):1–33, 2011.

Jack Hirshleifer. The private and social value of information and the reward to inventive activity. *The American Economic Review*, 61(4):561–574, 1971.

Craig W Holden and Avanidhar Subrahmanyam. Long-lived private information and imperfect competition. *The Journal of Finance*, 47(1):247–270, 1992.

Zsolt Katona, Marcus O Painter, Panos N Patatoukas, and Jean Zeng. On the capital market consequences of big data: Evidence from outer space. *Journal of Financial and Quantitative Analysis*, 60(2):551–579, 2025.

Albert S Kyle. Continuous auctions and insider trading. *Econometrica*, pages 1315–1335, 1985.

Albert S Kyle, Anna A Obizhaeva, and Yajun Wang. Smooth trading with overconfidence and market power. *The Review of Economic Studies*, 85(1):611–662, 2018.

Timothy P. Lillicrap, Jonathan J. Hunt, Alexander Pritzel, Nicolas Manfred Otto Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. *arXiv: Learning*, 2015.

Bryan R Routledge. Adaptive learning in financial markets. *The Review of Financial Studies*, 12(5):1165–1202, 1999.

Cynthia Rudin. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5):206–215, 2019.

David Silver, Guy Lever, Nicolas Heess, Thomas Degris, Daan Wierstra, and Martin Riedmiller. Deterministic policy gradient algorithms. In *International conference on machine learning*, pages 387–395. Pmlr, 2014.

Avanidhar Subrahmanyam. Transaction taxes and financial market equilibrium. *Journal of Business*, 71(1):81–118, 1998.

Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, 2018.

Carla Tighe and Ron Michener. The political economy of insider-trading laws. *The American Economic Review*, 84(2):164–168, 1994.

James Tobin. A proposal for international monetary reform. *Eastern Economic Journal*, 4 (3/4):153–159, 1978.

Steven R. Umlauf. Transaction taxes and the behavior of the swedish stock market. *Journal of Financial Economics*, 33(2):227–240, 1993.

Robert E Verrecchia. Information acquisition in a noisy rational expectations economy. *Econometrica*, pages 1415–1430, 1982.

John Vickers. Delegation and the theory of the firm. *The Economic Journal*, 95(Supplement): 138–147, 1985.

Brian M Weller. Does algorithmic trading reduce information acquisition? *The Review of Financial Studies*, 31(6):2184–2226, 2018.

Marek Weretka. Endogenous market power. *Journal of Economic Theory*, 146(6):2281–2306, 2011.

Hongjun Yan, Liyan Yang, Xueyong Zhang, and Deqing Zhou. Kyle meets friedman: Informed trading when anticipating future information. 2024.

Christina Zhu. Big data as a governance mechanism. *The Review of Financial Studies*, 32 (5):2021–2061, 2019.

A Proofs

A.1 Proof of Proposition 1

Rearranging (2) yields

$$\lambda = \frac{\tau_H}{b\big(2+(K-1)\tau_H\big)+\beta\tau_H}, \tag{28}$$

and setting this equal to (1) and re-arranging yields

$$\Gamma\sigma_v^2\big(b(2+(K-1)\tau_H)+\beta\tau_H\big) = \tau_H\left(\Gamma^2\sigma_v^2+Kb^2\sigma_\varepsilon^2+\sigma_z^2\right).$$

Since $\Gamma = Kb+\beta$, note that this is a quadratic in b, which simplifies to:

$$K\sigma_v^2\, b^2+\beta(2-\tau_H)\sigma_v^2\, b-\tau_H\sigma_z^2 = 0. \tag{29}$$

The second-order condition requires $\lambda > 0$, and this in turn requires that we pick the positive root of the quadratic:

$$b^* = \frac{-(2-\tau_H)\beta\sigma_v^2+\sqrt{(2-\tau_H)^2\beta^2\sigma_v^4+4K\tau_H\sigma_z^2\sigma_v^2}}{2K\sigma_v^2}.$$

The equilibrium λ follows by substituting b^* into (28).

Turning to expected profits, note that $\mathbb{E}[s_i\sum_j s_j] = K\sigma_v^2+\sigma_\varepsilon^2$ and $\mathbb{E}[s_i v] = \sigma_v^2$, which implies:

$$\begin{aligned}\Pi_H &= b\big(\sigma_v^2-\lambda\big(b(K\sigma_v^2+\sigma_\varepsilon^2)+\beta\sigma_v^2\big)\big)\\ &= b\big(\sigma_v^2(1-\lambda\Gamma)-\lambda b\sigma_\varepsilon^2\big).\end{aligned}$$

Note that $\sigma_\varepsilon^2 = \sigma_v^2(1-\tau_H)/\tau_H$, and that the FOC for human traders implies: $1-\lambda\Gamma = \lambda b(2-\tau_H)/\tau_H$. This implies:

$$\Pi_H = \lambda b^2\left(\frac{(2-\tau_H)\sigma_v^2}{\tau_H}-\sigma_\varepsilon^2\right) = \lambda b^2\cdot\frac{\sigma_v^2}{\tau_H}.$$

Substituting (28), gives

$$\Pi_H = \frac{b^2\sigma_v^2}{b\big(2+(K-1)\tau_H\big)+\beta\tau_H}. \tag{30}$$

Similarly, we have:

$$\Pi = \beta\sigma_v^2(1-\lambda\Gamma) = \frac{\beta b(2-\tau_H)\sigma_v^2}{b\big(2+(K-1)\tau_H\big)+\beta\tau_H}. \tag{31}$$

Next, we establish the comparative statics results. Let $D \equiv b(2+(K-1)\tau_H)+\beta\tau_H$ be the common denominator in (30)–(31).

Human trading intensity and profits are decreasing in β. Differentiating (29) with respect to β yields:

$$b' \equiv \frac{db^*}{d\beta} = -\frac{(2-\tau_H)\,b}{2Kb+(2-\tau_H)\beta} < 0.$$

This implies that differentiating $\Pi_H = b^2\sigma_v^2/D$ yields:

$$\frac{d\Pi_H}{d\beta} = \frac{\sigma_v^2 b}{D^2}\Big[\underbrace{b'(D+\beta\tau_H)}_{<0}\underbrace{-b\tau_H}_{<0}\Big] < 0.$$

Both terms are strictly negative, so

$$\frac{d\Pi_H}{d\beta} < 0 \quad \text{for all } \beta > 0,\ K \geq 1. \tag{32}$$

Human trading intensity and profits are decreasing in K. Differentiating (29) with respect to K:

$$b'_K \equiv \frac{db^*}{dK} = -\frac{b^2}{2Kb+(2-\tau_H)\beta} < 0.$$

This implies that differentiating $\Pi_H = b^2\sigma_v^2/D$ yields:

$$\frac{d\Pi_H}{dK} = \frac{\sigma_v^2 b}{D^2}\Big[\underbrace{b'_K(D+\beta\tau_H)}_{<0}\underbrace{-b^2\tau_H}_{<0}\Big] < 0.$$

Both terms are strictly negative (using $b'_K < 0$ and $D+\beta\tau_H > 0$), so

$$\frac{d\Pi_H}{dK} < 0 \quad \text{for all } K \geq 1,\ \beta \geq 0.$$

A.2 Proof of Proposition 2

Combining Equation (9) with Equation (7) implies that

$$\mathbb{E}[(\pi_{A,t} - q_t(1-\gamma) - w_t\varepsilon_{A,t})\varepsilon_{A,t}|\Psi_t = \Psi^*] = 0.$$

Since q_t is a number, this expression simplifies to

$$w_t = \frac{\mathbb{E}[\pi_{A,t}\varepsilon_{A,t}|\Psi_t = \Psi^*]}{\sigma_A^2}.$$

In the small noise limit, as $\sigma_A \to 0$, the algorithmic trader's profit is given by

$$\pi_{A,t} = (\beta_{A,t} + \varepsilon_{A,t})\, v_t \times \left(v_t - \lambda\left(\beta_{A,t}\right) \times \left(\left(\beta_{A,t} + \varepsilon_{A,t}\right) v_t + b(\beta_{A,t}) \sum_{i=1}^{K(\beta_{A,t})} \left(v_t + \varepsilon_{i,t}\right) + z_t \right) \right),$$

where $b(\beta_{A,t})$, $\lambda(\beta_{A,t})$, and $K(\beta_{A,t})$ are given in Equations (3), (4), and (19). Importantly, the temporary perturbation is not observed before trading and therefore does not induce contemporaneous changes in b, λ or K. This implies that

$$\frac{E\left[\pi_{A,t}\varepsilon_{A,t}|\Psi_t = \Psi^*\right]}{\sigma_A^2} \to \sigma_v^2 \left(1 - \lambda\left(\beta_A^*\right)\left(2\beta_A^* + K(\beta_A^*) b\left(\beta_A^*\right)\right)\right)$$

in the small noise limit. By Definition 1, $w^* = 0$, which implies that

$$\beta_A^* = \frac{1 - \lambda(\beta_A^*) K(\beta_A^*) b(\beta_A^*)}{2\lambda(\beta_A^*)}.$$

It remains to verify the equivalence with the rational-trader benchmark. Consider a rational trader who observes v_t and faces the equilibrium strategies $(b_A^*, \lambda_A^*, K_A^*)$. Given the linear strategies of the human traders and the market maker's pricing rule, its expected profit from submitting a trade x is

$$\mathbb{E}\left[x \left(v_t - \lambda_A^* \left(x + b_A^* \textstyle\sum_{i=1}^{K_A^*} s_{i,t} + z_t \right) \right) \middle| v_t \right] = x\left(v_t - \lambda_A^*\left(x + K_A^* b_A^* v_t\right)\right),$$

since $\mathbb{E}[s_{i,t}|v_t] = v_t$ and $\mathbb{E}[z_t] = 0$. The objective is strictly concave in x, and the unique maximizer is

$$x = \frac{1 - \lambda_A^* K_A^* b_A^*}{2\lambda_A^*}\, v_t,$$

which is a linear strategy with intensity given by Equation (20). Hence, the steady state trading intensity of the model-free algorithm is the unique best response of a rational trader who observes v_t. Since conditions (2)–(4) of Definition 2 coincide with the equilibrium conditions for the human traders, the market maker, and participation in the economy with a rational insider, the tuple $(\beta_A^*, b_A^*, \lambda_A^*, K_A^*)$ is an algorithmic trading equilibrium if and only if it is a Bayes-Nash equilibrium of the economy in which the algorithmic trader is replaced by a rational strategic trader who observes v_t.

A.3 Proof of Proposition 3

The steady state equations (18) imply that in the small noise limit, $p_{0,t} = 0$,

$$\mathbb{E}\left[v_t\left(v_t - p_{0,t} - 2\widehat{\lambda}_t(\beta_{M,t} + \varepsilon_{M,t})v_t\right)|\Psi_{M,t} = \Psi_M^*\right] = 0, \tag{33}$$

$$\mathbb{E}\left[\lambda_t\left((\beta_{M,t} + \varepsilon_{M,t})v_t + b_t\sum_{i=1}^{K_t} s_{i,t} + z_t\right) - p_{0,t} - \widehat{\lambda}_t(\beta_{M,t} + \varepsilon_{M,t})v_t\,|\Psi_{M,t} = \Psi_M^*\right] = 0$$

and

$$\mathbb{E}\left[\left(\lambda_t\left(b_t\sum_{i=1}^{K_t} s_{i,t} + z_t\right) - p_{0,t} + \left(\lambda_t - \widehat{\lambda}_t\right)(\beta_{M,t} + \varepsilon_{M,t})\,v_t\right)\varepsilon_{M,t}v_t\,|\Psi_{M,t} = \Psi_M^*\right] = 0, \tag{34}$$

where we used $p_t = \lambda_t y_t$. Equation (33) implies that

$$\beta_M^* = \frac{1}{2\widehat{\lambda}^*}.$$

Equation (34) implies that

$$\sigma_M^2\sigma_v^2\left(\lambda_t - \widehat{\lambda}_t\right) = 0$$

and in particular that $\widehat{\lambda}^* = \lambda_M^*$. Hence, the equilibrium is characterized by $\beta_M^* = 1/2\lambda_M^*$, $\lambda_M^* = \lambda(\beta_M^*)$, where $\lambda(.)$ is given in Equation (4), and $b_M^* = b(\beta_M^*)$, where $b(.)$ is given in Equation (3).

A.4 Proof of Proposition 4

For expositional clarity, let $\Pi_A(K)$ and $\Pi_M(K)$ denote the algorithmic trader's expected per-period profit, and $\Pi_H^A(K)$ and $\Pi_H^M(K)$ the human traders' expected per-period profit, in the algorithmic trading equilibrium with K active human traders under the model-free and model-based architectures, respectively.

A.4.1 Proof of Part (1)

Fix K. Using Equations (3), (4), and (20) implies that

$$\beta_A^* = \frac{\sigma_z\left(\sigma_v^2 + 2\sigma_\varepsilon^2\right)}{\sigma_v\sqrt{(K+1)\sigma_v^4 + (K+4)\sigma_v^2\sigma_\varepsilon^2 + 4\sigma_\varepsilon^4}},$$

after some algebra. We can similarly use Equations (3), (4), and (21) to obtain

$$\beta_M^* = \frac{\sigma_z\left((K+1)\sigma_v^2 + 2\sigma_\varepsilon^2\right)}{\sqrt{(2K+1)\sigma_v^6 + (3K+4)\sigma_v^4\sigma_\varepsilon^2 + 4\sigma_v^2\sigma_\varepsilon^4}}.$$

It then follows that $\beta_A^* < \beta_M^*$ in equilibrium.

To show that $\Pi_M(K) > \Pi_A(K)$, we can calculate

$$\Pi_A(K) = \frac{\sigma_v\sigma_z\left(\sigma_v^2 + 2\sigma_\varepsilon^2\right)^2}{\left((K+2)\sigma_v^2 + 4\sigma_\varepsilon^2\right)\sqrt{(K+1)\sigma_v^4 + (K+4)\sigma_v^2\sigma_\varepsilon^2 + 4\sigma_\varepsilon^4}} \tag{35}$$
$$\Pi_M(K) = \frac{\sigma_v\sigma_z\left(\sigma_v^2 + 2\sigma_\varepsilon^2\right)}{2\sqrt{(2K+1)\sigma_v^4 + (3K+4)\sigma_v^2\sigma_\varepsilon^2 + 4\sigma_\varepsilon^4}}.$$

After some algebra, it follows that $\Pi_M(K) > \Pi_A(K)$.

Under the model-free architecture, combining the human traders' first-order condition, Equation (28), with the steady-state condition (20) yields

$$u_A \equiv \beta_A^*/b^* = (2-\tau_H)/\tau_H, \tag{36}$$

and substituting into Equation (29) and using Equation (5) gives

$$\Pi_H^A(K) = \frac{\tau_H\sigma_v\sigma_z}{\left(4+(K-2)\tau_H\right)\sqrt{K\tau_H + (2-\tau_H)^2}}.$$

Under the model-based architecture, combining Equation (28) with the steady-state condition $\beta_M = \frac{1}{2\lambda}$ yields

$$u_M \equiv \beta_M^*/b^* = \frac{2+(K-1)\tau_H}{\tau_H}. \tag{37}$$

Substituting $\beta_M = u_M b$ into Equation (29) gives

$$b_M(K) = \frac{\tau_H\sigma_z}{\sigma_v\sqrt{K\tau_H + (2-\tau_H)\left(2+(K-1)\tau_H\right)}},$$

and, using $\Pi_H = \lambda b^2\sigma_v^2/\tau_H = b\,\sigma_v^2/\left(2+(K-1)\tau_H + u_M\tau_H\right)$ with $2+(K-1)\tau_H + u_M\tau_H = 2\left(2+(K-1)\tau_H\right)$,

$$\Pi_H^M(K) = \frac{\tau_H\sigma_v\sigma_z}{2\left(2+(K-1)\tau_H\right)\sqrt{K\tau_H + (2-\tau_H)\left(2+(K-1)\tau_H\right)}}.$$

The above expressions immediately imply that $\Pi_H^A(K)$ and $\Pi_H^M(K)$ are strictly decreasing in K. Moreover, for every $K \geq 1$,

$$2\left(2+(K-1)\tau_H\right)-\left(4+(K-2)\tau_H\right)=K\tau_H>0$$

and

$$\left(2-\tau_H\right)\left(2+(K-1)\tau_H\right)-\left(2-\tau_H\right)^2=(2-\tau_H)K\tau_H>0,$$

so the denominator of $\Pi_H^M(K)$ strictly exceeds that of $\Pi_H^A(K)$, and therefore

$$\Pi_H^M(K)<\Pi_H^A(K) \quad \text{for all } K \geq 1.$$

A.4.2 Proof of Part (2)

Since $\Pi_H^M(K)$ and $\Pi_H^A(K)$ are strictly decreasing in K and for each K, $\Pi_H^M(K) < \Pi_H^A(K)$, we have $\Pi_H^A(K_M^*) > \Pi_H^M(K_M^*) \geq c$, and hence $K_M^* \leq K_A^*$.

If $K_A^* = 0$, then $K_M^* = 0$, and Lemma 1 implies that the two equilibria coincide, so $\Pi(\beta_M^*; K) = \Pi(\beta_A^*; K)$. If $K_A^* \geq 1$, then

$$\Pi(\beta_M^*; K_M^*) \geq \Pi(\beta_M^*; K_A^*) > \Pi(\beta_A^*; K_A^*),$$

where the first inequality follows from $K_M^* \leq K_A^*$ and the monotonicity of $\Pi_M(\cdot)$, and the second inequality is the fixed-K comparison from part (1).

A.5 Proof of Lemma 2

Setting $K = 1$, the equilibrium quadratic (29) becomes

$$\sigma_v^2 b^2+\beta(2-\tau_H)\sigma_v^2 b-\tau_H\sigma_z^2=0, \tag{38}$$

and

$$\Pi_H(\beta; K=1)=\frac{b(\beta)^2\,\sigma_v^2}{2b(\beta)+\beta\tau_H}.$$

Let $u \equiv \beta/b$. Then, solving (38) yields

$$\begin{aligned} b &= \frac{\sqrt{\tau_H}\,\sigma_z}{\left(\sigma_v\sqrt{1+u(2-\tau_H)}\right)} \\ \Pi_H &= \frac{b\sigma_v^2}{2+u\tau_H}=\frac{\sqrt{\tau_H}\,\sigma_v\sigma_z}{(2+u\tau_H)\sqrt{1+u(2-\tau_H)}}, \end{aligned}$$

The threshold is characterized by $\Pi_H = c$, which implies:

$$\frac{\sqrt{\tau_H}\,\sigma_v\sigma_z}{(2+u\tau_H)\sqrt{1+u(2-\tau_H)}} = c$$
$$\Leftrightarrow (2+u\tau_H)\sqrt{1+u(2-\tau_H)} = \sqrt{\tau_H}\,\sigma_v\sigma_z/c > 2, \tag{39}$$

where the last inequality follows from Assumption 1. Note that the LHS is increasing in u and is equal to 2 when $u = 0$. This implies there exists a unique u which satisfies (39), which corresponds to a unique threshold intensity $\underline{\beta}$. Moreover, since $u = \beta/b$ and b is decreasing in β (from Proposition 1), the comparative statics results follow from implicitly differentiating (39).

A.6 Proof of Proposition 5

When $K = 1$, Lemma 2 implies that we can characterize the expected profit of a human as

$$\Pi_H(u) = \frac{\sqrt{\tau_H}\,\sigma_v\sigma_z}{(2+u\tau_H)\sqrt{1+u(2-\tau_H)}},$$

where $u = \beta/b(\beta)$. From (36) and (37), we know that $u_A = (2-\tau_H)/\tau_H$ and $u_M = (2+(K-1)\tau_H)/\tau_H$. This implies

$$\Pi_H(u_A) = \frac{\tau_H\sigma_v\sigma_z}{(4-\tau_H)\sqrt{\tau_H+(2-\tau_H)^2}}$$
$$\Pi_H(u_M) = \frac{\tau_H\sigma_v\sigma_z}{4\sqrt{4-\tau_H}}$$

When these quantities are smaller than c, entry by a human trader is not profitable. Finally $c_M < c_A$ follows from algebra.

A.7 Proof of Lemma 3

Using $u \equiv \beta/b$, $\Gamma = b(K+u)$ and $\lambda = \tau_H/(b(2+(K-1)\tau_H+u\tau_H))$, we can express price informativeness as:

$$PI(u) = \lambda\Gamma = \frac{\tau_H(K+u)}{2+(K-1)\tau_H+u\tau_H}. \tag{40}$$

Differentiating:

$$\frac{dPI}{du} = \frac{\tau_H(2-\tau_H)}{(2+(K-1)\tau_H+u\tau_H)^2} > 0,$$

using $\big(2+(K-1)\tau_H\big) - K\tau_H = 2-\tau_H > 0$. Since $b(\beta)$ is decreasing in β (Proposition 1), we have $u = \beta/b(\beta)$ is increasing in β and this implies for a fixed K, price informativeness is increasing in β.

A.8 Proof of Proposition 6

In the equilibrium with K participating humans and no algorithmic trader, we can show that

$$b^H = \frac{\sqrt{\tau_H}\,\sigma_z}{\sqrt{K}\,\sigma_v}, \qquad \lambda^H = \frac{\sqrt{K\tau_H}\,\sigma_v}{\sigma_z\big((K-1)\tau_H+2\big)},$$

and price informativeness

$$PI_K = Kb^H\lambda^H = \frac{K\tau_H}{(K-1)\tau_H+2}.$$

In contrast, when $K=0$, the algorithmic trader is a Kyle monopolist. From Lemma 1, this implies

$$PI_0 = \lambda\beta = \frac{1}{2}.$$

Finally, we can show condition (24) by setting $PI_K > PI_0$ and plugging in $\tau_H = \sigma_v^2/(\sigma_v^2+\sigma_\varepsilon^2)$.

A.9 Proof of Proposition 7

Again, for expositional clarity, let $\Pi_A(K)$ and $\Pi_M(K)$ denote the algorithmic trader's expected per-period profit, and $\Pi_H^A(K)$ and $\Pi_H^M(K)$ the human traders' expected per-period profit, in the algorithmic trading equilibrium with K active human traders under the model-free and model-based architectures, respectively.

Under the model-free architecture, we have $\beta_A^*/b^* = (2-\tau_H)/\tau_H$ (see Equation (36)); substituting into Equation (40) yields

$$PI_A(K) = \frac{\tau_H\left(K+\frac{2-\tau_H}{\tau_H}\right)}{2+(K-1)\tau_H+(2-\tau_H)} = \frac{2+(K-1)\tau_H}{4+(K-2)\tau_H}.$$

Under the model-based architecture, the proof of Proposition 4 establishes $u_M \equiv \beta_M^*/b^* = (2+(K-1)\tau_H)/\tau_H$. Substituting into Equation (40),

$$PI_M(K) = \frac{\tau_H\,(K+u_M)}{2+(K-1)\tau_H+u_M\tau_H} = \frac{2+(2K-1)\tau_H}{2\,(2+(K-1)\tau_H)}.$$

Part 1. Cross-multiplying,

$$(2 + (2K - 1)\tau_H)(4 + (K - 2)\tau_H) - 2(2 + (K - 1)\tau_H)^2 = K\tau_H(2 - \tau_H) > 0$$

for $K \geq 1$, so $PI_M(K) > PI_A(K)$. Moreover, one can show that both are increasing in K by differentiating, and are equal to $1/2$ when $K = 0$.

Part 2. Equilibrium price informativeness under architecture a is $PI_a(K_a^*)$. Note that

$$PI_A(K_A^*) = \frac{2 + (K_A^* - 1)\tau_H}{4 + (K_A^* - 2)\tau_H} > \frac{2 + (2K_M^* - 1)\tau_H}{2(2 + (K_M^* - 1)\tau_H)} = PI_M(K_M^*)$$
$$\Leftrightarrow K_A^* > 2K_M^*$$

For the sufficient condition, recall from the proof of Proposition 4 that $\Pi_H(\beta_a; K)$ is strictly decreasing in K, so that $K_M^* = 0$ if and only if $c > \Pi_H^M(1)$, and $K_A^* \geq 1$ if and only if $c \leq \Pi_H^A(1)$, where

$$\Pi_H^M(1) = \frac{\tau_H \sigma_v \sigma_z}{4\sqrt{4 - \tau_H}} \quad \text{and} \quad \Pi_H^A(1) = \frac{\tau_H \sigma_v \sigma_z}{(4 - \tau_H)\sqrt{\tau_H + (2 - \tau_H)^2}},$$

and $\Pi_H^M(1) < \Pi_H^A(1)$ since $\Pi_H^M(K) < \Pi_H^A(K)$ for fixed K, so the interval in condition (26) is nonempty. Under condition (26), $PI_M(K_M^*) = PI_M(0) = \frac{1}{2}$, while

$$PI_A(K_A^*) - \frac{1}{2} = \frac{2 + (K_A^* - 1)\tau_H}{4 + (K_A^* - 2)\tau_H} - \frac{1}{2} = \frac{K_A^* \tau_H}{2(4 + (K_A^* - 2)\tau_H)} > 0$$

for $K_A^* \geq 1$. Hence price informativeness is strictly lower under the model-based architecture.

Part 3. Comparing $PI_a(K^*)$ with price informativeness in the economy with only K human traders, $PI_K = K\tau_H/(2 + (K - 1)\tau_H)$ from Equation (25), yields the result.

A.10 Proof of Proposition 8

Fix K and let $u \equiv \beta_a / b(\beta_a)$. When the cap binds, $\beta_a = \bar{\beta}$. Equation (29) implies

$$b(u) = \frac{\sqrt{\tau_H}\,\sigma_z}{\sigma_v \sqrt{K + u(2 - \tau_H)}},$$

so that $\beta_a = ub(u)$ is strictly increasing in u. Therefore, tightening the cap reduces u. Since $b(u)$ is strictly decreasing in u, human trading intensity increases as the cap is tightened. Equation (32) then implies that human profit also increases.

The algorithmic trader's expected profit from Equation (35) can be written as

$$\Pi(u) = \frac{ub(u)(2-\tau_H)\sigma_v^2}{2+(K-1)\tau_H+u\tau_H}.$$

Recall that for the model-free architecture, we have $u_A = (2-\tau_H)/\tau_H$ (Equation (36)) while for the model-based architecture, we have $u_M = (2+(K-1)\tau_H)/\tau_H$ (Equation (37)), where $u_A \leq u_M$. For $u \leq (2+(K-1)\tau_H)/\tau_H$, which is the range below the unconstrained intensity, its logarithmic derivative is

$$\frac{\Pi'(u)}{\Pi(u)} = \frac{1}{u} - \frac{\tau_H}{2+(K-1)\tau_H+u\tau_H} - \frac{2-\tau_H}{2(K+u(2-\tau_H))} > 0.$$

The inequality follows after multiplying through by positive denominators, since the resulting expression is

$$2\big(2+(K-1)\tau_H\big)K + (2-\tau_H)u\big(2+(K-1)\tau_H - u\tau_H\big),$$

which is strictly concave in u and is equal to $2K(2+(K-1)\tau_H)$ at $u=0$ and $u=u_M$. Thus, the algorithmic trader's profit falls as the cap is tightened.

Finally, price informativeness is

$$PI(u,K) = \lambda\Gamma = \frac{\tau_H(K+u)}{2+(K-1)\tau_H+u\tau_H},$$

which is strictly increasing in u, since $\partial PI/\partial u = \tau_H(2-\tau_H)/(2+(K-1)\tau_H+u\tau_H)^2 > 0$. Thus, holding K fixed, price informativeness falls as the cap is tightened. Human profits rise for every fixed K, so Equation (19) implies that the equilibrium number of human traders weakly increases as the cap is tightened. Moreover, $PI(u,K)$ is strictly increasing in K for fixed u, and $u = \bar{\beta}/b$ also increases when an additional human trader enters because b decreases in K. Therefore, entry generates an upward jump in price informativeness, whereas the intensity effect reduces it continuously between entry thresholds.

B Extensions

B.1 Uncertainty about informed trading

We now provide a micro-foundation for the model-based algorithm in Section 4.2. Suppose that the algorithmic trader is uncertain about the extent to which other traders also trade on information that is correlated with v_t.[34] We show that this uncertainty leads to a higher trading intensity than in the Cournot benchmark, unless the algorithmic trader happens to have exactly the right prior about other traders' behavior. In this sense, our results from Section 4.2 are robust.

Specifically, suppose that the algorithmic trader's price estimate is given by

$$\widehat{p}_t = p_{0,t} + \widehat{\lambda}_t x_{U,t} + \widehat{\phi}_t v_t. \tag{41}$$

Here, the term $\widehat{\phi}_t$ is an estimate for the combined price impact of all other informed traders.[35] The initial value $\widehat{\phi}_0$ represents the algorithmic trader's prior. When $\widehat{\phi}_0 = 0$, the algorithmic trader initially believes that no other traders have information about v_t, whereas when $\widehat{\phi}_0 = \lambda_A^* K_A^* b_A^*$, the algorithmic trader's prior anticipates other traders' equilibrium price impact in the Cournot equilibrium.

We assume that $\widehat{\phi}_0 \in (0, \lambda_A^* K_A^* b_A^*)$. That is, the algorithmic trader's prior does not fully anticipate the price impact of other traders, relative to the Cournot benchmark in Proposition 2. For example, the algorithmic trader initially believes that some traders have information about v_t, but does not know how many there are or how intensely they trade on v_t.

The algorithm can now be summarized by $\Psi_{U,t} = (\beta_{U,t}, p_{0,t}, \widehat{\lambda}_t, \widehat{\phi}_t)$, where $\beta_{U,t}$, $p_{0,t}$, and $\widehat{\lambda}_t$ follow Equations (15) – (17), with the modification that $\widehat{p}_t$ in Equation (41) is the price estimate. The algorithmic trader updates $\widehat{\phi}_t$ according to

$$\widehat{\phi}_{t+1} - \widehat{\phi}_t := \alpha \left(\kappa(p_t - \widehat{p}_t)v_t - (1-\kappa)(\widehat{\phi}_t - \widehat{\phi}_0) \right). \tag{42}$$

Here, $\kappa \in [0,1]$ is a hyperparameter that measures the algorithmic trader's confidence in the

[34] Because of this, the algorithmic trader will underestimate how other traders' information affects the price in equilibrium (see Proposition 9). This is similar to "dismissiveness" in Banerjee (2011) and Banerjee, Davis, and Gondhi (2024a) and "cursedness" in Eyster, Rabin, and Vayanos (2019). In our setting, the dismissiveness is an endogenous outcome of the learning algorithm, whereas in this earlier work, it stems from behavioral assumptions on traders' behavior.

[35] Intuitively, $\widehat{p}_t - \widehat{\lambda} x_{M,t} = \widehat{\phi}_t v_t$ whenever $p_{0,t} = 0$, and for the true price evolution, $p_t - \lambda_t x_{M,t} = \lambda_t \left(\sum_{i=1}^K b_{i,t}(v_t + \varepsilon_{i,t}) + z_t \right)$. In this sense, $\widehat{\phi}_t$ measures the portion of the residual demand that is correlated with v_t.

data. For $\kappa = 0$, the algorithmic trader entirely trusts the prior, and does not attempt to update it from price data. In that case, Equation (42) implies that the steady state must satisfy $\phi^* = \widehat{\phi}_0$. Conversely, for $\kappa = 1$, the algorithmic trader has no trust in the prior and only updates according to data.[36]

The proposition below characterizes the algorithmic trading equilibrium for arbitrary values of $\widehat{\phi}_0$ and κ.

Proposition 9. *In an algorithmic trading equilibrium, $\widehat{\phi}^*$ is given by*

$$\widehat{\phi}^* = \frac{\kappa K_U^* \lambda_U^* b_U^* \sigma_v^2 + (1-\kappa)\widehat{\phi}_0}{\kappa \sigma_v^2 + 1 - \kappa}, \tag{43}$$

β_U^* is given by

$$\beta_U^* = \frac{1 - \widehat{\phi}^*}{2\lambda_U^*}, \tag{44}$$

b_U^* and λ_U^* are given by Equations (3) and (4), and K_U^* satisfies Equation (19), with $\beta = \beta_U^*$.

We have $\beta_U^* = \beta_M^*$ for $\kappa = \widehat{\phi}_0 = 0$ and $\beta_U^* = \beta_A^*$ for $\kappa = 1$ or $\widehat{\phi}_0 = \lambda_A^* K_A^* b_A^*$, where β_A^* and β_M^* are the equilibrium trading intensities in the model-free (Proposition 2) and the model-based benchmark (Proposition 3) respectively.

Whenever $\kappa < 1$ and $\widehat{\phi}_0 < \lambda_A^* K_A^* b_A^*$, $\beta_U^* > \beta_A^*$ and $K_U^* \leq K_A^*$.

Intuitively, the equilibrium value $\widehat{\phi}^*$ is an estimate of human traders' price impact. Since the algorithmic trader is uncertain about whether the prior or the observed data are correct, this estimate is a weighted average between the algorithmic trader's prior $\widehat{\phi}_0$ and human traders' actual price impact, which is given by $K_U^* \lambda_U^* b_U^*$. The parameter κ measures the weight the algorithm places on the observed data vs. the prior. A higher κ moves $\widehat{\phi}^*$ closer to human traders' price impact, whereas a lower κ moves $\widehat{\phi}^*$ closer to the prior. The trading intensity β_U^* takes the price impact estimate $\widehat{\phi}^*$ as given and is a best response to it.[37]

When the prior is that no other trader is informed ($\widehat{\phi}_0 = 0$) and the algorithm disregards the data ($\kappa = 0$), we recover the equilibrium for the model-based algorithm in Proposition 3. If instead the algorithm places no weight on the prior ($\kappa = 1$) we recover the equilibrium for the model-free algorithm in Proposition 2.

The proposition also shows that the qualitative results from the model-based benchmark are robust. Unless the prior somehow anticipates human traders' price impact in the Cournot equilibrium or the algorithm entirely disregards the prior, we have $\beta_U^* > \beta_A^*$. Hence, generically, the algorithm trades more aggressively than a rational Cournot trader would.

[36] Equation (42) is a *constant gain* learning rule. See e.g. Evans and Honkapohja (2001), Ch. 3.3.

[37] Note that if $\widehat{\phi}^* = K_U^* \lambda_U^* b_U^*$, Equation (44) implies that $\beta_U^* = \frac{1 - K_U^* \lambda_U^* b_U^*}{2\lambda_U^*}$. That is, β_U^* is the trading intensity in the Bayes-Nash equilibrium.

B.2 Imperfect information

In the baseline model, we assumed that the algorithmic trader observes the asset value v_t perfectly. In this appendix, we consider an imperfectly informed algorithmic trader, who uses the model-free algorithm of Section 4.1. We focus on the model-free architecture for analytical tractability — we expect results to be qualitatively similar under the model-based architecture. We show that as the algorithm becomes better at extracting information, human traders' profits decrease, and more human traders exit the market. As a result, price informativeness may decrease.

Suppose that the algorithmic trader observes a signal $s_{A,t} = v_t + \eta_t$, where $\eta_t \sim \mathcal{N}(0, \sigma_\eta^2)$ is independent of v_t and $\{\varepsilon_{i,t}\}_{i=1}^K$ and is independent across time. We maintain the assumption that the algorithmic trader has better information than human traders, i.e., $\sigma_\eta^2 \leq \sigma_\varepsilon^2$. As in the baseline model, we assume that the algorithmic trader trades on the signal $s_{A,t}$, i.e., $x_{A,t} = \beta_{A,t} s_{A,t}$. To save notation, we define $\tau_A = \sigma_v^2/(\sigma_v^2 + \sigma_\eta^2)$.

Repeating the analysis in Section 3.3, we can show that for a given β_A and K, the equilibrium is characterized by

$$\lambda = \frac{(Kb + \beta_A)\,\sigma_v^2}{(Kb + \beta_A)^2\,\sigma_v^2 + Kb^2\sigma_\varepsilon^2 + \beta_A^2\sigma_\eta^2 + \sigma_z^2} \tag{45}$$

and

$$b = \frac{1 - \lambda\left((K-1)b + \beta_A\right)}{2\lambda}\tau_H, \tag{46}$$

which follows from adapting Equations (1) and (2).

Proposition 10. *The algorithmic trading equilibrium is given by Equations (45), (46), and*

$$\beta_A^* = \tau_A \frac{1 - b^*\lambda^* K^*}{2\lambda^*}, \tag{47}$$

where K^* is determined by Equation (19).

For fixed K, as σ_η^2 decreases, β_A^* increases, b^* decreases, price informativeness increases and human traders' profit decreases. As K decreases, price informativeness decreases. K^* decreases weakly as σ_η^2 decreases.

Figure 7 illustrates the results. Intuitively, as the algorithmic trader becomes more informed, i.e., σ_η^2 decreases, the algorithm trades more aggressively (β_A^* increases), which causes human traders to trade less aggressively and lowers their profits. Price informativeness increases as long as there is no additional crowding out. Eventually, however, lower profits cause human traders to exit the market, which decreases price informativeness.

Figure 7: Equilibrium profits, trading intensities, and price informativeness as a function of σ_η^2.

Panel (a) plots the algorithmic trader's equilibrium profit Π for different values of σ_η^2, given the strategies in Proposition 10. Panel (b) plots the total human trading profit $K\Pi_H$ (dashed) and the average human trading profit Π_H (solid). Panel (c) plots the equilibrium trading intensities and price impact, while panel (d) plots the equilibrium price informativeness.

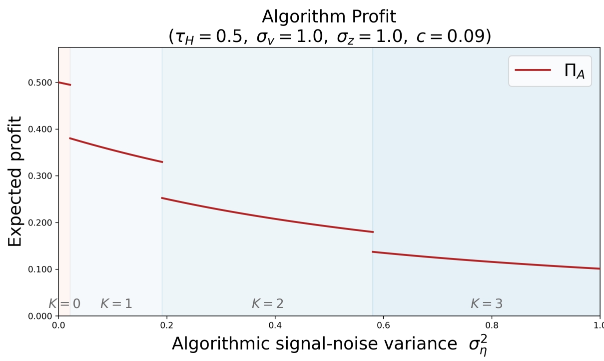

(a) Π versus σ_η^2

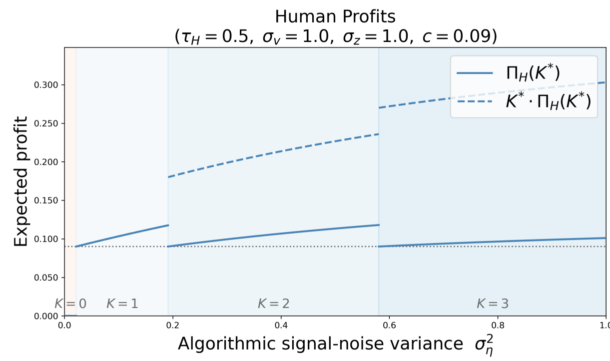

(b) Π_H versus σ_η^2

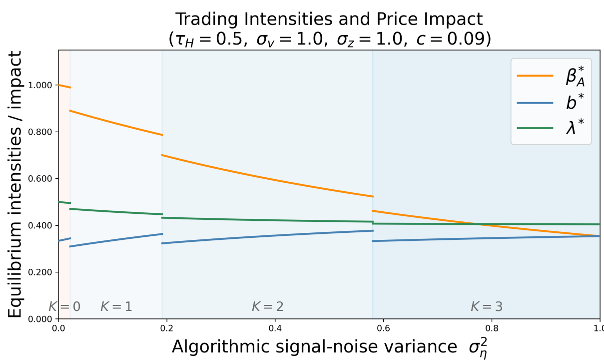

(c) β_A, b, and λ versus σ_η^2

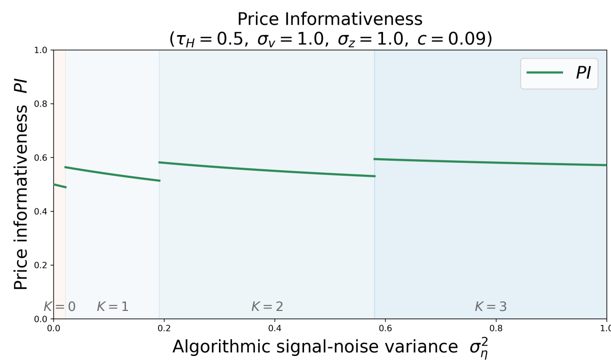

(d) PI versus σ_η^2

B.3 Targeted versus blanket transaction taxes

Transaction taxes are the oldest instrument in this debate (Tobin (1978)), and have been analyzed extensively in settings without algorithmic traders (e.g., Subrahmanyam (1998), Dávila (2023)). The distinguishing feature of recent proposals is that they target algorithmic trading specifically. France introduced a levy on high-frequency traders' canceled orders alongside its 2012 transaction tax,[38] and Italy's 2013 financial transaction tax includes a

[38] Article 235 ter ZD bis of French General Tax Code (introduced by Law No. 2012-354 of March 14, 2012) imposes a 0.01% levy on high-frequency traders' orders that are canceled or modified beyond a threshold ratio. The levy applies only to firms operating in France and exempts market making. See Colliard and Hoffmann (2017) for an analysis of the French tax and its effect on market composition.

levy on high-frequency order modifications.[39] Blanket taxes on all transactions, by contrast, have been proposed at the EU level and in the United States,[40] and were implemented in Sweden in the 1980s (Umlauf (1993)). We compare the two designs within our framework and focus on the model-free architecture — we expect the results under the model-based architecture to be analogous, but analytically less tractable.

We model a *targeted* tax at rate $\theta \geq 0$ by assuming that the algorithmic trader pays $(1+\theta)p_t$ per share, so that its per-period profit is $x_{A,t}\left(v_t - (1+\theta)p_t\right)$, while human traders are untaxed.[41] Under a *blanket* tax, human traders' per-period profits are likewise $x_{i,t}(v_t-(1+\theta)p_t) - c$, and noise traders also pay the tax (their demand is inelastic). The market maker's pricing rule is unaffected, since the tax is paid to the government rather than to the market maker. Definition 2 applies throughout with profits replaced by after-tax profits. Note that the model-free algorithm requires no knowledge of the tax: since the tax enters realized profits $\pi_{A,t}$, the actor-critic updating in Equations (9)–(11) internalizes it automatically, and the steady state (48) is the best response to (b^*, λ^*) given the tax.

The response of the model-based algorithm of Section 4.2, by contrast, depends on whether its conjectured profit model accounts for the tax. Predicting the market's response to a given policy therefore requires knowing the *architecture* of the algorithms being regulated — a version of the Lucas critique for algorithmic trading. As such, in what follows, we focus on the model-free architecture.

The following result characterizes the equilibrium with a targeted tax; recall that PI_K, given in Equation (25), denotes price informativeness in the economy with K human traders and no algorithmic trader.

Proposition 11. *Suppose the algorithmic trader pays a transaction tax at rate $\theta \geq 0$ and human traders are untaxed.*

1. *In an algorithmic trading equilibrium,*

$$\beta_A^* = \max\left\{\frac{1-(1+\theta)\,b^*\lambda^* K^*}{2(1+\theta)\lambda^*}, 0\right\}, \tag{48}$$

[39] Law No. 228 of December 24, 2012, imposes a 0.02% levy on orders modified or canceled within half a second beyond a threshold order-to-trade ratio. See Cappelletti, Guazzarotti, and Tommasino (2017) for an evaluation of the Italian tax.

[40] European Commission proposals COM(2011) 594 and COM(2013) 71 envision a broad-based tax of 0.1% on securities transactions. In the United States, see, e.g., the Wall Street Tax Act of 2019 (S. 647, 116th Congress), proposing a 0.1% tax, and the Inclusive Prosperity Act of 2019 (S. 1587, 116th Congress), proposing a 0.5% tax on stocks.

[41] Since prices are measured relative to the unconditional mean of the asset value, this specification is a tax on signed transaction value: the expected tax payment $\theta\mathbb{E}[x_{A,t}p_t] = \theta\lambda\beta_A\Gamma\sigma_v^2 > 0$ is proportional to the trader's contribution to price impact. A tax on absolute transaction value, $\theta\mathbb{E}[|x_{A,t}p_t|]$, has similar qualitative implications but breaks the linearity of best responses.

where b^* and λ^* are given by Equations (3) and (4), and K^* satisfies Equation (19).

2. *Fix $K \geq 1$ and let $\bar{\theta}(K) \equiv \frac{2-\tau_H}{K\tau_H}$. For $\theta < \bar{\theta}(K)$, we have $\beta_A^* > 0$; moreover, β_A^* and Π are strictly decreasing in θ, while b^* and Π_H are strictly increasing in θ. Price informativeness is given by*

$$PI(\theta, K) = \frac{(1+\theta)K\tau_H + 2 - \tau_H}{(1+\theta)\left(4 + (K-2)\tau_H\right)}, \tag{49}$$

   which is strictly decreasing in θ, with $PI\left(\bar{\theta}(K), K\right) = PI_K$. For $\theta \geq \bar{\theta}(K)$, $\beta_A^* = 0$ and the equilibrium coincides with the economy with K human traders only.

3. *With endogenous participation, $K^*(\theta)$ is weakly increasing in θ, and $PI(\theta, K)$ is strictly increasing in K for fixed θ.*

For a fixed number of human traders, the impact of the targeted tax is standard. The algorithm trades less aggressively, human traders trade more aggressively and earn higher profits, but the information lost from the (better informed) algorithm's trades is not fully replaced by the humans' trades: price informativeness falls. Holding participation fixed, taxing the better-informed trader therefore reduces price informativeness, in line with the standard intuition that discouraging informed trading makes prices less efficient.

The participation margin can overturn this conclusion. Because the tax raises human traders' profits, it induces entry, and each entry raises price informativeness discretely. Equilibrium price informativeness therefore follows a saw-tooth pattern in θ: it declines continuously between entry thresholds and jumps upward at each threshold. Figure 8 illustrates the saw-tooth pattern of $PI(\theta)$.

A targeted tax improves price informativeness if and only if the participation effect dominates the intensity effect. For instance, suppose Proposition 5 holds, so that absent the tax the algorithmic trader crowds out all human traders and $PI = \frac{1}{2}$. For taxes below the first entry threshold, informativeness falls (i.e., $PI(\theta, 0) = \frac{1}{2(1+\theta)}$), but it jumps up once the first human trader enters, and exceeds its laissez-faire level whenever entry occurs at $\theta < \frac{\tau_H}{4-3\tau_H}$, which obtains when c is close to the bound in Proposition 5. Conversely, a sufficiently large tax ($\theta \geq \bar{\theta}(K^*)$) drives the algorithm out of the market entirely, and whether this improves on laissez-faire is precisely the comparison in Proposition 6.

The next result characterizes the equilibrium when the tax cannot be targeted and applies to all traders.

Proposition 12. *Suppose all traders pay the transaction tax at rate $\theta \geq 0$.*

Figure 8: Price informativeness versus the targeted transaction tax

The solid line corresponds to the equilibrium price informativeness $PI(\theta)$ with endogenous participation, as a function of the tax rate θ levied on the algorithmic trader. The dot-dashed, horizontal line corresponds to price informativeness in the economy with only human traders, while the dashed horizontal line corresponds to the equilibrium price informativeness when only the algorithmic trader is present and the tax is zero ($PI = \frac{1}{2}$). Shaded bands indicate the equilibrium number of human traders K^*. Parameters are $\tau_H = 0.5$, $\sigma_v = \sigma_z = 1$, and $c = 0.09$, for which the algorithmic trader crowds out all human traders at $\theta = 0$.

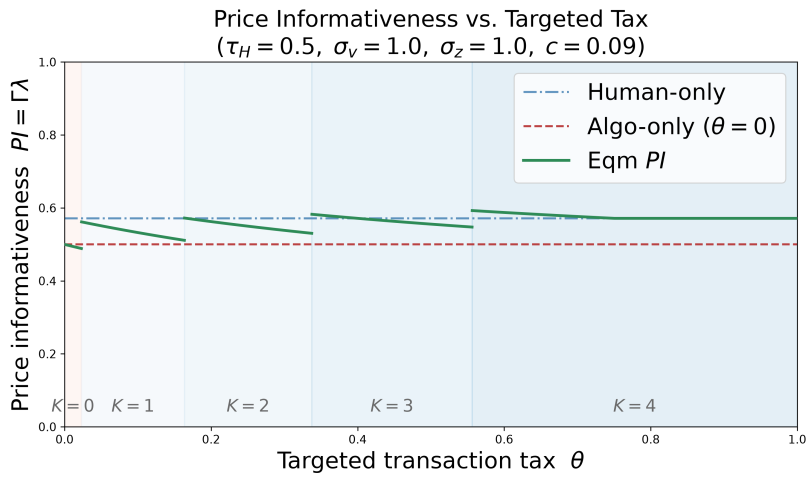

1. *For fixed K, the ratio of trading intensities is independent of the tax and equal to its value without taxes, $\frac{\beta_A^*}{b^*} = \frac{2-\tau_H}{\tau_H}$, and*

$$b^*(\theta) = \frac{\tau_H \sigma_z}{\sigma_v \sqrt{K\tau_H + (2-\tau_H)^2 + \theta\,(2+(K-1)\tau_H)\,(4+(K-2)\tau_H)}}. \tag{50}$$

2. *For fixed K, the intensities β_A^* and b^* and the profits Π and Π_H are strictly decreasing in θ, and price informativeness is given by*

$$PI(\theta, K) = \frac{1}{1+\theta}\,\frac{2+(K-1)\tau_H}{4+(K-2)\tau_H} = \frac{PI(0,K)}{1+\theta}. \tag{51}$$

3. *With endogenous participation, $K^*(\theta)$ is weakly decreasing in θ. Price informativeness therefore declines in the tax through both the intensity and the participation margin.*

The comparison with Proposition 11 isolates what a targeted tax accomplishes. A blanket tax scales down all informed trading uniformly so price informativeness falls mechanically at rate $\frac{1}{1+\theta}$. Because after-tax human profits fall, human traders exit, and informativeness falls further. Unlike a targeted tax, which changes the composition of informed trading in favor of human investors, a blanket tax unambiguously reduces price informativeness in our setting.[42]

In practice, the targeted taxes are difficult to implement. Regulation identifies algorithmic trading through observable proxies, e.g., message rates, order-to-trade ratios, order resting times, and technological infrastructure.[43] But sophisticated traders can modify their behavior to avoid restrictions: traders can slow message rates to just below the threshold, route order flow through exempt categories such as registered market makers, or relocate activity to untaxed venues, as documented for Sweden by Umlauf (1993) and for France by Colliard and Hoffmann (2017). Moreover, the boundary between algorithmic and human trading is increasingly blurred, since most human-originated institutional order flow is executed by algorithms, and so a "targeted" tax would reduce human trading too.

B.4 Proofs of Extensions

B.4.1 Proof of Proposition 9

Using Equation (41), we can write the profit estimate as

$$Q_t(\tilde{\beta}_{U,t}) = \tilde{\beta}_{U,t} v_t (v_t - p_{0,t} - \widehat{\lambda}_t \tilde{\beta}_{U,t} v_t - \widehat{\phi}_t v_t),$$

where we denote the perturbed policy by

$$\tilde{\beta}_{U,t} = \beta_{U,t} + \varepsilon_{U,t}$$

and $\varepsilon_{U,t}$ is the perturbation noise. As in Section 4.2, we focus on the small noise benchmark. The state of the algorithm consists of $\Psi_{U,t} = (\beta_{U,t}, p_{0,t}, \widehat{\lambda}_t, \widehat{\phi}_t)$ and we denote a steady state with Ψ_U^*.

The steady state consists of Equations (18), into which we plug in the price estimate in Equation (41), and

$$\mathbb{E}[\widehat{\phi}_{t+1} - \widehat{\phi}_t \mid \Psi_{U,t} = \Psi_U^*] = 0.$$

[42]In the model, noise trading is inelastic. If noise traders also reduced their trading in response to a blanket tax (a decline in σ_z), traders' profits would fall further (Equation (6)), amplifying exit. Declines in volume following broad transaction taxes are well documented empirically (Umlauf (1993), Colliard and Hoffmann (2017)).

[43]Directive 2014/65/EU, Article 4(1)(39)–(40), defines algorithmic and high-frequency trading via such proxies; see also Commission Delegated Regulation (EU) 2017/565, Article 19. The French levy applies only to firms operating in France and exempts market making; the Italian levy applies only above a threshold order-to-trade ratio.

Plugging in Equations (41) and (42) yields

$$\mathbb{E}\left[v_t\left(v_t - p_{0,t} - 2\widehat{\lambda}_t(\beta_{U,t}+\varepsilon_{U,t})v_t - \widehat{\phi}_t v_t\right)|\Psi_{U,t}=\Psi_U^*\right]=0,$$

$$\mathbb{E}\left[\lambda_t\left(b_t\sum_{i=1}^{K}s_{i,t}+z_t\right)-p_{0,t}+\left(\lambda_t-\widehat{\lambda}_t\right)\left(\beta_{U,t}+\varepsilon_{U,t}\right)v_t-\widehat{\phi}_t v_t\,|\Psi_{U,t}=\Psi_U^*\right]=0,$$

$$\mathbb{E}\left[\left(\lambda_t\left(b_t\sum_{i=1}^{K}s_{i,t}+z_t\right)-p_{0,t}+\left(\lambda_t-\widehat{\lambda}_t\right)\left(\beta_{U,t}+\varepsilon_{U,t}\right)v_t-\widehat{\phi}_t v_t\right)\varepsilon_{U,t}v_t\,|\Psi_{U,t}=\Psi_U^*\right]=0,$$

and

$$\begin{aligned}\mathbb{E}\Bigg[\Bigg(\kappa\left(\lambda_t\left(b_t\sum_{i=1}^{K}s_{i,t}+z_t\right)-p_{0,t}+\left(\lambda_t-\widehat{\lambda}_t\right)\left(\beta_{U,t}+\varepsilon_{U,t}\right)v_t-\widehat{\phi}_t v_t\right)v_t \\ -(1-\kappa)(\widehat{\phi}_t-\widehat{\phi}_0)\Bigg)\mid\Psi_{U,t}=\Psi_U^*\Bigg] &= 0.\end{aligned}$$

Combining these equations yields $p_0^*=0$, $\widehat{\lambda}^*=\lambda_U^*$, and $\widehat{\phi}^*$ and β_U^* as given in Equations (43) and (44).

We now derive a closed form solution for the equilibrium. Write Equation (43) as

$$\phi^* = uK_U^*\lambda_U^*b_U^* = \frac{\kappa K_U^*\lambda_U^*b_U^*\sigma_v^2+(1-\kappa)\widehat{\phi}_0}{\kappa\sigma_v^2+1-\kappa} \tag{53}$$

for some parameter $u\geq 0$. Then, combining Equations (43) and (44) with Equations (3) and (4) yields for fixed K:

$$\lambda(u)=\frac{\sigma_v\sqrt{(1+(2-u)K)\sigma_v^4+((3-2u)K+4)\sigma_v^2\sigma_\varepsilon^2+4\sigma_\varepsilon^4}}{\sigma_z\left(4\sigma_\varepsilon^2+2\sigma_v^2+\sigma_v^2(2-u)K\right)},$$

$$b(u)=\frac{\sigma_v\sigma_z}{\sqrt{(1+(2-u)K)\sigma_v^4+((3-2u)K+4)\sigma_v^2\sigma_\varepsilon^2+4\sigma_\varepsilon^4}},$$

and

$$\beta_U(u)=\frac{\sigma_z\left((1+K(1-u))\sigma_v^2+2\sigma_\varepsilon^2\right)}{\sigma_v\sqrt{(1+(2-u)K)\sigma_v^4+((3-2u)K+4)\sigma_v^2\sigma_\varepsilon^2+4\sigma_\varepsilon^4}}.$$

Equation (53) implies that

$$u=\frac{\kappa K\sigma_v^4+2(1-\kappa)\widehat{\phi}_0\left((K+1)\sigma_v^2+2\sigma_\varepsilon^2\right)}{\kappa K\sigma_v^4+(1-\kappa)K(1+\widehat{\phi}_0)\sigma_v^2} \tag{54}$$

in equilibrium.

For $\widehat{\phi}_0 = \kappa = 0$, we have $u = 0$ and plugging in $K = K_U^*$ via Equation (19) implies that $\beta_U(0) = \beta_M^*$. For $\kappa = 1$, we have $u = 1$ and similarly $\beta_U(1) = \beta_A^*$. For $\kappa < 1$, plugging in $\widehat{\phi}_0 = \lambda_A^* K_A^* b_A^*$ implies that $u = 1$ and $\beta_U^* = \beta_A^*$, which follows by guess-and-verify.

It remains to establish that $\beta_U^* > \beta_A^*$ whenever $\widehat{\phi}_0 < \lambda_A^* K_A^* b_A^*$ and $\kappa < 1$. To this end, we can calculate each human trader's profit as

$$\Pi_H^U(u, K) = \frac{\sigma_v^3 \sigma_z \left(\sigma_v^2 + \sigma_\varepsilon^2\right)}{\left(4\sigma_\varepsilon^2 + 2\sigma_v^2 + \sigma_v^2(2-u)K\right)} \frac{1}{\sqrt{(1+(2-u)K)\sigma_v^4 + ((3-2u)K+4)\sigma_v^2\sigma_\varepsilon^2 + 4\sigma_\varepsilon^4}}.$$

Plugging in u from Equation (54) yields a function $\Pi_H^U(\widehat{\phi}_0, K)$, which is strictly increasing in $\widehat{\phi}_0$ and strictly decreasing in K. For $\widehat{\phi}_0 = \lambda_A^* K_A^* b_A^*$, guess-and-verify implies that $K_U^* = K_A^*$, and since $\Pi_H^U(\widehat{\phi}_0, K)$ is increasing in $\widehat{\phi}_0$, we have $K_U^* \leq K_A^*$ for any $\widehat{\phi}_0 < \lambda_A^* K_A^* b_A^*$.

Next, suppose by way of contradiction that $\beta_U^* \leq \beta_A^*$. Plugging in u in Equation (54) into $\beta(u)$ and taking derivatives implies that β_U is strictly decreasing in $\widehat{\phi}_0$ for fixed K and $\kappa < 1$. Thus, for fixed K, $\beta_U^* > \beta_U(1) = \beta_A^*$. This immediately implies a contradiction for $K_U^* = K_A^*$. Suppose instead that $K_A^* > K_U^*$. Then, since Proposition 1 implies that $\Pi_H(\beta, K)$ is decreasing in both β and K, it holds that $\Pi_H(\beta_U^*, K_U^*) > \Pi_H(\beta_U^*, K_A^*) > \Pi_H(\beta_A^*, K_A^*)$, which in turn implies that $K_U^* \geq K_A^*$, another contradiction. Hence, it must be the case that $\beta_U^* > \beta_A^*$ and $K_U^* \leq K_A^*$.

B.4.2 Proof of Proposition 10

The steady state conditions in Equation (12) imply that in the small noise limit,

$$\frac{E\left[\pi_{A,t}\varepsilon_{A,t} | \Psi_t = \Psi^*\right]}{\sigma_A^2} \to \sigma_v^2 - \lambda\left(\beta_A^*\right)\left(2\beta_A^*\left(\sigma_v^2 + \sigma_\eta^2\right) + Kb\left(\beta_A^*\right)\sigma_v^2\right),$$

which implies Equation (47). Combining this equation with Equations (45) and (46) then yields the closed-form solutions

$$\beta_A^* = \sqrt{\frac{\left(\sigma_v^2 + 2\sigma_\varepsilon^2\right)^2 \sigma_z^2}{(K^*+1)\sigma_v^6 + \sigma_v^4\left(\left(4K^*+1\right)\sigma_\eta^2 + (K^*+4)\sigma_\varepsilon^2\right) + 4\left(\sigma_v^2\left(\sigma_\eta^2 + \sigma_\varepsilon^2\right) + \sigma_\eta^2\sigma_\varepsilon^2\right)\left(K^*\sigma_\eta^2 + \sigma_\varepsilon^2\right)}},$$

$$\lambda^* = \frac{\sigma_v^2 + 2\sigma_\varepsilon^2}{\sigma_v^2\left(2(K^*+1)\sigma_\eta^2 + (K^*+2)\sigma_v^2\right) + 4\sigma_\varepsilon^2\left(\sigma_\eta^2 + \sigma_v^2\right)} \frac{\sigma_v^2}{\beta_A^*},$$

and

$$b^* = \frac{\sigma_v^2 + 2\sigma_\eta^2}{\sigma_v^2 + 2\sigma_\varepsilon^2}\beta_A^*,$$

where K^* is again determined by Equation (19).

The results that β_A^* is decreasing in σ_η^2 and K^*, and that b^* is increasing in σ_η^2 and decreasing in K^* follow directly from the closed form expressions.

The price informativeness in Equation (23) becomes

$$PI = \frac{\left(\beta_A + K^* b\right)^2 \sigma_v^4}{\left(\beta_A + K^* b\right)^2 \sigma_v^2 + \beta_A^2 \sigma_\eta^2 + K^* b^2 \sigma_\varepsilon^2 + \sigma_z^2}$$

and plugging in the solution yields the closed form

$$PI = \frac{\sigma_v^2 \left(K^* \left(2\sigma_\eta^2 + \sigma_v^2\right) + \sigma_v^2 + 2\sigma_\varepsilon^2\right)}{\sigma_v^2 \left(2(K^*+1)\sigma_\eta^2 + (K^*+2)\sigma_v^2\right) + 4\sigma_\varepsilon^2 \left(\sigma_\eta^2 + \sigma_v^2\right)}.$$

We can then directly calculate

$$\frac{\partial PI}{\partial \sigma_\eta^2} < 0,$$

and the closed form implies that PI is increasing in K^*. Finally, the human traders' profit is given by

$$\Pi_H = \lambda b^2 \frac{\sigma_v^2}{\tau_H},$$

and plugging in the solution above yields the closed form

$$\Pi_H = \frac{\sigma_v^2 \left(2\sigma_\eta^2 + \sigma_v^2\right)^2 \left(\sigma_v^2 + \sigma_\varepsilon^2\right) \sqrt{\frac{\sigma_z^2}{K^* \sigma_v^2 \left(2\sigma_\eta^2 + \sigma_v^2\right)^2 + K^* \sigma_\varepsilon^2 \left(2\sigma_\eta^2 + \sigma_v^2\right)^2 + \sigma_v^4 \left(\sigma_\eta^2 + \sigma_v^2\right) + 4\sigma_v^2 \sigma_\varepsilon^2 \left(\sigma_\eta^2 + \sigma_v^2\right) + 4\sigma_\varepsilon^4 \left(\sigma_\eta^2 + \sigma_v^2\right)}}}{\sigma_v^2 \left(2(K^*+1)\sigma_\eta^2 + (K^*+2)\sigma_v^2\right) + 4\sigma_\varepsilon^2 \left(\sigma_\eta^2 + \sigma_v^2\right)}.$$

Then, differentiating yields

$$\frac{\partial \Pi_H}{\partial \sigma_\eta^2} > 0,$$

which implies that K^* must weakly decrease as σ_η^2 decreases.

B.4.3 Proof of Proposition 11

Throughout, let $u \equiv \beta_A / b$, and note that $4 + (K-2)\tau_H = \big(2 + (K-1)\tau_H\big) + \big(2 - \tau_H\big)$.

Part 1. The algorithmic trader's after-tax realized profit is $\pi_{A,t} = x_{A,t}(v_t - (1+\theta)p_t)$. Repeating the argument in the proof of Proposition 2, in the small noise limit

$$\frac{\mathbb{E}\left[\pi_{A,t}\varepsilon_{A,t} | \Psi_t = \Psi^*\right]}{\sigma_A^2} \to \sigma_v^2 \left(1 - (1+\theta)\lambda(\beta_A^*) \left(2\beta_A^* + K b(\beta_A^*)\right)\right),$$

and the steady-state condition $w^* = 0$ yields the interior solution in Equation (48). Since human traders and the market maker are untaxed, Equations (1) and (2), and therefore

Equations (29) and (28), continue to hold. For the corner, note that at $\beta_A = 0$ the equilibrium is the K-human economy, in which Equation (28) implies $\lambda b = \tau_H / \big(2 + (K-1)\tau_H\big)$; the expected drift of $\beta_{A,t}$ at zero is then proportional to $1-(1+\theta)K\tau_H/\big(2+(K-1)\tau_H\big)$, which is non-positive if and only if $\theta \geq \bar{\theta}(K) = (2-\tau_H)/(K\tau_H)$ (using $\big(2+(K-1)\tau_H\big)-K\tau_H = 2-\tau_H$). Since $\beta_A \in \mathbb{R}_+$, $\beta_A^* = 0$ is a steady state exactly when $\theta \geq \bar{\theta}(K)$.

Part 2. Fix $K \geq 1$. Writing Equation (28) as $\lambda b\big(2 + (K-1)\tau_H + u\tau_H\big) = \tau_H$ and the interior steady-state condition as $(1+\theta)\lambda b(2u+K) = 1$, and dividing the two, we obtain $(1+\theta)\tau_H(2u+K) = 2 + (K-1)\tau_H + u\tau_H$, i.e.

$$u^*(\theta) = \frac{2 - \tau_H - \theta K \tau_H}{(1+2\theta)\,\tau_H}, \tag{55}$$

which is strictly positive if and only if $\theta < \bar{\theta}(K)$. Differentiating,

$$\frac{du^*}{d\theta} = -\frac{K\tau_H + 2(2-\tau_H)}{(1+2\theta)^2\tau_H} < 0.$$

From Equation (29) with $\beta_A = ub$,

$$b(u) = \frac{\sqrt{\tau_H}\,\sigma_z}{\sigma_v\sqrt{K + u(2-\tau_H)}}, \tag{56}$$

which is strictly decreasing in u; hence b^* is strictly increasing in θ. Next, $\beta_A^* = u\,b(u)$ satisfies

$$\frac{d}{du}\left[\frac{u}{\sqrt{K+u(2-\tau_H)}}\right] = \frac{2K + u(2-\tau_H)}{2\left(K+u(2-\tau_H)\right)^{3/2}} > 0,$$

so β_A^* is strictly increasing in u and therefore strictly decreasing in θ.

Since human traders are untaxed, the derivation of Equation (5) applies, and using Equation (28),

$$\Pi_H = \frac{\lambda b^2 \sigma_v^2}{\tau_H} = \frac{b\,\sigma_v^2}{2 + (K-1)\tau_H + u\tau_H}. \tag{57}$$

As θ increases, b increases and u decreases, so Π_H is strictly increasing in θ. For the algorithmic trader, the after-tax expected profit is $\Pi = \beta_A\sigma_v^2\,(1-(1+\theta)\lambda\Gamma)$, and the steady-state condition $(1+\theta)\lambda(2\beta_A + Kb) = 1$ implies $1-(1+\theta)\lambda\Gamma = (1+\theta)\lambda\beta_A$, so that

$$\Pi = (1+\theta)\lambda\,\beta_A^2\,\sigma_v^2 = \frac{u^2\,b\,\sigma_v^2}{2u+K} = \sqrt{\tau_H}\,\sigma_v\sigma_z\,g(u), \qquad g(u) \equiv \frac{u^2}{(2u+K)\sqrt{K+u(2-\tau_H)}},$$

using $(1+\theta)\lambda = 1/\left(b(2u+K)\right)$ and Equation (56). Since

$$\frac{g'(u)}{g(u)} = \frac{2}{u} - \frac{2}{2u+K} - \frac{2-\tau_H}{2(K+u(2-\tau_H))} \geq \frac{2}{u} - \frac{1}{u} - \frac{1}{2u} = \frac{1}{2u} > 0,$$

where the inequality uses $2u+K \geq 2u$ and $K+u(2-\tau_H) \geq u(2-\tau_H)$, the profit Π is strictly increasing in u and therefore strictly decreasing in θ.

For price informativeness, the market maker is untaxed, so as in the proof of Lemma 3, $PI = \lambda\Gamma = \tau_H(K+u)/\big(2+(K-1)\tau_H + u\tau_H\big)$. Substituting Equation (55) and simplifying, using

$$K+u^* = \frac{(1+\theta)K\tau_H + (2-\tau_H)}{(1+2\theta)\tau_H} \qquad \text{and} \qquad 2+(K-1)\tau_H+u^*\tau_H = \frac{(1+\theta)\left(4+(K-2)\tau_H\right)}{1+2\theta},$$

yields Equation (49). Rewriting it as

$$PI(\theta,K) = \frac{K\tau_H}{4+(K-2)\tau_H} + \frac{2-\tau_H}{(1+\theta)\left(4+(K-2)\tau_H\right)}$$

shows that PI is strictly decreasing in θ. At $\theta = \bar{\theta}(K)$, we have $u^* = 0$ and $PI = K\tau_H/\big(2+(K-1)\tau_H\big) = PI_K$, which coincides with Equation (25), so PI is continuous at the exit threshold.

Part 3. Substituting Equations (55) and (56) into Equation (57) yields the closed form

$$\Pi_H(\theta,K) = \frac{\tau_H\,\sigma_v\sigma_z\,(1+2\theta)^{3/2}}{(1+\theta)\left(4+(K-2)\tau_H\right)\sqrt{K\tau_H(1+\theta\,\tau_H)+(2-\tau_H)^2}}.$$

Both K-dependent terms in the denominator are strictly increasing in K, so $\Pi_H(\theta,K)$ is strictly decreasing in K; and it is strictly increasing in θ by Part 2. Since Equation (19) defines $K^*(\theta)$ as the largest K with $\Pi_H(\theta,K) \geq c$, it follows that $K^*(\theta)$ is weakly increasing in θ. Finally, treating K as continuous in Equation (49),

$$\frac{\partial PI}{\partial K} = \frac{\tau_H\,(2-\tau_H)\,(1+2\theta)}{(1+\theta)\left(4+(K-2)\tau_H\right)^2} > 0.$$

B.4.4 Proof of Proposition 12

We retain the notation $u \equiv \beta_A/b$ from the previous proof.

Part 1. Under the blanket tax, human trader i's conditional expected profit is

$$\mathbb{E}\left[x_{i,t}\left(v_t - (1+\theta)p_t\right) \mid s_{i,t}\right] = x_{i,t}\left(\tau_H s_{i,t} - (1+\theta)\lambda\left((\Gamma-b)\,\tau_H s_{i,t} + x_{i,t}\right)\right),$$

as in Section 3.3 with λ replaced by $(1+\theta)\lambda$. The first-order condition and matching with $x_{i,t} = bs_{i,t}$ yield

$$(1+\theta)\,\lambda\,b\,(2+(K-1)\tau_H + u\tau_H) = \tau_H, \tag{58}$$

where we used $b(2-\tau_H)+\Gamma\tau_H = b(2+(K-1)\tau_H+u\tau_H)$. The algorithmic trader's steady-state condition is derived exactly as in the proof of Proposition 11:

$$(1+\theta)\,\lambda\,b\,(2u+K) = 1. \tag{59}$$

Dividing Equation (58) by Equation (59) gives $\tau_H(2u+K) = 2+(K-1)\tau_H+u\tau_H$, i.e., $u\tau_H = \big(2+(K-1)\tau_H\big) - K\tau_H = 2-\tau_H$, so

$$u^* = \frac{2-\tau_H}{\tau_H}$$

for every $\theta \geq 0$, which is the ratio of intensities without taxes (Equation (55) at $\theta = 0$).

The market maker's pricing rule (1) is unchanged. Substituting λ from Equation (58) into Equation (1) and rearranging,

$$b^2\left\{(1+\theta)(K+u)(2+(K-1)\tau_H+u\tau_H)\sigma_v^2 - \tau_H\left((K+u)^2\sigma_v^2 + K\sigma_\varepsilon^2\right)\right\} = \tau_H\sigma_z^2.$$

Evaluating at $u^* = (2-\tau_H)/\tau_H$ and rearranging yields

$$(b^*)^2 = \frac{\tau_H^2\,\sigma_z^2}{\sigma_v^2\left(K\tau_H + (2-\tau_H)^2 + \theta\left(2+(K-1)\tau_H\right)(4+(K-2)\tau_H)\right)},$$

which is Equation (50).

Part 2. From Equation (50), b^* is strictly decreasing in θ, and so is $\beta_A^* = u^*b^*$ since u^* is constant. For after-tax profits, the same steps as in the proof of Proposition 11 (with the human trader's first-order condition (58) in place of (28)) give

$$\Pi_H = (1+\theta)\,\lambda\,b^2\,\frac{\sigma_v^2}{\tau_H} = \frac{b\,\sigma_v^2}{2+(K-1)\tau_H+u^*\tau_H} = \frac{b\,\sigma_v^2}{4+(K-2)\tau_H}$$
$$\Pi = (1+\theta)\,\lambda\,\beta_A^2\,\sigma_v^2 = \frac{(u^*)^2\,b\,\sigma_v^2}{2u^*+K},$$

and both are strictly decreasing in θ through b^*. For price informativeness, using the market maker's (untaxed) price impact from Equation (58), $\lambda = \tau_H/\left((1+\theta)b\,(4+(K-2)\tau_H)\right)$, we obtain

$$PI = \lambda\Gamma = \frac{\tau_H\,(K+u^*)}{(1+\theta)\,(4+(K-2)\tau_H)} = \frac{1}{1+\theta}\,\frac{2+(K-1)\tau_H}{4+(K-2)\tau_H},$$

which is Equation (51), since $\big(2 + (K-1)\tau_H\big)/\,(4 + (K-2)\tau_H) = PI(0, K)$ by Equation (49) evaluated at $\theta = 0$.

Part 3. From Part 2, $\Pi_H(\theta, K) = b^*(\theta)\sigma_v^2/\,(4 + (K-2)\tau_H)$ is strictly decreasing in θ, and it is strictly decreasing in K since both $4 + (K-2)\tau_H$ and the denominator of Equation (50) are strictly increasing in K. By Equation (19), $K^*(\theta)$ is therefore weakly decreasing in θ. Since $PI(\theta, K)$ in Equation (51) is strictly increasing in K (Proposition 11, Part 3, at $\theta = 0$, scaled by $\frac{1}{1+\theta}$) and strictly decreasing in θ, price informativeness declines as the tax rises through both margins.

C Simulation results

In this section, we provide simulation results for the model under both the model-based and model-free architectures. The algorithms converge to the respective equilibrium in Propositions 2 and 3.

For the simulation results, we assume that the learning rate decays exponentially, so that $\alpha_t = \alpha e^{-\rho t}$. This specification is commonly used in work simulating economies with Q-learners, e.g., Colliard et al. (2022) and Dou et al. (2023).

For each specification, we pick a horizon of $T = 10^6$ episodes and we average the results over 1000 runs. We initialize the model-based algorithm at $\beta_{M,0} = 0.5$, $p_{0,0} = 0$, and $\widehat{\lambda}_0 = 0.5$, and the model-free algorithm at $\beta_{A,0} = 0.5$, $w_0 = 0$, and $q_0 = 0$. We pick the same parameter values as in Figure 6, i.e., $\sigma_v = \sigma_z = \sigma_\varepsilon = 1$, so that $\tau_H = 0.5$, and $c = 0.04$. We set the hyperparameters to $\alpha = 0.05$, $\gamma = 0.9$, and $\sigma_A^2 = \sigma_M^2 = 0.05$.

At each t, for a given β_t chosen by either algorithm, we approximate the human traders' intensity $b(\beta_t)$, the market maker's $\lambda(\beta_t)$, and human traders' expected profit $\Pi_H(\beta_t)$ from Proposition 1.[44] We then use Equation (19) to determine K_t. For the model-free algorithm, we use Equations (7)–(10) to calculate the evolution of the algorithm. For the model-based algorithm, we use Equations (15)–(17).

Figure 9 illustrates the results for the model-free algorithm. The algorithm's trading intensity $\beta_{A,t}$ converges to β_A^*, and human traders play their best response to this trading intensity, yielding the same crowding out and price informativeness as in Proposition 2.

Figure 10 illustrates the analog for the model-based algorithm, and shows that algorithmic trading intensity and human participation converge to the same values as in Proposition 3.

[44]Strictly, since the algorithm's trading strategy is stochastic, the order flow follows a mixture-of-normals distribution, and so we are no longer ensured a linear equilibrium. We use the equilibrium of the zero-noise limit as an approximation to the equilibrium in the finite, but small noise case.

Figure 9: Simulation results for model-free algorithm

The solid lines correspond to the mean values of $\beta_{A,t}$, and K_t, averaged across runs in each period. The dashed lines represent β_A^* and K^* in Proposition 2. The bands represent the 95% confidence bands around the means. We pick $\rho = 10^{-5}$.

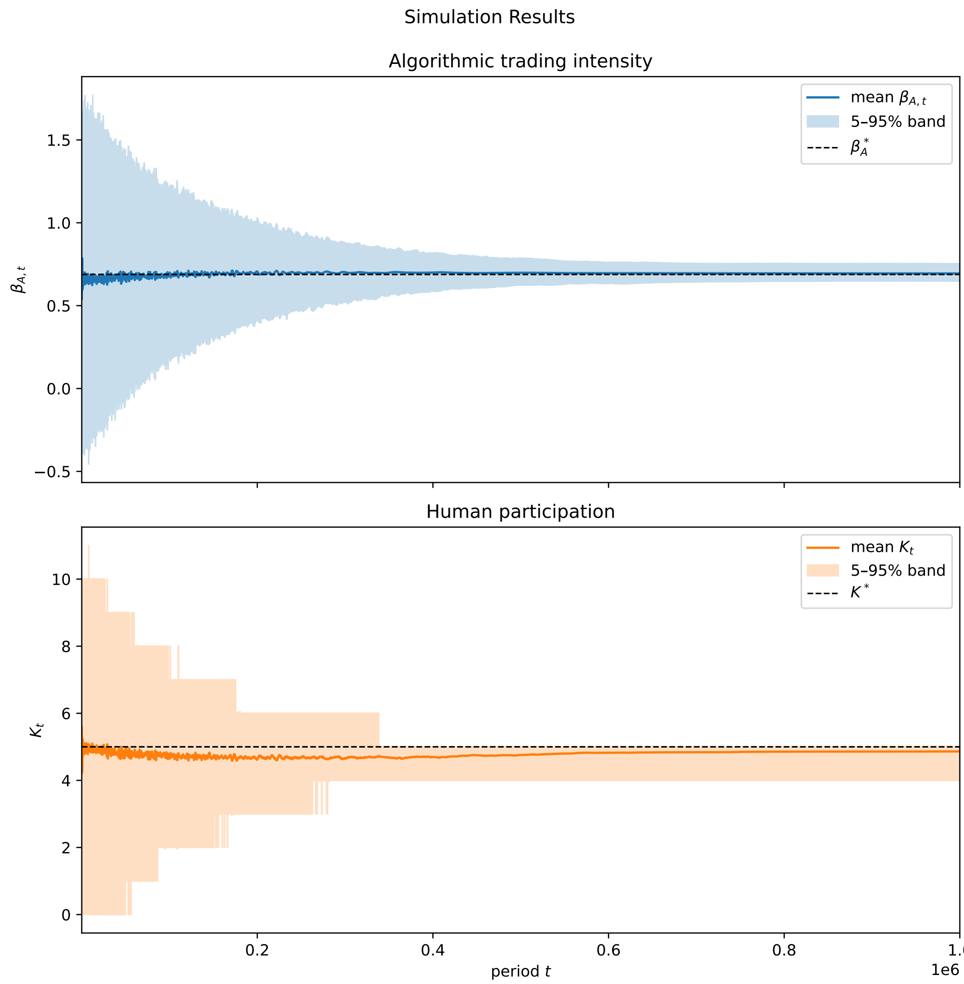

Figure 10: Simulation results for model-based algorithm

The solid lines correspond to the means over $\beta_{M,t}$, and K_t, averaged across runs in each period. The dashed lines represent β_M^* and K^* in Proposition 3. The bands represent the 95% confidence bands around the means. We pick $\rho = 10^{-5}$.