Optimal Execution and Belief Influence: A Mean Field Game Approach

Rodrigo Ribeiro, Yuri Saporito

March 23, 2025

Abstract

Beliefs play a central role in trading, influencing decision-making and driving market dynamics. This centrality gives rise to conflicts of interest, particularly in contexts such as analysts' recommendations and manipulative practices like spoofing. In this paper, we propose a game-theoretical framework to investigate these phenomena, where a major agent strategically influences the beliefs of a continuum of minor agents while engaging in an optimal execution problem. The major agent faces a cost associated with this belief manipulation, leading to a trade-off between influence and trading performance. We embed this interaction into a mean field game framework, which allows us to characterize the resulting Nash equilibrium through a coupled system of Hamilton-Jacobi-Bellman equations that are reduced to a tractable system of ordinary differential equations. Our analysis is complemented by numerical simulations, which illustrate how strategic belief manipulation affects trading outcomes and overall market dynamics. The insights derived from our model shed light on the interplay between information asymmetry and market behavior, with important implications for both market participants and regulatory policies.

1 Introduction

Financial markets are, at their core, belief-aggregation machines. In this context, the role of information is crucial, as it drives the decisions of investors and influences market dynamics. Traditional economic theories suggest that profitable trading strategies emerge from the ability to incorporate new and valuable information into market prices, whether through news, macroeconomic indicators, or proprietary analyses. However, the recognition that beliefs play a central role in financial markets has led market participants to develop strategies that exploit belief dynamics, rather than relying solely on superior information processing. This raises critical questions about how beliefs evolve in response to strategic actions, and how such dynamics impact markets.

Historically, documented strategies of exploiting belief manipulation date back at least to the seventeenth century, as seen in the Amsterdam Stock Exchange, where brokers engaged in *bear raids*: short-selling stocks to drive down prices, then repurchasing them at a discount. In Allen and Gale [1992], the history of stock price manipulation is discussed, and it is shown, under a discrete time trade model, that such stock manipulation can be profitable. A more recent anecdote is studied in Huberman and Regev [2001], which presents the case of EntreMed, whose stock price surged by over 200% following a New York Times article highlighting its cancer research, despite the fact that the underlying scientific breakthrough had been public knowledge for over five months. This price movement, triggered by the mere re-publication of old information, underscores the role of herding behavior and belief contagion in financial markets, where investor attention and media amplification can create significant price distortions independent of new fundamental developments.

In electronic markets, belief influence has grown increasingly sophisticated. A prominent example is *spoofing*, where traders place and cancel orders to create illusory liquidity, distorting others' perception on supply or demand. Spoofing is illegal, and there is a fine associated with it, which makes employing this strategy risky. For a better discussion on spoofing via stochastic control analysis, see Cartea et al. [2020], which models an agent spoofing the limit order book while performing liquidation. Unlike spoofing, which operates through observable market actions, subtler influences can occur through institutional trust. Firms operating on both the sell side and buy side face inherent conflicts of interest, leading to regulatory efforts to enforce *Ethical Walls* that separate these activities. Several scandals where large institutions got fined for breaching this information barrier have raised attention to this topic. Michaely and Womack [1999] studies stock recommendations and discusses these conflicts of interests, arguing that their market-wide effects remain underappreciated by regulators. These considerations suggest that beliefs are not only a byproduct of information, but strategic processes which can be optimized by agents.

We develop a Mean Field Game framework that considers the interaction of different market participants where belief has an endogenous component. This is done by considering the interaction between a major agent and a population of minor agents, both engaged in optimal execution. Inspired by real-world phenomena such as *spoofing* and biased stock recommendations, we introduce belief influence control as an explicit decision variable for the major agent, allowing us to investigate how the major agent balances the cost of influence against potential trading advantages. We characterize the optimal behavior for each agent in the resulting Nash Equilibrium.

The main interest of this modeling approach is to ask what are the effects of having the wrong model for asset dynamics when performing optimal execution with other agents, and what is the optimal strategy when one agent can influence the belief of others. The answer to the first question is inherently dependent the functional agents are maximizing. As argued in Section 3, we can choose a clever gain functional under which answering this question becomes explicitly clear. The second question, however, is more subtle, as it depends on how the influence in belief occurs. We propose a baseline model for this cost dynamics that preserves the linear-quadratic structure of the optimal control problem.

This approach is influenced by some ideas in the mathematical finance literature. We provide a brief summary of advancements in this field: Cartea and Jaimungal [2016] study how order flow information can be incorporated into optimal execution, Cardaliaguet and Lehalle [2018] propose a MFG of controls to study optimal execution and crowd trading; Huang et al. [2019] and Bergault and Sánchez-Betancourt [2024] study a mean field game where different populations perform optimal execution under different models. Cartea et al. [2017] analyze this optimal execution problem under an ambiguity aversion framework. Building on this body of knowledge, our work extends literature by introducing belief control as an explicit decision variable under a MFG framework.

2 Model Setup

We consider the interaction of a major agent with a continuum of minor agents performing optimal execution of an asset. Each minor agent has an inventory process, $(Q_t)_{t\in[0,T]}$, and a wealth process, $(X_t)_{t\in[0,T]}$, evolving according to:

$$dQ_t = \nu_t dt,$$

and

$$dX_t = -\nu_t(S_t + b\nu_t)dt,$$

where $(S_t)_{t\in[0,T]}$ is the asset price process to be defined later, $b > 0$ is the instantaneous price impact coefficient of the minor agents, and $(\nu_t)_{t\in[0,T]}$ is the trading speed controlled individually by each minor agent. Similarly for the major agent, we have processes $(Q_t^M)_{t\in[0,T]}$ and $(X_t^M)_{t\in[0,T]}$ such that:

$$dQ_t^M = \nu_t^M dt,$$

and

$$dX_t^M = -\nu_t^M(S_t + b_M\nu_T^M)dt,$$

where $b_M > 0$ is the instantaneous price impact coefficient for the major agent and $(\nu_t^M)_{t\in[0,T]}$ is the trading speed process for the major agent.
We assume that the stock price follows the dynamics

$$dS_t = \mu_t dt + \sigma dW_t,$$

where $(\mu_t)_{t\geq 0}$ is the drift process, and $(W_t)_{t\geq 0}$ is a standard Brownian motion under the real-world probability measure $\mathbb{P}$. If $(\mathcal{F}_t)_{t\geq 0}$ is the natural filtration to which the Brownian motion is adapted and given a process $(\alpha_t)_{t\geq 0}$, we define an alternative probability measure $\mathbb{P}^\alpha$, absolutely continuous with respect to $\mathbb{P}$, such that:

$$\left.\frac{d\mathbb{P}^\alpha}{d\mathbb{P}}\right|_{\mathcal{F}_t} = \exp\left\{\int_0^t \frac{\alpha_t}{\sigma} dW_t - \frac{1}{2}\int_0^t \left(\frac{\alpha_t}{\sigma}\right)^2 dt\right\}.$$

By Girsanov's theorem,

$$dS_t = (\mu_t + \alpha_t)dt + \sigma d\widetilde{W}_t,$$

where $(\widetilde{W}_t)_{t\geq 0}$ is a Brownian motion under $\mathbb{P}^\alpha$.

2.1 Minor Agents' Belief and Herding

We assume the major agent controls the level of influence such that the minor agents' belief stochastically reverts towards this level. More formally, we write:

$$d\alpha_t = \kappa(c_t - \alpha_t)dt + \sigma_b dW_t^b,$$

$$\alpha_0 = 0,$$

where c_t is the major agent's control, and $(W_t^b)_{t\geq 0}$ is a Brownian motion independent of the ones previously defined. The choice of the control $(c_t)_{t\geq 0}$ incurs in a cost for the major agent, as we will define also in Section 3.
Here, κ is a parameter that encompasses how influential the major agent is, and one of the challenges of this modeling approach is to estimate reasonable values for it.

3 Objective Functionals

For the minor agents, we use the same functional as Cartea and Jaimungal [2016], which is shown to incorporate robustness with regards to model uncertainty in Cartea et al. [2017]. This robustness allows us to explicitly state how a bias in belief affects agents' optimization.
We define for the minor agent:

$$J(t,x,s,q^M,q,\alpha,\nu) = \mathbb{E}_t^{\mathbb{P}^\alpha}\left[X_T + Q_T(S_T - \gamma Q_T) - \psi \int_t^T (Q_u)^2 du\right],$$

where this expectation is conditional on $X_t = x, S_t = s, Q_t^M = q^M, Q_t = q, \alpha_t = \alpha$, $\gamma \geq 0$ is a penalty coefficient for the remaining inventory and ψ is a urgency parameter that incentivizes quick liquidation. Using Girsanov's Theorem and integration by parts, one can show that:

$$\mathbb{E}_t^{\mathbb{P}^\alpha}\left[X_T + Q_T(S_T - \gamma Q_T) - \int_t^T \psi(Q_u)^2 du\right] = \mathbb{E}_t^{\mathbb{P}}\left[X_T + Q_T(S_T - \gamma Q_T) - \int_t^T \left(\psi - \frac{\alpha_u}{Q_u}\right)(Q_u)^2\, du\right],$$

which answers one of our questions: performing optimal execution while having the wrong model is equivalent to performing optimal execution with the correct model, but with a different (and stochastic) urgency process instead of the original urgency parameter. This equation by itself is already insightful, as it provides intuition to what values of α_t the major agent aims to induce given the value of the minor agents inventory Q_t.
The objective functional for the major agent is very similar, but incorporates a quadratic cost for c_t, resulting in

$$J^M(t,x,s,q^M,\alpha,\nu^M) = \mathbb{E}_t^{\mathbb{P}}\left[X_T^M + Q_T^M(S_T - \gamma^M Q_T^M) - \psi^M \int_t^T \left(Q_u^M\right)^2 du - \zeta \int_t^T (c_u)^2\, du\right],$$

where $\gamma^M \geq 0$ and $\zeta > 0$ is the cost parameter for the influence of the major agent.

4 Mean Field Setting

All minor agents are identical in their preferences, differing only in their inventory holdings. At each time t, the inventory holdings of minor agents are distributed according to a probability measure $M(t,\cdot)$ defined on the Borel σ-algebra $\mathcal{B}(\mathbb{R})$. For any Borel set $A \subseteq \mathbb{R}$, $M(t,A)$ represents the proportion of minor agents whose inventory holdings lie within A at time t. We assume the measure M has a density with respect to the lebesgue measure, m, such that when minor agents are performing their optimal execution, the following transport equation holds:

$$\partial_t m + \partial_q\left(m\, \nu^*(t,q)\right) = 0,$$

where we abuse the notation here to denote by $\nu^*(t,q)$ the optimal trading speed of a minor agent as a function of q, implicitly conditioned on the values of q^M and α.
We then define the mean field quantity $\bar{\nu}_t$, the average trading speed of the minor agents as

$$\bar{\nu}(t) = \int_{\mathbb{R}} \nu^*(t,q) m(t,q) dq,$$

from which we define the market clearing condition:

$$\mu_t = \lambda_1 \nu_t^M + \lambda_2 \bar{\nu}(t),$$

imposing a linear[1] permanent price impact with different coefficients for the different agents. We also define the average holding for the minor agent population at time t,

$$E(t) = \int_{\mathbb{R}} q m(t,q) dq,$$

and, using integration by parts,

$$E'(t) = \int_{\mathbb{R}} \nu^*(t,q) m(t,q) dq = \bar{\nu}(t).$$

5 Value Functions and HJB equations

Each minor agent computes its optimal behavior given a aggregate trading speed $\bar{\nu}$, and the major agent trading speed ν^M. They have the value function

$$\tilde{V}(t,x,s,q^M,q,\alpha) = \sup_{\nu} J(t,x,s,q^M,q,\alpha,\nu;\bar{\nu}),$$

which leads to the HJB equation:

$$\frac{\partial \tilde{V}}{\partial t} + \sup_{v} \Bigg[-v(s+bv)\frac{\partial \tilde{V}}{\partial x} + (\lambda_1 \nu_u^M + \lambda_2 \bar{\nu}_u)\frac{\partial \tilde{V}}{\partial s} + \frac{1}{2}\sigma^2 \frac{\partial^2 \tilde{V}}{\partial s^2} + \mathcal{L}^\alpha \tilde{V} + v\frac{\partial \tilde{V}}{\partial q} + \nu_u^M \frac{\partial \tilde{V}}{\partial q^M} - \psi q^2 \Bigg] = 0.$$

Again using integration by parts, we can write:

$$\tilde{V}(t,x,s,q^M,q,\alpha;\bar{\nu},\nu^M) = x + qS + V(t,q^M,q,\alpha;\bar{\nu},\nu^M).$$

Using this substitution, we get the following HJB equation:

$$\frac{\partial V}{\partial t} + \sup_{v} \left[v\frac{\partial V}{\partial q} + \nu^M \frac{\partial V}{\partial q^M} + \mathcal{L}^\alpha V - bv^2 + q(\lambda_1 \nu_u^M + \lambda_2 \bar{\nu}_u) - \psi q^2 \right] = 0,$$

which allows us state the optimal control as:

$$\nu^*(t,q) = \frac{\partial_q V}{2b},$$

and

$$\partial_t V + \frac{(\partial_q V)^2}{2b} + \nu^M \partial_{q^M} V + \mathcal{L}^\alpha V - \frac{(\partial_q V)^2}{4b} + q(\lambda_1 \nu_u^M + \lambda_2 \bar{\nu}) - \psi q^2 = 0,$$

with terminal condition

$$V(T,q^M,q,\alpha;\bar{\nu},\nu^M) = -\gamma q^2.$$

[1] We are not incorporating transient price impact in this model. Doing it would make the equations more complex, but the framework we develop here can still be applied.

The major player faces a similar problem, with the HJB equation:

$$\partial_t V^M + \sup_{v,c}\left[\frac{\partial V^M}{\partial q^M}v + \mathcal{L}^\alpha V^M + q^M(\lambda_1 v + \lambda_2\bar{\nu}) - b_M v^2 - \psi^M (q^M)^2 - \zeta a^2\right] = 0,$$

which implies optimal controls

$$\nu^{*,M} = \frac{\partial_{q^M} V^M + q^M\lambda_1}{2b_M} \quad \text{and} \quad c^* = \frac{\kappa\partial_\alpha V^M}{2\zeta},$$

where

$$\partial_t V^M + \partial_{q^M} V^M\left(\frac{\partial_{q^M} V^M + q^M\lambda_1}{2b_M}\right) + \mathcal{L}^\alpha V^M + q^M\left(\lambda_1\frac{\partial_{q^M} V^M + q^M\lambda_1}{2b_M} + \lambda_2\bar{\nu}\right) - b_M\left(\frac{\partial_{q^M} V^M + q^M\lambda_1}{2b_M}\right)^2$$
$$-\psi^M(q^M)^2 - \zeta\left(\frac{\kappa\partial_\alpha V^M}{2\zeta}\right)^2 = 0,$$

with terminal condition

$$V(T, q^M, \alpha; \bar{\nu}) = -\gamma^M (q^M)^2.$$

6 Mean Field Equilibrium

When acting optimally, each agent adjusts their speed of trading as the best response to the actions of all others. Equilibrium is achieved when this best response function converges to a fixed point for all agents. This fixed point is the solution to the following system of coupled PDEs:

$$\begin{cases} \partial_t V + \frac{(\partial_q V)^2}{4b} + \left(\frac{\partial_{q^M} V^M + q^M\lambda_1}{2b_M}\right)\partial_{q^M} V + \mathcal{L}^\alpha V + q\left(\lambda_1\frac{\partial_{q^M} V^M + q^M\lambda_1}{2b_M} + \lambda_2 E'(t)\right) - \psi q^2 = 0, \\ E'(t) = \int_Q \frac{\partial_q V}{2b} m(t,q)dq, \\ \partial_t V^M + \left(\frac{\partial_{q^M} V^M + q^M\lambda_1}{2b_M}\right)\partial_{q^M} V^M + \mathcal{L}^\alpha V^M + q^M\left(\lambda_1\frac{\partial_{q^M} V^M + q^M\lambda_1}{2b_M} + \lambda_2 E'(t)\right) - b_M\left(\frac{\partial_{q^M} V^M + q^M\lambda_1}{2b_M}\right)^2 \\ -\psi^M(q^M)^2 - \zeta(\frac{\kappa\partial_\alpha V^M}{2\zeta})^2 = 0, \end{cases}$$

with appropriate boundary conditions.

Due to the linear quadratic structure of this model, we expect the solution to be a quadratic polynomial in q (for V), q^M and α. We then write the following substitutions:

$$V = \phi_1 q^2 + \phi_2 q\, q^M + \phi_3 (q^M)^2 + \phi_4 q + \phi_5 q^M + \phi_6\alpha + \phi_7\alpha^2 + \phi_8\alpha q + \phi_9\alpha q^M + \phi_{10},$$
$$V^M = \tilde{\phi}_1 (q^M)^2 + \tilde{\phi}_2 q^M\alpha + \tilde{\phi}_3(\alpha)^2 + \tilde{\phi}_4 q^M + \tilde{\phi}_5\alpha + \tilde{\phi}_6.,$$

and

$$E(t) = \bar{\phi}_1 q^M + \bar{\phi}_2\alpha + \bar{\phi}_3.$$

Plugging these formulas into the system of coupled PDEs and collecting terms knowing that the resulting polynomial is identically null polynomial results in a system of 19 ODEs, see A. These

ODEs can be solved analytically, using integrating factor and Ricatti ODE theory. Due to the high dimensionality of this system, in order to solve it analytically, it would be appropriate to use a symbolic calculator software. While symbolic computation could, in principle, yield analytical solutions, we anticipate that the key dynamical insights will emerge most transparently from numerical simulations. Moreover, the structure of the ODEs, sufficiently tractable for modern numerical solvers, ensures robustness in their integration, allowing us to focus directly on presenting and interpreting the results that underpin our claims.

7 FBSDE Approach

One of the theoretical questions that emerge from the approach used up to this point is whether one can guarantee that a mean field equilibrium does exist. One of the ways to answer this question is to pose the problem in the FBSDE framework, where we can rely on known results for linear FBSDEs that guarantee existence and uniqueness of the solution. We begin by rewriting the gain functional for the minor agent. We have, slightly abusing the notation for J,

$$J(\nu) = \mathbb{E}_t^{\mathbb{P}^\alpha}\left[X_T + Q_T(S_T - \gamma Q_T) - \psi \int_0^T (Q_u)^2 du\right].$$

We can rewrite this definition using integration by parts to get

$$\begin{aligned} J(\nu) &= x + q_t S_t + q_t^2\, \mathbb{E}_t^{\mathbb{P}^\alpha}\left[\int_0^T q_u dS_u + \int_0^T -b(\nu_u)^2 - 2\gamma q_u \nu_u - \psi (q_u)^2 du\right] \\ &= C_t + \mathbb{E}_t^{\mathbb{P}^\alpha}\left[\int_0^T q_u dS_u - \int_0^T \begin{pmatrix}\nu_u \\ q_u\end{pmatrix}^T \begin{pmatrix} b & \gamma \\ \gamma & \psi \end{pmatrix} \begin{pmatrix}\nu_u \\ q_u\end{pmatrix} du\right], \end{aligned}$$

where $C_t = x_t + q_t S_t + q_t^2$. For the Major agent, we similarly have:

$$J^M((\nu^M, c)) = \mathbb{E}_t^{\mathbb{P}}\left[X_T^M + Q_T^M(S_T - \gamma^M Q_T^M) - \psi^M \int_0^T \left(Q_u^M\right)^2 du - \zeta \int_0^T (c_u)^2\, du\right].$$

such that

$$J^M((\nu^M, c)) = Z_t + \mathbb{E}_t^{\mathbb{P}}\left[\int_0^T q_u^M dSu - \int_0^T \begin{pmatrix}\nu_u^M \\ q_u^M \\ c_u\end{pmatrix}^T \begin{pmatrix} b_M & \gamma^M & 0 \\ \gamma_M & \psi^M & 0 \\ 0 & 0 & \zeta \end{pmatrix} \begin{pmatrix}\nu_u^M \\ q_u^M \\ c_u\end{pmatrix} du\right],$$

where Z_t is defined analogously.

We now define the functionals $L(\nu)$ and $L^M((\nu^M, c))$ as

$$\begin{aligned} L(\nu) &= \mathbb{E}_t^{\mathbb{P}^\alpha}\left[\int_0^T q_u dS_u - \int_0^T \begin{pmatrix}\nu_u \\ q_u\end{pmatrix}^T \begin{pmatrix} b & \gamma \\ \gamma & \psi \end{pmatrix} \begin{pmatrix}\nu_u \\ q_u\end{pmatrix} du\right], \\ L^M((\nu^M, c)) &= \mathbb{E}_t^{\mathbb{P}}\left[\int_0^T q_u^M dSu - \int_0^T \begin{pmatrix}\nu_u^M \\ q_u^M \\ c_u\end{pmatrix}^T \begin{pmatrix} b_M & \gamma^M & 0 \\ \gamma^M & \psi^M & 0 \\ 0 & 0 & \zeta \end{pmatrix} \begin{pmatrix}\nu_u^M \\ q_u^M \\ c_u\end{pmatrix} du\right], \end{aligned}$$

such that the maximizers of Js are the same as the ones for Ls, as both C_t and Z_t do not depend on the controls. We will now show that both L and L^M are strictly concave, which allows us to determine optimal controls straightforwardly.

Lemma 1. *For any $\rho \in (0,1)$, and optimal controls ν_1, ν_2, c_1, c_2 and $v_1 = (\nu_1, c_1), v_2 = (\nu_2, c_2)$, the following inequalities hold:*

$$L(\rho\nu_1 + (1-\rho)\nu_2) > \rho L(\nu_1) + (1-\rho)L(\nu_2),$$
$$\textit{and}$$
$$L^M(\rho v_1 + (1-\rho)v_2) > \rho L^M(v_1) + (1-\rho)L^M(v_2).$$

Proof. See Appendix B.1 □

We define the Gateaux derivative of $L(\nu)$ by $\mathcal{D}L(\nu)(\cdot)$, as an operator such that

$$\mathcal{D}L(\nu)(\omega) = \lim_{\epsilon \to 0} \frac{L(\nu + \epsilon\omega) - L(\nu)}{\epsilon},$$

and the analogous definition holds for the Gateaux derivative of $L^M(v)$.

Lemma 2. *The following formulas hold for the Gateaux derivatives of $L(v)$ and $L^M(v)$:*

$$\mathcal{D}L(v)(\omega) = \mathbb{E}^{\mathbb{P}^\alpha}\left[\int_0^T \omega_s \left(-2b\nu_s - \int_0^s \mathbb{E}^{\mathbb{P}^\alpha}\left[\lambda_1 \nu_u^M + \lambda_2 \bar{\nu}(u) \mid \mathcal{F}_u\right] - 2\psi q_u^\nu \, du + M_s\right) ds\right],$$

$$\begin{aligned}\mathcal{D}L^M(v^M, c)(\omega, \chi) = \mathbb{E}^{\mathbb{P}}&\left[\int_0^T \omega_s \left(-2b_M \nu_s^M - \int_0^s \mathbb{E}^{\mathbb{P}^\alpha}\left[\lambda_1 \nu_u^M + \lambda_2 \bar{\nu}(u) \mid \mathcal{F}_u\right] - 2\psi^M q_u^{\nu^M} \, du + N_s\right) ds\right] \\ &+ \mathbb{E}^{\mathbb{P}}\left[\int_0^T 2\zeta \chi_s c_s ds\right].\end{aligned}$$

Where M_s and N_s are $\mathbb{P}^\alpha$ and $\mathbb{P}$ martingales respectively, defined by:

$$M_s = \mathbb{E}^{\mathbb{P}^\alpha}\left[-2\gamma q_T^\nu + \int_0^T \lambda_1 \nu_u^M + \lambda_2 \bar{\nu}(u) - 2\psi q_u^\nu du \mid \mathcal{F}_s\right],$$

and

$$N_s = \mathbb{E}^{\mathbb{P}}\left[-2\gamma^M q_T^{\nu^M} + \int_0^T \lambda_1 \nu_u^M + \lambda_2 \bar{\nu}(u) - 2\psi^M q_u^{\nu^M} du \mid \mathcal{F}_s\right].$$

Proof. See Appendix B.2 □

8 Simulations

To gain insights into the model's mechanics, we perform simulations using reasonable parameter values. Many of these values are informed by related literature, such as Cartea and Jaimungal [2016], which provides empirical estimates for several key parameters, such as price impact coefficients. However, the parameters κ and ζ lack established empirical validation, presenting a challenge for accurate calibration. To address this, we explore a range of plausible values for these parameters and analyze their impact on the interactions studied. We present here the main results we find under reasonable value parameters.
In Figure 1 we find a summary of the results of simulations under a set of specific parameters. In order to understand the histogram of terminal wealth, we define the benchmark to this model being the case where $\kappa = 0$, as this implies no belief control by the major agent. By comparing the

terminal wealth with the benchmark, we get an idea of the excess wealth achieved by controlling belief. In order to define the excess terminal wealth, we first introduce

$$C_T = X_T + S_T Q_T.$$

The excess terminal wealth, Ξ, is then defined as

$$\Xi = \frac{C_T^{\text{model}} - C_T^{\text{benchmark}}}{C_T^{\text{benchmark}}}.$$

In Figure 2[2], we can see plots about inventory dynamics under this model, as well as comparisons between trading under belief control, and the benchmark model ($\kappa = 0$).

We also look at the effect of the parameters ζ and κ on the dynamics. For that, we calculate the excess terminal wealth varying the values of these parameters and plot the results in Figure 3. In Figure 3, we see histograms for the major agent's excess terminal wealth for different parameter values. The plot in 4 does the same for the minor agents. When comparing Figures 4 and 3,

[2]One interesting conjecture that emerges from these plots is that the relative difference in inventories curves and the $y = 0$ line are concurrent across all simulations. This might be a trivial fact, but it remains to further examine it.

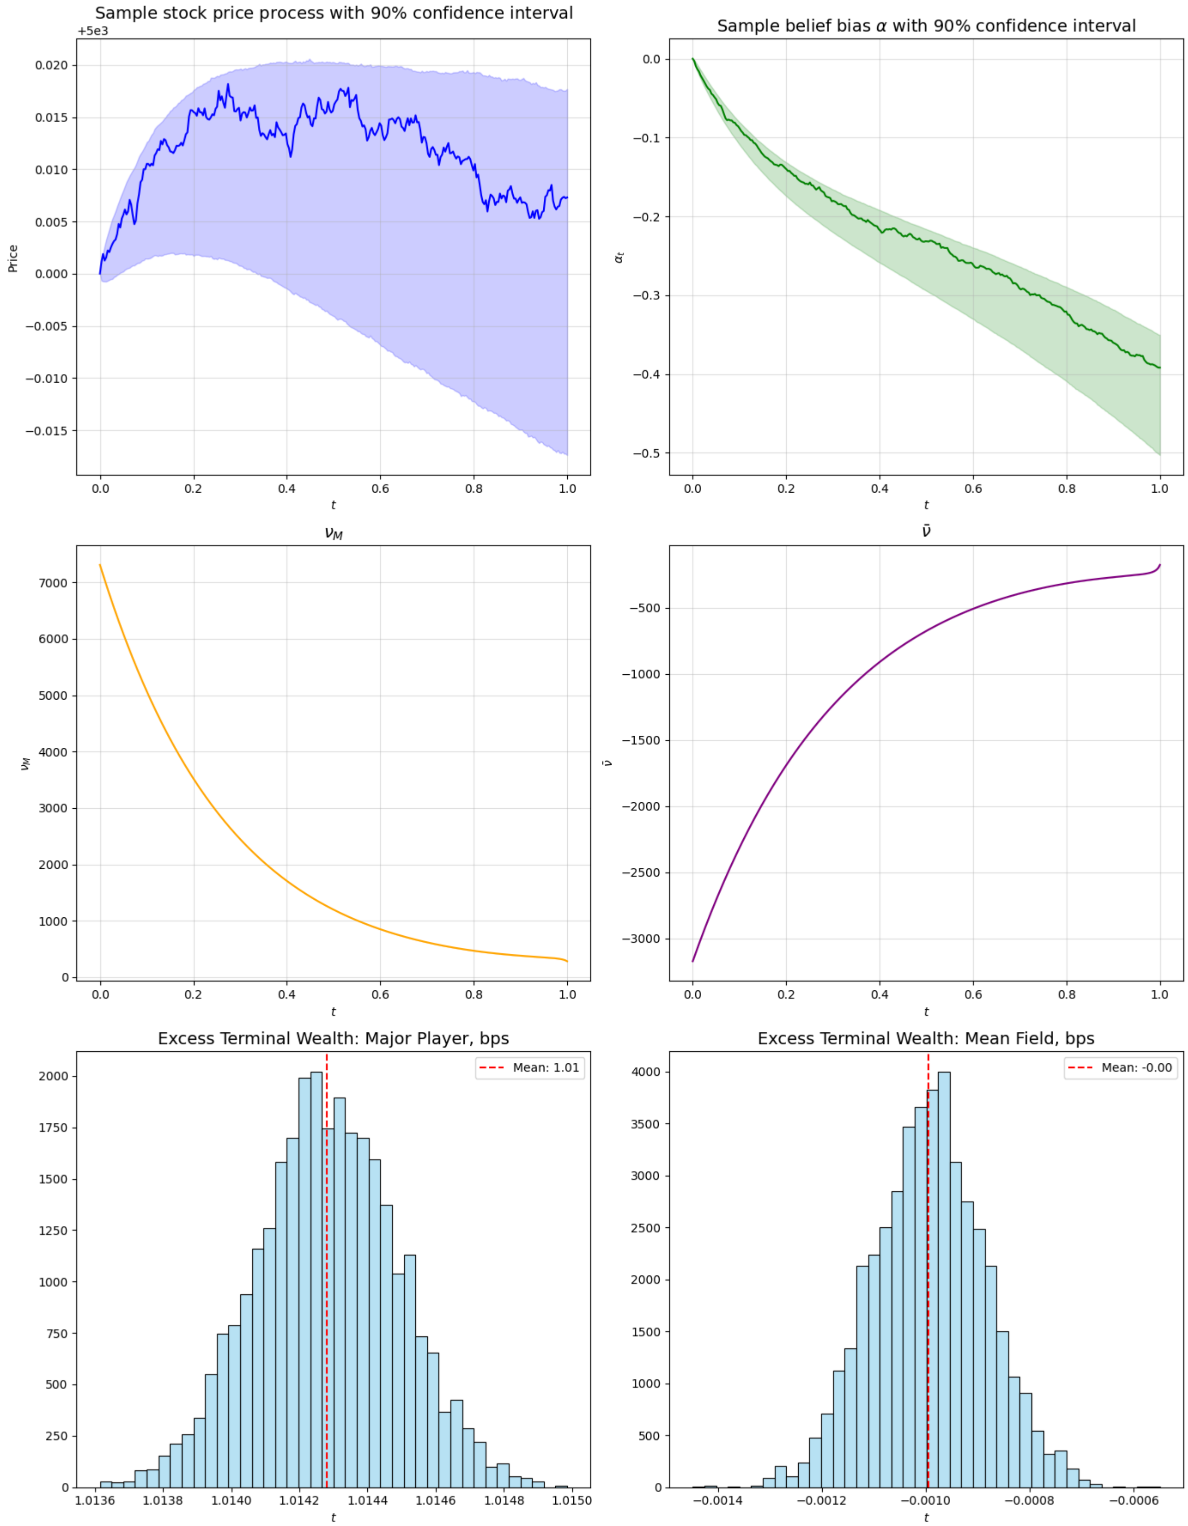

Figure 1: Overview of processes extracted from simulations.

the variation in wealth for the minor agents is smaller compared to the wealth variation for the major agent. One possible explanation is the fact that the ζ parameter is directly impacting the final wealth for the major agent, but not for the minor agent. However, the fact that this same

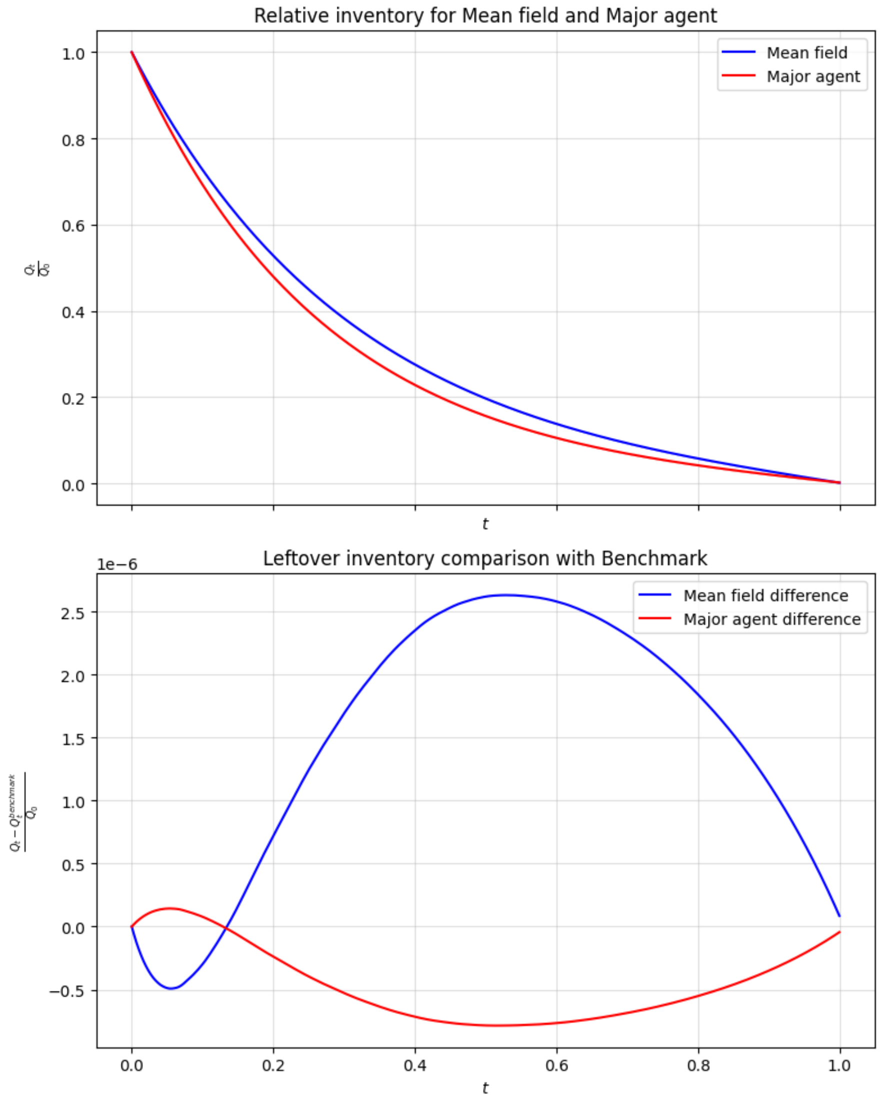

Figure 2: Inventory dynamics and comparison with benchmark.

observation also holds true for varying the κ suggests that this observation is more fundamental about this system. It is still necessary to understand if this is governed by regimes that change under different parametric regions, and this can be better understood in the future. From what we can observe from this data, interesting questions about the nature of this game are raised. It suggests that the wealth under different beliefs does not form a 0-sum game, as the wealth gains for the major agent does not generate a wealth loss of the same magnitude for the minor agents[3]. Another interesting visualization here is the 3d plot of the mean excess wealth varying both ζ and κ. From the tests we performed, the deviation from this mean is not significant compared to the magnitude of the mean itself. The plot on Figure 5 summarizes this surface in a subset of the parameter values. We have analogous plot for the mean field of minor agents in Figure 6.

[3] In order to understand the aggregate effects on wealth for the collective of all players, it is important to understand what is the relative wealth size of the major agent in relation to the population of minor agents. From the model for μ_t, we expect that the *wealth-size* of the minor agents is $\frac{\lambda_2}{\lambda_1+\lambda_2}$, and, under this hypothesis, the aggregate wealth can increase due to belief control.

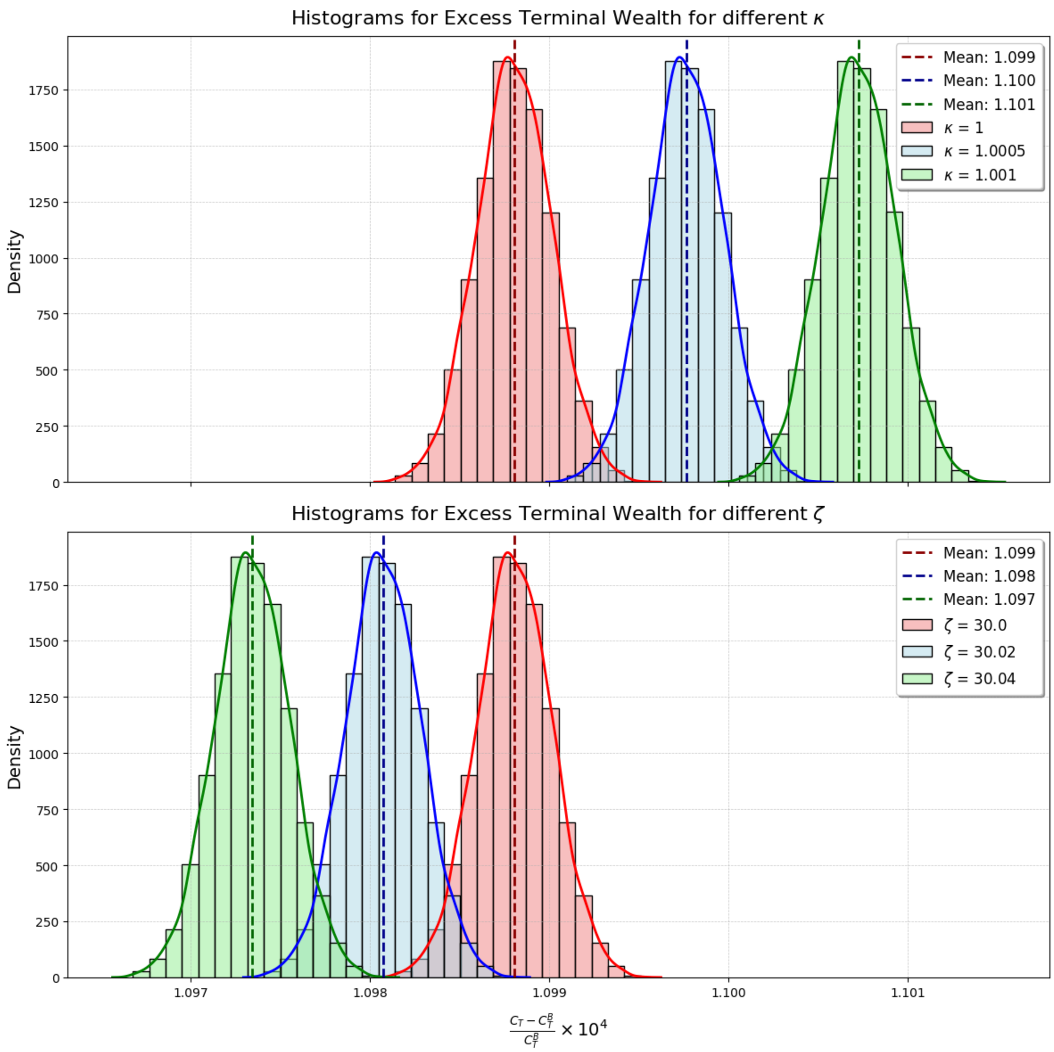

Figure 3: Excess terminal wealth for the major agent given different parameter choices

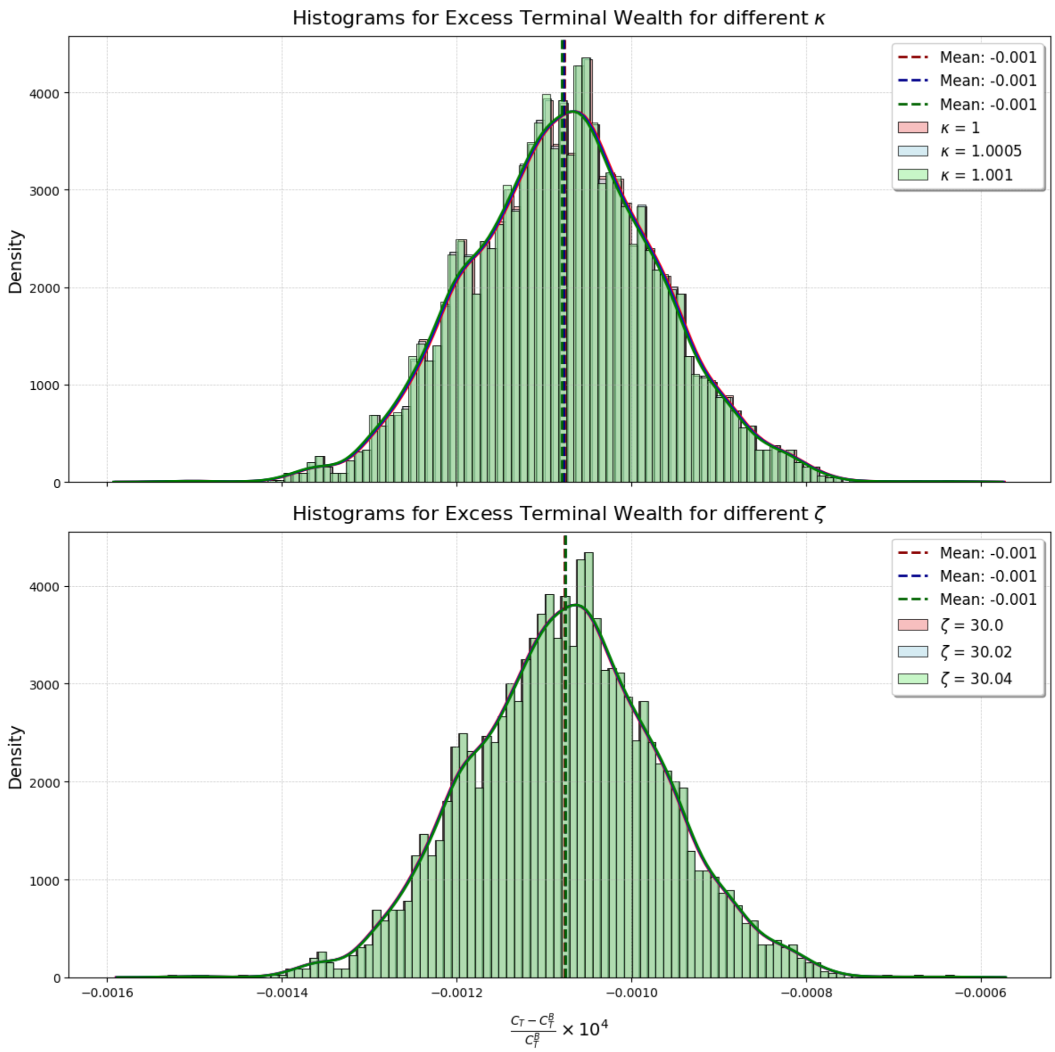

Figure 4: Excess terminal wealth for the mean field of minor agents given different parameter choices

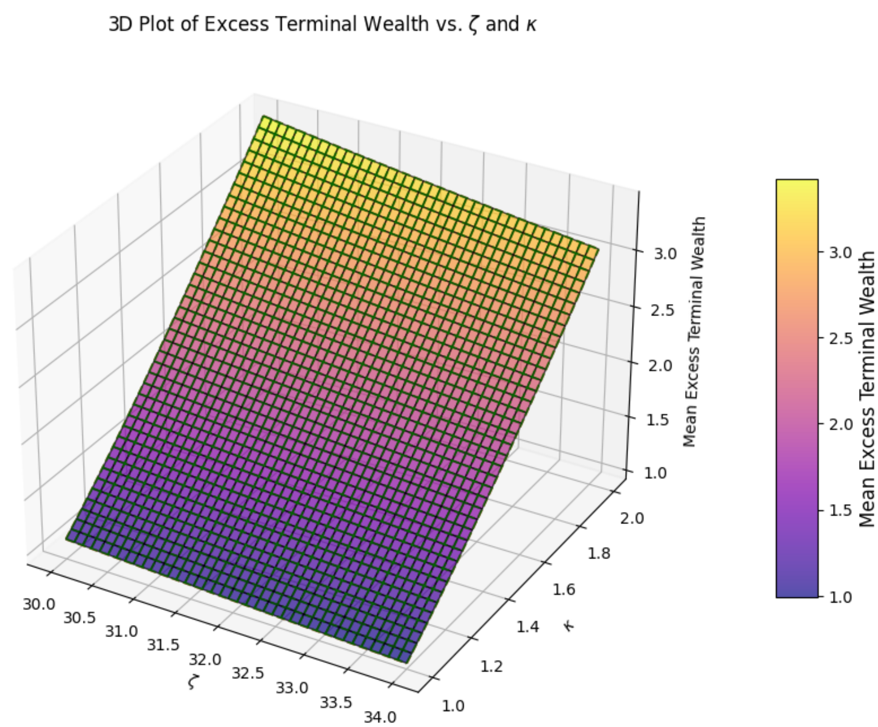

Figure 5: 3D plot of the mean excess wealth surface for the major agent

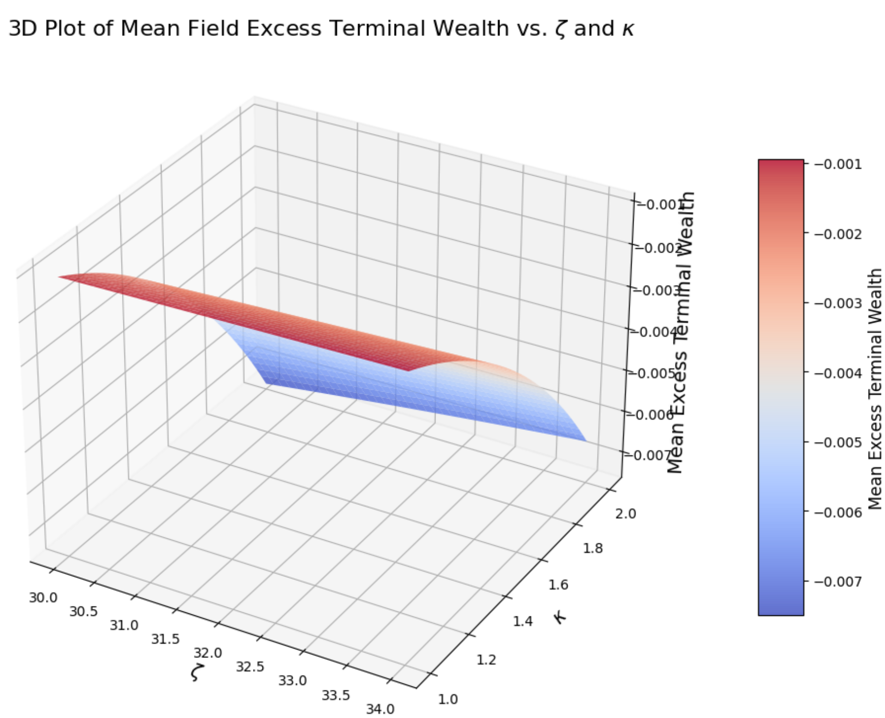

Figure 6: 3D plot of the mean excess wealth surface for the mean field of minor agents

9 Future Works

Potential future works include the extension of this approach to multiple populations of both major and minor agents, understanding the cost structure of belief manipulation in measurable strategies such as spoofing, or modeling beliefs with a path-dependent component, which could enrich the mathematical analysis, while improving the modeling of behavioral phenomena such as trend seeking.

Another interesting direction of study would be to investigate the aggregate effects of this belief influence. From the simulations, there are cases where both the major agent and the minor agents are better off under belief control when compared to the benchmark. This phenomenon raises questions about the wealth creation that seems to be possible and, for instance, one could ask how does this type of wealth creation relates to bubble formation.

A System of Ricatti ODEs

We have the following system of ODEs obtained by collecting monomials of the quadratic polynomials on q, q^M and α:

$$
\begin{aligned}
\frac{d\phi_1}{dt} &= \frac{\phi_1^2}{b} - \psi_m, \\
\frac{d\phi_2}{dt} &= \frac{\phi_1\phi_2}{b} + \frac{2\tilde{\phi}_1\phi_2 + \lambda_1\phi_2 + 2\lambda_1\tilde{\phi}_1 + \lambda_1^2}{2b_M} + \lambda_2\frac{2\phi_1\bar{\phi}_1 + \phi_2}{2b} + \frac{\kappa^2\tilde{\phi}_2\phi_8}{2\zeta}, \\
\frac{d\phi_3}{dt} &= \frac{\phi_2^2}{4b} + \frac{4\tilde{\phi}_1\phi_3 + 2\lambda_1\phi_3}{2b_M} + \frac{\kappa^2\phi_9\tilde{\phi}_2}{2\zeta}, \\
\frac{d\phi_4}{dt} &= \frac{\phi_1\phi_4}{b} + \frac{\tilde{\phi}_4(\phi_2 + \lambda_1)}{2b_M} + \frac{\kappa^2\phi_8\tilde{\phi}_5}{2\zeta} + \lambda_2\frac{2\phi_1\bar{\phi}_3 + \phi_4}{2b}, \\
\frac{d\phi_5}{dt} &= \frac{\phi_2\phi_4}{2b} + \frac{2\tilde{\phi}_1\phi_5 + 2\phi_3\tilde{\phi}_4}{2b_M} + \frac{\kappa^2(\phi_9\tilde{\phi}_5 + \tilde{\phi}_2\phi_6)}{2\zeta}, \\
\frac{d\phi_6}{dt} &= \frac{\phi_8\phi_4}{2b} + \frac{\tilde{\phi}_2\phi_5 + \tilde{\phi}_4\phi_9}{2b_M} + \frac{\kappa^2(\phi_7\tilde{\phi}_5 + \tilde{\phi}_3\phi_6)}{\zeta} - \kappa\phi_6, \\
\frac{d\phi_7}{dt} &= \frac{\phi_8^2}{4b} + \frac{\tilde{\phi}_2\phi_9}{2b_M} + 2\frac{\kappa^2\phi_7\tilde{\phi}_3}{\zeta} - 2\kappa\phi_7, \\
\frac{d\phi_8}{dt} &= \frac{\phi_1\phi_8}{b} + \frac{\tilde{\phi}_2\phi_2 + \lambda_1\tilde{\phi}_2}{2b_M} + \frac{\kappa^2 2\tilde{\phi}_3\phi_8}{2\zeta} - \kappa\phi_8 + \lambda_2\frac{2\phi_1\bar{\phi}_2 + \phi_8}{2b}, \\
\frac{d\phi_9}{dt} &= \frac{\phi_2\phi_8}{2b} + \frac{2\tilde{\phi}_1\phi_9 + 2\tilde{\phi}_2\phi_3 + \lambda_1\phi_9}{2b_M} + \frac{\kappa^2(\phi_7\tilde{\phi}_2 + \phi_9\tilde{\phi}_3)}{\zeta} - \kappa\phi_9, \\
\frac{d\phi_{10}}{dt} &= \frac{\phi_4^2}{4b} + \frac{\tilde{\phi}_4\phi_5}{2b_M} + \frac{\kappa^2\phi_6\tilde{\phi}_5}{2\zeta} + \sigma_b^2\phi_7, \\
\frac{d\tilde{\phi}_1}{dt} &= \frac{(\lambda_1 + 2\tilde{\phi}_1)^2}{4b_M} + \frac{(\kappa\tilde{\phi}_2)^2}{4\zeta} + \lambda_2\frac{\phi_2 + 2\phi_1\bar{\phi}_1}{2b} - \psi_M, \\
\frac{d\tilde{\phi}_2}{dt} &= \frac{4\tilde{\phi}_1\tilde{\phi}_2 + 2\tilde{\phi}_2\lambda_1}{4b_M} + \frac{\kappa^2\tilde{\phi}_3\tilde{\phi}_2}{\zeta} - \kappa\tilde{\phi}_2 + \lambda_2\frac{\phi_8 + 2\phi_1\bar{\phi}_2}{2b}, \\
\frac{d\tilde{\phi}_3}{dt} &= \frac{\tilde{\phi}_2^2}{4b_M} + \frac{(\kappa\tilde{\phi}_3)^2}{\zeta} - 2\kappa\tilde{\phi}_3, \\
\frac{d\tilde{\phi}_4}{dt} &= \frac{4\tilde{\phi}_1\tilde{\phi}_4 + 2\lambda_1\tilde{\phi}_4}{4b_M} + \frac{\kappa^2\tilde{\phi}_5\tilde{\phi}_2}{2\zeta} + \lambda_2\frac{2\phi_1\bar{\phi}_3 + \phi_4}{2b}, \\
\frac{d\tilde{\phi}_5}{dt} &= \frac{\tilde{\phi}_2\tilde{\phi}_4}{2b_M} + \frac{\kappa^2\tilde{\phi}_3\tilde{\phi}_5}{\zeta} - \kappa\tilde{\phi}_5, \\
\frac{d\tilde{\phi}_6}{dt} &= \frac{\tilde{\phi}_4^2}{4b_M} + \frac{\kappa^2(\tilde{\phi}_5)^2}{4\zeta} + \sigma_b^2\tilde{\phi}_3, \\
\frac{d\bar{\phi}_1}{dt} &= -\frac{2\phi_1\bar{\phi}_1 + \phi_2}{2b}, \\
\frac{d\bar{\phi}_2}{dt} &= -\frac{2\phi_1\bar{\phi}_2 + \phi_8}{2b}, \\
\frac{d\bar{\phi}_3}{dt} &= -\frac{2\phi_1\bar{\phi}_3 + \phi_4}{2b}.
\end{aligned}
$$

This is the system of ODEs we solve numerically.

B Proofs

B.1 Proof of Lemma 1

Proof. As we will see, the proof for these statements relies heavily in the integrated quadratic form structure of the functionals. We will present the proof for the case of the minor agent, and the major agent's case is analogous.
We begin by defining

$$M(\nu) = \mathbb{E}_t^{\mathbb{P}^\alpha}\left[\int_0^T q_u dS_u\right],$$
$$N(\nu) = \mathbb{E}_t^{\mathbb{P}^\alpha}\left[-\int_0^T \begin{pmatrix}\nu_u \\ q_u\end{pmatrix}^T \begin{pmatrix} b & \gamma \\ \gamma & \psi\end{pmatrix}\begin{pmatrix}\nu_u \\ q_u\end{pmatrix} du\right],$$

such that

$$L(\nu) = M(\nu) + N(\nu).$$

We will show that $M(\nu)$ is concave and $N(\nu)$ is strictly concave.
Let $A = \begin{pmatrix} b & \gamma \\ \gamma & \psi\end{pmatrix}$ and $y_u^\nu = \begin{pmatrix}\nu_u \\ q_u\end{pmatrix}$. The inventory dynamics implies that

$$y_u^{\rho\nu_1+(1-\rho)\nu_2} = \rho y_u^{\nu_1} + (1-\rho)y_u^{\nu_2} \quad \forall\, 0 \le u \le T.$$

Notice that, as $b > 0$, $\gamma > 0$ and $\psi > 0$, the matrix A is symmetric with positive entries, and therefore it is positive definite. It is known that for a positive definite matrix A, the function $g(y) = y^T A y$ is strictly convex, therefore $N(\nu)$ is strictly concave. To finish the proof, we can use the same argument we used for the convex combination $y_u^{\rho\nu_1+(1-\rho)\nu_2}$ to see that:

$$M(\rho\nu_1 + (1-\rho)\nu_2) = \rho M(\nu_1) + (1-\rho)M(\nu_2).$$

This implies $M(\nu)$ is concave, and therefore $L(\nu)$ is strictly concave. □

B.2 Proof of Lemma 2

Proof. We begin by noticing that

$$q^{\nu+\epsilon\omega} = (1-\epsilon)q^{\frac{\nu}{(1-\epsilon)}} + \epsilon q^\omega,$$

and $q_u^{\frac{\nu}{1-\epsilon}} = Q_0 - \frac{1}{1-\epsilon}\int_0^u \nu_s ds$, such that:

$$q^{\nu+\epsilon\omega} - q^\nu = -\epsilon Q_0 + \epsilon q^\omega.$$

Before proceeding, we state the following identity. For x and y vectors of appropriate dimensions, and letting $A = \begin{pmatrix} b & \gamma \\ \gamma & \psi\end{pmatrix}$, we have:

$$x^T A x - y^T A y = (x-y)^T A (x-y) + 2(x-y)^T A y.$$

Using this identity, we write:

$$L(\nu+\epsilon\omega)-L(\nu)=\mathbb{E}^{\mathbb{P}^\alpha}\left[\int_0^T q_u^{\nu+\epsilon}-q_u^\nu dS_u-\int_0^T \begin{pmatrix}\nu_u+\epsilon\omega_u\\ q_u^{\nu+\epsilon\omega}\end{pmatrix}^T A\begin{pmatrix}\nu_u+\epsilon\omega_u\\ q_u^{\nu+\epsilon\omega}\end{pmatrix}-\begin{pmatrix}\nu_u\\ q_u^\nu\end{pmatrix}^T A\begin{pmatrix}\nu_u\\ q_u^\nu\end{pmatrix}du\right]$$
$$=\mathbb{E}^{\mathbb{P}^\alpha}\left[\epsilon\int_0^T q_u^\omega-Q_0\,dS_u-\int_0^T\begin{pmatrix}\epsilon\omega_u\\ \epsilon(q_u^\omega-Q_0)\end{pmatrix}^T A\begin{pmatrix}\epsilon\omega_u\\ \epsilon(q_u^\omega-Q_0)\end{pmatrix}+2\begin{pmatrix}\epsilon\omega_u\\ \epsilon(q_u^\omega-Q_0)\end{pmatrix}^T A\begin{pmatrix}\nu_u\\ q_u^\nu\end{pmatrix}du\right]$$

Notice that the first term in the integrand is proportional to ϵ^2 and therefore it vanishes in the limit for the Gateaux derivative. We then write, using the dynamics for S_t:

$$\begin{aligned}
\lim_{\epsilon\to 0}\frac{L(\nu+\epsilon\omega)-L(\nu)}{\epsilon}&=\mathbb{E}^{\mathbb{P}^\alpha}\left[\int_0^T q_u^\omega-Q_0\,dS_u-2\int_0^T\begin{pmatrix}\omega_u\\ q_u^\omega-Q_0\end{pmatrix}^T A\begin{pmatrix}\nu_u\\ q_u^\nu\end{pmatrix}du\right]\\
&=\mathbb{E}^{\mathbb{P}^\alpha}\left[\int_0^T(q_u^\omega-Q_0)(\lambda_1\nu_u^M+\lambda_2\bar\nu(u)-2\gamma\nu_u-2\psi q_u^\nu)du-\int_0^T 2\omega_u(b\nu_u+\gamma q_u^\nu)du\right]\\
&=\mathbb{E}^{\mathbb{P}^\alpha}\left[\int_0^T\int_0^u\omega_s ds\;(\lambda_1\nu_u^M+\lambda_2\bar\nu(u)-2\gamma\nu_u-2\psi q_u^\nu)du-\int_0^T 2\omega_u(b\nu_u+\gamma q_u^\nu)du\right]\\
&=\mathbb{E}^{\mathbb{P}^\alpha}\left[\int_0^T\omega_s\int_s^T(\lambda_1\nu_u^M+\lambda_2\bar\nu(u)-2\gamma\nu_u-2\psi q_u^\nu)duds-\int_0^T 2\omega_u(b\nu_u+\gamma q_u^\nu)du\right]\\
&=\mathbb{E}^{\mathbb{P}^\alpha}\left[\int_0^T\omega_s\left(\int_s^T(\lambda_1\nu_u^M+\lambda_2\bar\nu(u)-2\gamma\nu_u-2\psi q_u^\nu)du-2(b\nu_s+\gamma q_s^\nu)\right)ds\right]\\
&=\int_0^T\mathbb{E}^{\mathbb{P}^\alpha}\left[\omega_s\left(\int_s^T(\lambda_1\nu_u^M+\lambda_2\bar\nu(u)-2\psi q_u^\nu)du-2(b\nu_s+\gamma q_T^\nu)\right)\right]ds\\
&=\int_0^T\mathbb{E}^{\mathbb{P}^\alpha}\left[\omega_s\left(-2(b\nu_s+\gamma q_T^\nu)+\int_s^T\mathbb{E}^{\mathbb{P}^\alpha}\left[\lambda_1\nu_u^M+\lambda_2\bar\nu(u)\mid\mathcal{F}_u\right]-2\psi q_u^\nu du\right)\right]ds\\
&=\mathbb{E}^{\mathbb{P}^\alpha}\left[\int_0^T\omega_s\left(-2(b\nu_s+\gamma q_T^\nu)+\int_s^T\mathbb{E}^{\mathbb{P}^\alpha}\left[\lambda_1\nu_u^M+\lambda_2\bar\nu(u)\mid\mathcal{F}_u\right]-2\psi q_u^\nu du\right)ds\right]
\end{aligned}$$

We have not yet mentioned, but we assume the class of admissible controls to be the space $\mathcal{H}^2$, defined as

$$\mathcal{H}^2=\left\{\phi\in\mathcal{M}\;;\;\mathbb{E}\left[\int_0^T\phi^2(\omega,t)dt\right]<+\infty\right\},$$

where $\mathcal{M}$ is the space of functions $\phi:\Omega\times[0,T]\to\mathbb{R}$ such that ϕ is measurable with respect to $\mathcal{F}\otimes\mathcal{B}([0,T])$ and adapted with respect to $(\mathcal{F}_t)_{t\geq 0}$. $\mathcal{H}^2$ is a vector space over $\mathbb{R}$, and we can define the inner product $\langle f,g\rangle=\mathbb{E}\left[\int_0^T fg\,dt\right]$. Notice now that by the deduction above, $\mathcal{D}L(\nu)(\omega)\in\mathcal{H}^{2*}$, as it is a linear operator on $\mathcal{H}^2$. From the Riesz Representation Theorem, there exists a representative element $h_{L(\nu)}\in\mathcal{H}^2$, which is $(\mathcal{F}_t)_{t\geq 0}$-adapted, such that $\mathcal{D}L(\nu)(\omega)=\langle\omega,h_{L(\nu)}\rangle$. We write an explicit formula for $h_{L(\nu)}$ noticing that

$$
\begin{aligned}
\mathcal{D}L(\nu)(\omega) &= \mathbb{E}^{\mathbb{P}^\alpha}\left[\int_0^T \omega_s\left(-2(b\nu_s+\gamma q_T^\nu)+\int_s^T \mathbb{E}^{\mathbb{P}^\alpha}\left[\lambda_1\nu_u^M+\lambda_2\bar{\nu}(u)\mid\mathcal{F}_u\right]-2\psi q_u^\nu du\right)\right]\\
&= \int_0^T \mathbb{E}^{\mathbb{P}^\alpha}\Big[\omega_s\Big(-2(b\nu_s+\gamma q_T^\nu)-\int_0^s \mathbb{E}^{\mathbb{P}^\alpha}\left[\lambda_1\nu_u^M+\lambda_2\bar{\nu}(u)\mid\mathcal{F}_u\right]-2\psi q_u^\nu\,du\\
&+\int_0^T \mathbb{E}^{\mathbb{P}^\alpha}\left[\lambda_1\nu_u^M+\lambda_2\bar{\nu}(u)\mid\mathcal{F}_u\right]-2\psi q_u^\nu\,du\Big)\Big]\\
&= \int_0^T \mathbb{E}^{\mathbb{P}^\alpha}\Big[\omega_s\Big(-2b\nu_s-\int_0^s \mathbb{E}^{\mathbb{P}^\alpha}\left[\lambda_1\nu_u^M+\lambda_2\bar{\nu}(u)\mid\mathcal{F}_u\right]-2\psi q_u^\nu\,du\\
&+\int_0^T \mathbb{E}^{\mathbb{P}^\alpha}\left[\lambda_1\nu_u^M+\lambda_2\bar{\nu}(u)-2\psi q_u^\nu\mid\mathcal{F}_s\right]du+\mathbb{E}^{\mathbb{P}^\alpha}\left[-2\gamma q_T^\nu\right]\Big)\Big]\\
&= \int_0^T \mathbb{E}^{\mathbb{P}^\alpha}\Big[\omega_s\Big(-2b\nu_s-\int_0^s \mathbb{E}^{\mathbb{P}^\alpha}\left[\lambda_1\nu_u^M+\lambda_2\bar{\nu}(u)\mid\mathcal{F}_u\right]-2\psi q_u^\nu\,du\\
&+\mathbb{E}^{\mathbb{P}^\alpha}\left[-2\gamma q_T^\nu+\int_0^T \lambda_1\nu_u^M+\lambda_2\bar{\nu}(u)-2\psi q_u^\nu du\mid\mathcal{F}_s\right]\Big]
\end{aligned}
$$

Finally, defining $M_s=\mathbb{E}^{\mathbb{P}^\alpha}\left[-2\gamma q_T^\nu+\int_0^T \lambda_1\nu_u^M+\lambda_2\bar{\nu}(u)-2\psi q_u^\nu du\mid\mathcal{F}_s\right]$, M_s is $\mathcal{F}_s$-measurable, and therefore

$$
h_{L(\nu)}=(-2b\nu_s-\int_0^s \mathbb{E}^{\mathbb{P}^\alpha}\left[\lambda_1\nu_u^M+\lambda_2\bar{\nu}(u)\mid\mathcal{F}_u\right]-2\psi q_u^\nu\,du+M_s)_{s\geq 0},
$$

such that

$$
\mathcal{D}L(\nu)(\omega)=\langle\omega,h_{L(\nu)}\rangle.
$$

The major agent has 2 independent controls: trading speed and influence level. Therefore, using similar arguments, we can write

$$
\mathcal{D}L^M(\nu^M,c)(\omega,\chi)=\langle\omega,h_{L^M(\nu,c)}\rangle+\langle\chi,\tilde{h}_{L^M(\nu^M,c)}\rangle,
$$

where

$$
h_{L^M(\nu)}=(-2b_M\nu_s^M-\int_0^s \mathbb{E}^{\mathbb{P}}\left[\lambda_1\nu_u^M+\lambda_2\bar{\nu}(u)\mid\mathcal{F}_u\right]-2\psi^M q_u^{\nu^M}\,du+N_s)_{s\geq 0},
$$

where

$$
N_s=\mathbb{E}^{\mathbb{P}}\left[-2\gamma^M q_T^{\nu^M}+\int_0^T \lambda_1\nu_u^M+\lambda_2\bar{\nu}(u)-2\psi^M q_u^{\nu^M}du\mid\mathcal{F}_s\right].
$$

And for the influence level part, we have

$$
\tilde{h}_{L^M(\nu^M,c)}=(2\zeta c_s)_{s\geq 0}.
$$

□

References

Franklin Allen and Douglas Gale. Stock-price manipulation. *The Review of Financial Studies*, 5 (3):503–529, 1992.

Philippe Bergault and Leandro Sánchez-Betancourt. A mean field game between informed traders and a broker. *arXiv preprint arXiv:2401.05257*, 2024.

Pierre Cardaliaguet and Charles-Albert Lehalle. Mean field game of controls and an application to trade crowding. *Mathematics and Financial Economics*, 12:335–363, 2018.

Álvaro Cartea and Sebastian Jaimungal. Incorporating order-flow into optimal execution. *Mathematics and Financial Economics*, 10:339–364, 2016.

Álvaro Cartea, Ryan Donnelly, and Sebastian Jaimungal. Algorithmic trading with model uncertainty. *SIAM Journal on Financial Mathematics*, 8(1):635–671, 2017.

Álvaro Cartea, Sebastian Jaimungal, and Yixuan Wang. Spoofing and price manipulation in order-driven markets. *Applied Mathematical Finance*, 27(1-2):67–98, 2020.

Xuancheng Huang, Sebastian Jaimungal, and Mojtaba Nourian. Mean-field game strategies for optimal execution. *Applied Mathematical Finance*, 26(2):153–185, 2019.

Gur Huberman and Tomer Regev. Contagious speculation and a cure for cancer: A nonevent that made stock prices soar. *The Journal of Finance*, 56(1):387–396, 2001.

Roni Michaely and Kent L Womack. Conflict of interest and the credibility of underwriter analyst recommendations. *The review of financial studies*, 12(4):653–686, 1999.