Whence LASSO? A Rational Interpretation

Wen Chen, Bo Hu, Liyan Yang *

July 29, 2023

Abstract

This paper rationalizes the LASSO algorithm based on uncertain fat-tail priors and max-min robust optimization. Our rationalization excludes heuristic learning or restrictive prior assumptions in the original interpretation of LASSO (Tibshirani (1996)). In our setting, economic agents (arbitrageurs) face ambiguity about fat-tail shocks and in equilibrium, they ignore a reasonable range of ambiguous signals but respond linearly to almost unambiguous signals. With this LASSO equivalent strategy, arbitrageurs can amass extra market power which induces a "cartel" to protect their aggregate profit from being competed away. This result shows a new mechanism for limited arbitrage.

Keywords: LASSO, fat tails, model risk, robust optimization, limits to arbitrage

*Chen is from Chinese University of Hong Kong (Shenzhen) School of Economics and Management, email: wenchen@cuhk.edu.cn. Hu is from George Mason University School of Business, email: bhu5@gmu.edu. Yang is from University of Toronto Rotman School of Management. Please send correspondence to Liyan Yang, email: liyan.yang@rotman.utoronto.ca. The authors are grateful to Andrea Buffa, John Chao, Hui Chen, Alex Chinco, Steve Heston, Pete Kyle, Ye Li, Mark Loewenstein, Semyon Malamud, Richmond Mathews, Alberto Rossi, Shrihari Santosh, Duane Seppi, Hyun-Song Shin, Yajun Wang, Wei Xiong, and Guofu Zhou for helpful comments. We also thank the audiences at 2021 WFA, 2023 Kentucky Finance Conference, 2023 FTG Summer School, 2023 CICF, and 2023 XJTLU AI and Big Data in Accounting and Finance Conference.

1 Introduction

Machine learning is widely used in economics and finance to make predictions, classifications, and decisions based on sample data (Jordan and Mitchell (2015), Athey (2018), Nagel (2021), and Hastie, Tibshirani, Friedman, and Friedman (2009)). Most techniques were developed in the fields of statistics and computer science. They have been proven, by numerous applications, to be superior in prediction accuracy and computational efficiency. From an economic perspective, there are at least two important questions to be addressed: (1) Are machine-learning methods rational choices for economic agents? (2) What would happen if machine learning becomes a new doctrine in financial markets? Answering both questions can provide insights into new topics in risk management and asset pricing.

This paper first rationalizes a widely used machine learning algorithm, the *Least Absolute Shrinkage and Selection Operator* (LASSO), invented by the renowned statistician Robert Tibshirani in 1996. This is a linear regression with an l_1 norm penalty in its loss function. The l_1-penalty is critical for the LASSO to achieve both variable selection and shrinkage. It may be explained by the argument of *Occam's Razor* or the principle of parsimony, which is akin to behavioral economics. Tibshirani (1996) formalized a statistical interpretation that the LASSO regression can be derived using the *maximum a posteriori* (MAP) estimate under a Laplace (double exponential) prior on the estimated parameter. This interpretation is not Bayesian rational. The MAP is a heuristic learning rule that uses the posterior *mode* as the point estimate, instead of integrating all the posterior information. This argument is also restrictive, as it only works with a pure and fixed Laplace prior.

An economic setting is required to rationalize an algorithm. We design an arbitrage trading model to show that the LASSO algorithm can be an equilibrium strategy chosen by Bayesian-rational agents (traders) when they have uncertain fat-tail priors (model risks). Our proof does not rely on a heuristic (MAP) learning rule or the restrictive assumption of a pure, fixed Laplace prior. In this regard, we provide the first economic rationale for the LASSO. The argument highlights the robustness of LASSO which stems from an inaction region endogenously chosen by agents to ignore directionally ambiguous signals. It sacrifices little optimality, because the robust strategy also responds in a nearly optimal manner to the directional, most profitable signals. We find that the robust LASSO strategy allow agents to accumulate extra market power that protects their total profit even when their population grows to infinity, providing an interesting channel for inefficient markets.

As a statistical tool, LASSO and its variants have gained tremendous popularity in empirical finance; see the surveys by Giglio, Kelly, and Xiu (2022) and Kelly and Xiu (2023).[1]

[1]For instance, Rapach, Strauss, and Zhou (2013) apply the LASSO to study lead-lag relationships among

However, there are only limited theoretical studies that examine the LASSO from an economic perspective. Gabaix (2014) proposes a sparsity model, in the spirit of LASSO, for the *anchoring-and-adjustment* bias (Tversky and Kahneman (1974)) and the limited attention (Sims (2003)). Martin and Nagel (2022) develop a model to study cross-sectional return predictability due to sparsity or shrinkage when investors face a high-dimensional prediction problem. Both papers take the LASSO as an outcome of *bounded rationality*, agreeing with Tibshirani's heuristic interpretation.

The MAP method uses the posterior mode (instead of the mean) as the point estimate. A rational agent should not use the MAP estimate since it can ignore valuable information. Tibshirani's interpretation assumes the agent's prior is a pure and fixed Laplace distribution. With a sharp peak at the origin, this prior allows the MAP learning rule to generate sparse solutions. All too often, the Laplace prior assumption is not carefully verified or justified. It remains unclear whether this prior is valid in many applications of LASSO. If the Laplace prior is just an approximation, it is unclear how good the approximation is. A further issue is model uncertainty. For example, even though the Laplace distribution can fit the histogram of past stock returns,[2] our prior knowledge about future stock returns can be ambiguous and we cannot simply apply the MAP estimate when the prior parameter(s) are uncertain. As a bottom line, if the applicability of LASSO hinges on a fixed Laplace prior, then it may work under certain circumstances but cause unexpected problems in other conditions.

Why is it important to look for an economic rationale for a simple algorithm? This can be pedagogical in our view. Such a theoretical attempt is missing in the literature, perhaps for two reasons:

1. The MAP-based interpretation theoretizes the LASSO as a scientific tool using Bayes' rule and hence lends support to its *mathematical* legitimacy. This is undoubtedly the merit of Tibshrina's work. Given the power of LASSO, users may not care much about its *economic* legitimacy and thus leave some fundamental questions open. This is not a unique example. Researchers are endowed with a large library of machine learning

monthly international stock returns. Goto and Xu (2015) use the graphical LASSO to solve a sparse estimator of the inverse covariance matrix in mean-variance portfolio optimization. Chinco, Clark-Joseph, and Ye (2019) apply the LASSO to select short-term predictors and forecast individual stock returns one-minute ahead, given cross-sectional returns over the past few minutes. Gu, Kelly, and Xiu (2020) apply a zoo of machine learning tools, including the LASSO, to study the time-series predictability of monthly individual stock returns. Freyberger, Neuhierl, and Weber (2020) apply the adaptive group LASSO to identify the relationships between numerous firm characteristics and cross-section of expected returns. Kozak, Nagel, and Santosh (2020) utilize the elastic net, a variant of LASSO, to construct a robust stochastic discount factor which integrate the explanatory power of a large number of cross-sectional return predictors. Dong, Li, Rapach, and Zhou (2022) employ a variety of LASSO-related shrinkage techniques to extract predictive signals from long-short anomaly portfolio returns in a high-dimensional setting. Huang and Shi (2022) apply the adaptive group LASSO to government bonds and construct a macro factor based on 30 predictors.

[2]See Mantegna and Stanley (1999), Lillo and Mantegna (2000), Silva, Prange, and Yakovenko (2004).

tools. The literature has also fed us with abundant technical knowledge about them. However, there is only meager understanding of whether and why those tools are economically sensible. A similar phenomenon, as noted by McQueen and Vorkink (2004), is in statistical models of volatility clustering, such as the *autoregressive conditional heteroskedasticity* (ARCH) model (Engle (1982)) and the generalized ARCH model (Bollerslev (1986)). The popularity of using these models stems from their power of fitting the data, but "*our theoretical knowledge of why volatility clusters is paltry.*"

2. Different fields have developed different systems of topics, methods, and standards. For example, the MAP estimate is instructed in most statistics and machine learning courses but rarely mentioned in economics or finance textbooks. This kind of difference can delay scientific discoveries and call for interdisciplinary efforts. As a related implication, the financial industry has hired many quants with strong training in math, physics, and engineering. Their similar backgrounds may reinforce the quant mindset but compromise the economic thinking. This may build up systemic risks, as probably exemplified by the *quant meltdown* in 2007 (Khandani and Lo (2011), Mussalli (2018)).

Therefore, more theoretical research is needed at the interface between machine learning and financial economics. Though the emerging literature leans toward a behavioral argument (Gabaix (2014), Mullainathan and Spiess (2017), and Camerer (2018)), it does not preclude the possibility of rational theories or improvements. For example, can the LASSO algorithm be an equilibrium strategy? If no, there should be profitable deviations which may improve the original LASSO. If yes, it is something new and needs to be formally addressed.

We have asked two general questions at the beginning. To address them for any particular technique (e.g., LASSO), we need a concrete setting to specify the agents' objective functions and define the equilibrium. This approach has two benefits. First, it can help the audience understand the economic mechanism. Second, it allows us to perform normative analysis of important issues such as market stability, price efficiency, and the performance of strategies.

This paper shows that the LASSO method can be a rational choice for economic agents with robust optimization.[3] The robustness comes from the inaction zone of LASSO, as it can filter out most ambiguous signals but keep unambiguous ones. Economically, if agents are uncertain about the direction of some fat-tailed signals, they may only respond to strong

[3]The machine learning literature only takes robustness as a secondary property of LASSO; see Hastie, Tibshirani, and Wainwright (2015). The argument is provided by modifying the loss function in the LASSO estimation problem (Xu, Caramanis, and Mannor (2008)). This *ad hoc* method is widely used in statistics but not adopted in economics. For example, the MAP rule may be "theoretized" by using the "hit-or-miss" loss function (Robert et al. (2007) [p. 166]). However, this derivation does not *rationalize* the MAP. To establish an economic rationale for machine learning, we have to formulate the utility functions that agents optimize, rather than tampering with the loss function for their learning rules.

signals and ignore vague ones. Inaction can avoid betting on the wrong side and suffering from fat-tailed losses. To formalize this intuition, we focus on two related features: uncertain fat-tail priors and robust optimization. Both issues concern arbitrageurs in financial markets.

Specifically, we develop a trading model where ambiguity-averse arbitrageurs predict and exploit pricing errors caused by random fat-tail shocks. We use a general Gaussian-Laplacian mixture distribution for the stock value. Given a linear pricing rule,[4] random fat-tail shocks can produce disproportionate pricing errors, the frequency and magnitude of which are tuned by the mixing weight and the fat-tail scale parameter, respectively. Arbitrageurs are uncertain about the scale parameter. Each of them makes robust trading decisions by optimizing the *max-min expected utility*, as axiomatized by Gilboa and Schmeidler (1989).

We show that the equilibrium robust strategy chosen by arbitrageurs is equivalent to the LASSO estimate of pricing errors, conditional on the order flow (or the price change) observed in a time window right before their trading. Specifically, the strategy has an endogenous threshold inversely related to the averaged scale of fat-tail shocks but independent of their frequency, whereas the response intensity beyond the inaction zone is proportional to the frequency but independent of the scale parameter. Thus, under fairly general conditions, we show that the use of LASSO is a Bayesian rational strategy optimally chosen by agents who are concerned about fat-tailed model risks. This economic interpretation does not use any heuristic learning rule, nor make the restrictive assumption of a pure and fixed Laplace prior, which is an extreme case of our assumed general mixture distribution.

We also find that the robust LASSO strategy can outperform the optimal (benchmark) strategy, a nonlinear smooth response, which ignores the model risk and optimizes the subjective expected utility.[5] The benchmark strategy can easily lose profits if the estimate deviates from the true parameter or if the number of competitors increases. In contrast, the LASSO strategy is much more robust to the estimate bias because its inaction zone avoids small but frequent mistakes. More importantly, its performance is much less susceptible to traders' competition. We prove that even as the number of arbitrageurs goes to infinity, their aggregate profit does not vanish but converges to a positive level. This effect is induced by the under-trading (shrinkage) of the LASSO strategy. Its conservativeness mitigates traders' competition, allowing them to amass extra market power to protect their profit from being

[4]The empirical price impact function, which measures the average price change in response to the size of an incoming order, is sublinear with some concavity. See Loeb (1983), Grinold and Kahn (2000) [p. 453], Gabaix, Gopikrishnan, Plerou, and Stanley (2006), and Kyle and Obizhaeva (2016). The linear pricing rule can be endogenized in multi-period Kyle-type models (e.g., Kyle (1985), Holden and Subrahmanyam (1992), Foster and Viswanathan (1994, 1996)) by assuming that market makers adhere to the Gaussian belief or restrict their considerations to linear pricing strategies perhaps for simplicity and robustness.

[5]In our setup, the MAP-based strategy differs from the LASSO strategy if agents' prior is not exactly Laplacian. This heuristic MAP strategy is not an equilibrium outcome and thus can incur significant losses.

competed away. Even an infinite number of traders can act as if they were a monopolist who buys or sell the asset at a better price than the fair one. This seemingly collusive behavior does not involve any trading or financial constraints, nor require any communication device or explicit agreement. The "cartel" is facilitated tacitly by traders' uncoordinated exercise of risk management. Thus, our model describes a novel channel for the limits to arbitrage. This agrees with the finding of Da, Nagel, and Xiu (2022) that arbitrage strategies based on LASSO or ridge methods can only achieve a relatively low Sharpe ratio out of sample, despite a much higher but infeasible Sharpe ratio. Alphas can survive in equilibrium when ambiguity-averse arbitrageurs lack perfect knowledge about the trading environment.

Our two-period trading model describes a market with *short-lived* and *infrequent* return predictability. This agrees with the findings of Chinco et al. (2019) who apply the LASSO to select a small set of short-term predictors from thousands of candidate stocks. Yet they acknowledged that "*the LASSO identifies predictors that are not easy to intuit.*" With Tibshirani's interpretation, their LASSO regressions would implicitly assume a Laplace prior on the predictive power (i.e., the regression coefficient) of each predictor. Our model can provide a different view. An econometrician may form a general mixture prior on the predictive power of each predictor in the context of Chinco et al. (2019). If she has little knowledge about the scale of fat-tailed outliers, it sounds reasonable to invoke a robust estimate (i.e., the LASSO regression) to shrink most ambiguous estimates to zero.

In our model, traders apply the LASSO estimate to the stock value, not to the stock's predictive power for other stocks. This follows the stylized fact that the distribution of stock returns has a sharp peak with fat tails on both sides; see Fama (1963, 1965), Granger and Ding (1995), and Mantegna and Stanley (1999). It is error-prone to predict extreme events (e.g., Embrechts, Klüppelberg, and Mikosch (2013)). This leads to model risks which can motivate traders to implement robust optimization, voluntarily or mandatorily. Extending our setup to a large number of stocks, we are able to provide an intuitive explanation for the sparse, cross-sectional return predictability documented in Chinco et al. (2019). This is also consistent with the framework of principal portfolios developed by Kelly, Malamud, and Pedersen (2023), as all assets' signals may help predict each individual asset return.

Our work attempts to bridge the gap between machine learning (e.g., the LASSO) and financial economics. The basic model is inspired by the fat tails in asset prices. Unexpected fat-tail shocks can cause temporarily inefficient prices. We integrate this feature into the classic framework of Kyle (1985).[6] We also adopt the max-min criterion to model ambiguity

[6]This framework has been studied by many others; see Back (1992), Holden and Subrahmanyam (1992), Foster and Viswanathan (1994, 1996), Vayanos (1999, 2001), Back, Cao, and Willard (2000), Collin-Dufresne and Fos (2016), Yang and Zhu (2020), among others.

aversion within the framework of Gilboa and Schmeidler (1989). See also Schmeidler (1989), Dow and Werlang (1992), Hansen and Sargent (2001, 2008), Garlappi, Uppal, and Wang (2006), Epstein and Schneider (2008, 2010), Easley and O'Hara (2009, 2010), Illeditsch (2011), Banerjee, Davis, and Gondhi (2019), among others. While Klibanoff, Marinacci, and Mukerji (2005) propose a smooth preference model of decision making under ambiguity, we adopt the kinked preference as it is compatible with the experimental results of Bossaerts, Ghirardato, Guarnaschelli, and Zame (2010) and Ahn, Choi, Gale, and Kariv (2014).

Our results shed light on a new mechanism for limited arbitrage. The literature has studied various market frictions, such as short-selling costs, leverage constraints, and wealth effects, which directly limit arbitrageurs' *ability* to trade; see the survey of Gromb and Vayanos (2010) and the work of Shleifer and Vishny (1997), Xiong (2001), Abreu and Brunnermeier (2002), Gabaix, Krishnamurthy, and Vigneron (2007), Kondor (2009), among others. By excluding those frictions, our model can highlight a mechanism that only affects arbitrageurs' *willingness* to trade. It is the prior uncertainty about the fat-tailed alphas that can deter arbitrageurs from eliminating all possible mispricings. The resulted conservatism also mitigates trading competition and prevents asset prices from being efficient. This inefficiency persists even when the economy hosts an infinite number of (risk-neutral) arbitrageurs.

Finally, our model framework may help us better understand algorithmic trading. The literature has addressed many market implications of algorithmic traders; see Hendershott, Jones, and Menkveld (2011), Brogaard, Hendershott, and Riordan (2014), and Van Kervel and Menkveld (2019). While simple algorithms have an advantage in speed (Lewis (2014), Biais, Foucault, and Moinas (2015), Budish, Cramton, and Shim (2015)), they may also appear at odds with the Bayesian learning. In our model, agents rationally use a "mechanical" trading algorithm in equilibrium, to be distinguished from behavioral traders who extrapolate price trends (DeLong, Shleifer, Summers, and Waldmann (1990), Barberis, Greenwood, Jin, and Shleifer (2015, 2018)). Recently, Colliard, Foucault, and Lovo (2022) and Dou, Goldstein, and Ji (2023) advocate the Q-learning algorithm which enables traders to learn about the environment by trial and error with almost no prior knowledge. This is interesting and useful, though different from the standard Bayesian framework followed by our model.

The rest of this paper proceeds as follows. Section 2 discusses the definition and interpretation of LASSO. Section 3 describes our model. Section 4 solves the equilibrium problems. Our main results are presented in Section 5, with extensions and applications in Section 6. We make the concluding remarks in Section 7. All proofs are provided in the Appendix.

2 Definition and Interpretation of the LASSO

The LASSO can perform variable selection and regularization simultaneously. In statistics, *regularization* refers to a technical process that can help simplify the solutions or models, for example, to obtain approximate solutions for ill-posed problems or to prevent overfitting (Hastie et al. (2009)). There is a statistical argument which takes many regularization methods as equivalent to imposing some prior distributions on parameters. The LASSO features an l_1 regularization, equivalent to imposing a Laplace prior. This l_1 penalty is essential for its ability to improve both prediction accuracy and model selection. In contrasts, the *ridge regression* involves an l_2 regularization which is equivalent to imposing a Gaussian prior.

In its general form, LASSO is applied to a sample of $I \geq 1$ pairs of predictor-response observations, $\{x_i, y_i\}_{i=1}^I$, where x_i is a vector of $J \geq 1$ covariates (independent variables). After demeaning of the data, the LASSO estimates of $\mathbf{v} = (v_1, ..., v_J)$ are defined by

$$\widehat{\mathbf{v}}^{\text{lasso}} := \arg\min_{\mathbf{v}} \left\{ \frac{1}{2I} \|\mathbf{y} - \mathbf{X} \cdot \mathbf{v}\|_2^2 + \rho \|\mathbf{v}\|_1 \right\}, \tag{1}$$

where $\mathbf{y} = (y_1, ..., y_I)$ is the I-vector of dependent variables, $\mathbf{X}$ is an $I \times J$ covariate matrix, and the positive scalar ρ is the l_1 regularization parameter tuned exogenously. This l_1 penalty forces many insignificant coefficients to zero. Prediction accuracy is improved by sacrificing some bias to mitigate prediction errors. LASSO can also reduce a high-dimensional prediction problem to a much simpler one with a sparse subset of predictors.

The simplest version of LASSO corresponds to the setting of $I = J = 1$. Given a single predictor-response pair $\{x_1, y_1\}$, the optimization problem (1) becomes

$$\widehat{v}^{\text{lasso}} := \arg\min_{v} \left\{ \frac{1}{2} |y_1 - x_1 v|^2 + \rho |v| \right\}. \tag{2}$$

The solution to the above problem is given by

$$\widehat{v}^{\text{lasso}} = \begin{cases} (x_1 y_1 - \rho)/x_1^2, & \text{if} \quad x_1 y_1 > \rho, \\ 0, & \text{if} \quad |x_1 y_1| \leq \rho, \\ (x_1 y_1 + \rho)/x_1^2, & \text{if} \quad x_1 y_1 < -\rho. \end{cases} \tag{3}$$

It is convenient to introduce and define the *soft-thresholding* operator:

$$\mathcal{S}(y; K) = \text{sign}(y) \max(|y| - K, 0) = [y - \text{sign}(y) K] \mathbf{1}_{|y| > K}. \tag{4}$$

Then the LASSO solution (3) can be concisely written as

$$\widehat{v}^{\text{lasso}} = \mathcal{S}(y_1/x_1; \rho/x_1^2) = \mathcal{S}(y_1; \rho/x_1)/x_1. \tag{5}$$

Tibshirani proposes that the LASSO can be viewed as the MAP (i.e., posterior mode) estimates of the linear regression coefficients with Laplace priors. For the simplest version (2), assume that the prior on v follows a Laplace distribution

$$f_L(v) = \frac{1}{2\xi} \exp\left(-\frac{|v|}{\xi}\right). \tag{6}$$

This density function is sharply peaked since its first derivative is discontinuous at zero. It decays on both sides at the exponential rate ξ and has a raw kurtosis always equal to 6. The likelihood of observing extreme events under a Laplace distribution is much higher than that under the Gaussian distribution with an identical variance. For a linear model $y_1 = x_1\tilde{v} + \tilde{u}_1$ with a Gaussian noise term $\tilde{u}_1 \sim \mathcal{N}(0, \sigma_u^2)$, the posterior distribution of $\tilde{v}$ is

$$f(v|y_1) = \frac{f(y_1|v)f_L(v)}{f(y_1)} = \frac{1}{2\xi f(y_1)\sqrt{2\pi\sigma_u^2}} \exp\left\{-\frac{(y_1 - x_1 v)^2}{2\sigma_u^2} - \frac{|v|}{\xi}\right\}. \tag{7}$$

One can then solve for the MAP estimate under the Laplace prior $\tilde{v} \sim \mathcal{L}(0, \xi)$,

$$\widehat{v}^{\text{map},\mathcal{L}} = \arg\max_v f(v|y_1) = \arg\min_v \left\{\frac{(y_1 - x_1 v)^2}{2\sigma_u^2} + \frac{|v|}{\xi}\right\} = \frac{1}{x_1}\mathcal{S}(y_1; \sigma_u^2/(x_1\xi)) = \widehat{v}^{\text{lasso}}. \tag{8}$$

This coincides with the LASSO estimate (3) if we assign $\rho = \sigma_u^2/\xi$. The MAP coupled with a sharply peaked Laplace prior can effectively shrink a range of estimates to zero. However, the definition of LASSO (in its Lagrangian form) does not require the specification of a prior. Only the interpretation of LASSO needs one to specify the prior. The above statistical interpretation of LASSO has two problems: First, a Bayesian rational agent should use the posterior mean estimate which integrates all the relevant information. The MAP mode estimate can miss useful posterior information. Second, the MAP-based argument can only work with a pure and fixed Laplace prior. If we assume a slightly different prior or consider any uncertainty about the prior, the MAP argument will not produce the LASSO objective.[7]

In summary, the statistical interpretation of LASSO is both heuristic (in learning) and restrictive (in prior). This does not mean that the LASSO itself is heuristic or restrictive. We will show that it admits a Bayesian rational interpretation under flexible, uncertain priors.

[7] A pure and fixed Laplace prior seems too restrictive. One might think of this as an approximation or a simple device to model sparsity. The point is that the general definition (1) of LASSO does not enforce such a prior. LASSO is invented as a tool to solve sparse learning problems, not a model to describe them.

3 Model

Arbitrage opportunities are often short-lived and unexpected. We will develop a model to show how traders may be motivated to use the LASSO to capitalize on such opportunities. We first describe a realistic trading environment which produces pricing errors occasionally. This is modelled in a two-period setup where random fat-tail shocks can occur and disrupt a presumably efficient market. We then allow arbitrageurs to exploit mispricings.

Our model is built on three key assumptions: (1) the general Laplacian-Gaussian mixture distribution for the asset's liquidation value; (2) the linear price adjustments in responses to total order flows, which contain informed trading proportional to the residual information; (3) the strategic arbitrageurs who secretly exploit pricing errors by optimizing their max-min expected utilities under uncertain fat-tail priors (as model risk). The linear pricing schedules and the non-Gaussian asset values can jointly imply the occurrences of pricing errors. Our focus is to study how arbitrageurs would trade in this uncertain fat-tailed environment.

Fat Tails. Consider a market with a single risky asset and two rounds of trading, indexed by $t = 1, 2$. The asset liquidation value follows a Laplacian-Gaussian mixture distribution, which is denoted as $\tilde{v} \sim \mathcal{LG}(\alpha, \xi_v)$ and described by the probability density function,

$$f(v) = \frac{\alpha}{2\xi_v} \exp\left(-\frac{|v|}{\xi_v}\right) + \frac{1-\alpha}{\sqrt{2\pi\sigma_v^2}} \exp\left(-\frac{v^2}{2\sigma_v^2}\right). \tag{9}$$

It depends on the mixing weight α (the frequency of fat-tail shocks), the fat-tail scale parameter ξ_v (the dispersion of fat-tail shocks), and the Gaussian variance σ_v^2. In simulations, $\tilde{v}$ is randomly drawn from either a Gaussian or a Laplacian distribution, since we can write

$$\tilde{v} = (1-\tilde{s}) \cdot \tilde{v}_G + \tilde{s} \cdot \tilde{v}_L, \quad \text{with} \quad \tilde{v}_G \sim \mathcal{N}(0, \sigma_v^2) \quad \text{and} \quad \tilde{v}_L \sim \mathcal{L}(0, \xi_v), \tag{10}$$

where $\tilde{s}$ is a Bernoulli random variable that is equal to 1 with probability α and to 0 with probability $1-\alpha$. In Eq. (10), the distributional type of $\tilde{v}$ is encoded by the value of $\tilde{s}$.

Empirically, the distribution of stock returns exhibits one sharp peak and two fat tails. It can be reasonably characterized by the above mixture distribution (Lillo and Mantegna (2000), Silva et al. (2004), Haas, Mittnik, and Paolella (2006), Behr and Pötter (2009)). The picture is also consistent with the stylized fact that the stock market experiences jumps.

Theoretically, the mixture prior combines two well-known distributions which are simple and stable.[8] It may be viewed as a microstructure snapshot of the jump-diffusion process

[8] As discussed by Fama (1963, 1965) and Rachev and SenGupta (1993), it is appealing to model stock returns as realizations of some stable distribution. The Gaussian distribution is Lévy stable, whereas the Laplace distribution is geometric stable. See Kotz, Kozubowski, and Podgórski (2001).

assumed in the option pricing model of Kou (2002). The fat-tail component ($\tilde{v}_L$) in Eq. (10) can create sparse arbitrages if the general market believes that both informed and noise demands follow Gaussian distributions (e.g., Kyle (1985)). The general structure of Eq. (9) covers the Laplace distribution as a special case (i.e., $\alpha = 1$). This is convenient when we refer to the statistical interpretation of LASSO since it hinges on a pure Laplace prior.

Linear Pricing. Price movements are assumed to be linear in the total order flows $\tilde{y}_t$:

$$\tilde{p}_1 - p_0 = \lambda_1 \tilde{y}_1, \qquad \tilde{p}_2 - \tilde{p}_1 = \lambda_2 \tilde{y}_2, \tag{11}$$

where $\lambda_t > 0$ is the price impact per unit of order flow at time t. We can set the initial price $p_0 = 0$ without loss of generality. The total order flow $\tilde{y}_t$ contains private information which, by dynamical consistency, is proportional to the residual information $\tilde{v} - \tilde{p}_{t-1}$. This is a common feature of dynamic trading models following Kyle (1985). To avoid the information from being fully revealed, the order flows are contaminated by noise trading demands, which obey the Gaussian law, $\tilde{u}_1 \sim \mathcal{N}(0, \sigma_u^2)$ and $\tilde{u}_2 \sim \mathcal{N}(0, \gamma\sigma_u^2)$, with time-varying volatilities tuned by the parameter $\gamma > 0$. All the random variables $\tilde{v}$, $\tilde{u}_1$, and $\tilde{u}_2$ are mutually independent. Thus, before considering the arbitrage trading, the total order flows are

$$\tilde{y}_1 = \beta_1(\tilde{v} - p_0) + \tilde{u}_1, \qquad \tilde{y}_2 = \beta_2(\tilde{v} - \tilde{p}_1) + \tilde{u}_2, \tag{12}$$

where $\beta_t > 0$ is the trading intensity on the remaining information $\tilde{\theta}_t := \tilde{v} - \tilde{p}_{t-1}$ at time t. Eq. (11) and Eq. (12) together describe the trading environment of our model. This can be embedded in a linear subgame-perfect Markov equilibrium of a classical trading model, such as Kyle (1985), Holden and Subrahmanyam (1992), Foster and Viswanathan (1996). An example (microfoundation) is presented below, with details discussed in Appendix A.1.

Example. Consider the two-period case of the multi-period model in Holden and Subrahmanyam (1992), where $M \geq 1$ informed traders privately observe the value of $\tilde{v}$ at $t = 0$. Market makers believe that $\tilde{v}$ is drawn from the Gaussian distribution $\mathcal{N}(0, \sigma_v^2)$. There exists a unique linear equilibrium where the asset price moves linearly as in Eq. (11) with

$$\lambda_1 = \frac{\sqrt{M(M+1)^2[(M+1)^2 - 2/\delta]}}{(M+1)^3 - 2M/\delta} \cdot \frac{\sigma_v}{\sigma_u}, \qquad \lambda_2 = \delta\lambda_1 = \sqrt{\frac{\delta M/\gamma}{\delta(M+1)^3 - 2M}} \cdot \frac{\sigma_v}{\sigma_u}. \tag{13}$$

Here, the ratio $\delta := \lambda_2/\lambda_1$ is the root of the cubic equation, $(M+1)^4\gamma\delta^3 - 2(M+1)^2\gamma\delta^2 - (M+1)^3\delta + 2M = 0$, subject to the constraint $\delta(M+1)^2 > 2$ and the second-order conditions. The aggregate order flows in this example obey the form of Eq. (12). The intensities of aggregate informed trading in Eq. (12) are found to be

$$\beta_1 = \frac{\delta M(M+1)^2 - 2M}{\lambda_1[\delta(M+1)^3 - 2M]}, \qquad \beta_2 = \frac{M}{\lambda_2(M+1)}. \tag{14}$$

When $M = 1$, Eq. (13) and Eq. (14) reproduce the solution for the two-period Kyle model.

In general, the linear pricing rule (11) is efficient only when $\tilde{v}$ is Gaussian (i.e., $\alpha = 0$). If $\tilde{v}$ actually follows the mixture distribution (9) with $\alpha \in (0, 1]$ and if its fat-tailed part is unexpected by the market, then pricing errors will take place with probability α. The generic results in this paper do not depend on the content of λ_t or β_t. For this reason, both λ_t and β_t will be treated as exogenously given and commonly known. The above example provides a microfoundation which is not unique but general enough for numerical purposes.

The assumption of linear price changes (11) has both theoretical and empirical relevance. Huberman and Stanzl (2004) show that if the price impact of trades is both permanent and time-independent, then only linear price impact functions can rule out quasi-arbitrage and support viable market prices. Empirically, the price impact function is found to be sublinear for small and medium-size orders, with moderate concavity for large orders.[9]

The linear pricing rule implies that market makers may have Gaussian beliefs or other concerns (e.g., robustness) that effectively restrict themselves to the linear pricing strategy. It is perhaps the case that the market has ignored some sparse, short-lived arbitrages (Chinco et al. (2019)) or that some information is too subtle to be noticed by average traders (Deng, Gao, Hu, and Zhou (2020)). Price inefficiency caused by incorrect beliefs may not persist in the long term because the market may gradually learn and fix the problem. If we follow the standard assumption in Kyle-type models that market makers know the true distributions and set prices as efficient as possible, then their optimal pricing strategy should be nonlinear and convex given the fat tails of $\tilde{v}$. This is at odds with the empirical price impact function. Another possibility is that the objective of market makers in reality may be different from what has been assumed in standard models. When the trading environment receives various uncertain shocks, market makers may become concerned more about the robustness than about the efficiency of their pricing schedules. The linear pricing rule is simple and robust, allowing them to readily tune the price impact parameters λ_t and avoid losses on average. This argument can sustain inefficient prices, The linear pricing rule can lose efficiency when the asset value is fat-tailed: it tends to underestimate the information content in large orders. The frequency and the magnitude of mispricings are controlled by α and ξ_v, respectively. This is how our setup generates arbitrage opportunities and opens the door to arbitrageurs.

[9] See Loeb (1983), Lillo and Mantegna (2000), Grinold and Kahn (2000) [p. 453], Hasbrouck and Seppi (2001), Plerou, Gopikrishnan, Gabaix, and Stanley (2002), Gabaix et al. (2006), Kyle and Obizhaeva (2016).

Ambiguity-Averse Arbitrageurs. Suppose a number of arbitrageurs, indexed by $n = 1, ..., N$, know the distributional structure of $\tilde{v}$. They do not observe $\tilde{v}$ until it is revealed at $t = 3$. Each arbitrageur can secretly place two orders, $\tilde{z}_{1,n}$ and $\tilde{z}_{2,n}$, to exploit the short-term mispricings.[10] Their strategy profile is denoted by a matrix of real-valued functions, $\mathbf{Z} = [\mathbf{Z}_1, ..., \mathbf{Z}_N]$ where $\mathbf{Z}_n = \langle Z_{1,n}, Z_{2,n} \rangle$ is the n-th arbitrageur's strategy. Then the quantities traded by the n-th arbitrageur are $\tilde{z}_{1,n} = Z_{1,n}(p_0)$ and $\tilde{z}_{2,n} = Z_{2,n}(\tilde{p}_1)$, and her profit is $\tilde{\pi}_{z,n} \coloneqq \sum_{t=1}^{2} (\tilde{v} - \tilde{p}_t)\tilde{z}_{t,n}$. With their trading activities, the actual order flow at time t is

$$\tilde{y}_t = \beta_t(\tilde{v} - \tilde{p}_{t-1}) + \sum\nolimits_{n=1}^{N} \tilde{z}_{t,n}(\tilde{p}_{t-1}) + \tilde{u}_t. \tag{15}$$

Arbitrageurs are strategic institutional traders who know the number (N) of competitors. In our model, they could infer α from the sample kurtosis of realized stock values. Note that our key results remain when traders are uncertain about α.[11] Given the difficulty in predicting jump events (Bollerslev and Todorov (2011a,b)), we assume that arbitrageurs do not know exactly when the pricing errors would take place (i.e., they do not observe $\tilde{s}$ in Eq. (10)) or how large the errors would be (i.e., they are uncertain about ξ_v).

To formalize the model risk, we express arbitrageurs' uncertain priors as $\tilde{v} \sim \mathcal{LG}(\alpha, \tilde{\xi})$, where $\tilde{\xi} \in [\xi_L, \xi_H]$ can take any value between the lowest and highest priors. Arbitrageurs seek robust trading strategies under this model risk. Gilboa and Schmeidler (1989) axiomatize the *max-min expected utility* theory as a standard, rational framework for modeling ambiguity-averse preferences. We follow this classic theory to assume that each arbitrageur's objective is to maximize the minimum expected trading profit over all possible priors.

Definition of Equilibrium. Given the price function (11) and the aggregate order flows (15), we define a sequential trading (partial) equilibrium among arbitrageurs who have uncertain fat-tail priors $\mathcal{LG}(\alpha, \tilde{\xi})$ with $\tilde{\xi} \in [\xi_L, \xi_H]$. The equilibrium is described by a matrix of their strategies $\mathbf{Z}$ such that for all $n = 1, ..., N$ and any alternative strategy profile $\mathbf{Z}'$ that differs from $\mathbf{Z}$ only in the n-th entry $\mathbf{Z}'_n = \langle Z'_{1,n}, Z'_{2,n} \rangle$, the strategy profile $\mathbf{Z}$ yields a utility level (i.e., the minimum expected profit over all possible priors) no less than $\mathbf{Z}'$, and $Z_{2,n}$ yields a utility level in the second period no less than that produced by any single deviation $Z'_{2,n}$:

$$\min_{\xi} \mathrm{E}^{\mathcal{A}}[\tilde{\pi}_{z,n}(\mathbf{Z})|\xi_v = \xi] \quad \geq \quad \min_{\xi} \mathrm{E}^{\mathcal{A}}[\tilde{\pi}_{z,n}(\mathbf{Z}')|\xi_v = \xi], \tag{16}$$

$$\min_{\xi} \mathrm{E}^{\mathcal{A}}[(\tilde{v} - \tilde{p}_2(\cdot, Z_{2,n}))Z_{2,n}|\tilde{p}_1, \xi_v = \xi] \quad \geq \quad \min_{\xi} \mathrm{E}^{\mathcal{A}}[(\tilde{v} - \tilde{p}_2(\cdot, Z'_{2,n}))Z'_{2,n}|\tilde{p}_1, \xi_v = \xi]. \tag{17}$$

[10] This agrees with the assumption of Malikov and Pasquariello (2021) that other market participants (e.g., market makers) do not anticipate the trading of quants or strategically respond to it.

[11] For the mixture distribution (9), the kurtosis of $\tilde{v}$ is $3 + 3(4+\alpha)(1-\alpha)/(2-\alpha)^2$ which only depends on α; see Haas et al. (2006). Section 6.1 discusses the case when traders are uncertain about both α and ξ_v.

4 Equilibrium Strategies

4.1 Optimal strategy under a fixed prior

Arbitrageurs' trading strategies are driven by their estimates of the extent to which the asset has been mispriced. Conditional on past prices, arbitrageurs' expectations of $\tilde{v}$ depend on their fat-tail priors $\mathcal{LG}(\alpha, \tilde{\xi})$ with $\alpha \in (0, 1]$. When there is no model risk about ξ at all, arbitrageurs become the standard *subjective expected utility* optimizers, under a fixed fat-tail prior $\mathcal{LG}(\alpha, \xi)$. This leads to the benchmark trading strategy in this paper.

Theorem 1. *Suppose arbitrageurs have the same fixed Laplacian-Gaussian prior $\mathcal{LG}(\alpha, \xi)$ where $\alpha \in [0, 1]$ and $\xi \in (0, \infty)$. There exists a symmetric equilibrium where they choose to watch the market without any betting at $t = 1$, i.e., $Z^o_{1,n} = 0$, and their optimal strategy at $t = 2$ is proportional to their posterior mean estimate $\widehat{\theta}$ of the pricing error $\tilde{\theta} := \tilde{v} - p_1$, i.e.,*

$$Z^o_{2,n}(p_1; \alpha, \xi) = \frac{1 - \beta_2\lambda_2}{\lambda_2(N + 1)} \cdot \widehat{\theta}(p_1; \alpha, \xi) = \frac{1 - \beta_2\lambda_2}{\lambda_2(N + 1)} \cdot [\widehat{v}(p_1; \alpha, \xi) - p_1]. \tag{18}$$

Here, $\widehat{v} := \mathrm{E}^{\mathcal{A}}[\tilde{v}|p_1, p_0]$ is the posterior mean estimate of $\tilde{v}$ under arbitrageurs' prior $\mathcal{LG}(\alpha, \xi)$ and conditional on the price history. Given the linear pricing rule (11), $\widehat{v}$ only depends on the order flow y_1 which drives the price change $p_1 - p_0$. If we measure order flows in units of the noise volatility, $y := y_1/\sigma_u$, and define $\kappa := \sigma_u/(\beta_1\xi)$, then the posterior mean of $\tilde{v}$ is

$$\widehat{v}(y) = \frac{\alpha\sigma_u(y - \kappa)\mathrm{erfc}\left(\frac{\kappa - y}{\sqrt{2}}\right)}{\beta_1\mathrm{erfc}\left(\frac{\kappa - y}{\sqrt{2}}\right) + \beta_1 e^{2\kappa y}\mathrm{erfc}\left(\frac{\kappa + y}{\sqrt{2}}\right)} + \frac{\alpha\sigma_u(y + \kappa)\mathrm{erfc}\left(\frac{\kappa + y}{\sqrt{2}}\right)}{\beta_1\mathrm{erfc}\left(\frac{\kappa + y}{\sqrt{2}}\right) + \beta_1 e^{-2\kappa y}\mathrm{erfc}\left(\frac{\kappa - y}{\sqrt{2}}\right)} + (1 - \alpha)\lambda_1\sigma_u y. \tag{19}$$

This function strictly increases with the order flow and has two shape parameters α and $\kappa(\xi)$. Asymptotically, $\widehat{v}$ becomes linear in the total order flow realized at $t = 1$:

$$\widehat{v} \to \alpha[y_1 - \mathrm{sign}(y_1)\kappa\sigma_u]/\beta_1 + (1 - \alpha)\lambda_1 y_1, \quad \text{as} \quad |y_1| \to \infty. \tag{20}$$

Proof. See Appendix A.2. □

Given Eq. (19) and $p_1 = \lambda_1 y_1$, one can factor out α from the strategy function (18):

$$Z^o_{2,n}(y_1; \alpha, \xi) = \alpha Z^o_{2,n}(y_1; \alpha = 1, \xi) = \frac{\alpha(1 - \beta_2\lambda_2)[\widehat{v}(y_1; \alpha = 1, \xi) - \lambda_1 y_1]}{\lambda_2(N + 1)}. \tag{21}$$

The trading intensity is exactly proportional to the frequency of fat-tail shocks. Therefore, the key mathematical properties of $Z^o_{2,n}(y_1;\alpha,\xi)$ is independent of α. This scaling rule is an important feature of all Bayesian rational strategies discussed in this paper.

Because arbitrageurs' prior $\mathcal{LG}(\alpha,\xi)$ is a symmetric distribution, they tend to postpone trading until they can distinguish the direction of signals (as the posterior becomes skewed). When solving the equilibrium, we conjecture first and verify later that arbitrageurs do not trade at $t=1$. This is confirmed by Theorem 1 and explains our choice of a two-period setup. The no-trade conjecture holds if the market is not too crowded for arbitrageurs; otherwise, it can be profitable for a trader to trade at $t=1$ in the hope that other traders get misled.

Proposition 1. *For $\xi>0$ and $\alpha\in[0,1]$, the symmetric equilibrium in Theorem 1 exists if the following market condition holds:*

$$1+\frac{\alpha(1-\beta_1\lambda_1)}{\beta_1\lambda_1}\cdot\frac{N-1}{N+1}<\frac{2\sqrt{\lambda_2/\lambda_1}}{1-\beta_2\lambda_2}. \tag{22}$$

Proof. See Appendix A.3. □

The condition (22) turns out to be a universal condition for all the equilibria analyzed in this paper. Unless otherwise specified, this (sufficient) condition is assumed to hold hereafter. This precludes unilateral deviations from the no-trade strategy at $t=1$. According to (22), the equilibrium can accommodate an infinite number of arbitrageurs when the trading environment satisfies: $\alpha+(1-\alpha)\beta_1\lambda_1<2\beta_1\sqrt{\lambda_1\lambda_2}/(1-\beta_2\lambda_2)$.[12]

Corollary 4.1. *The strategy function $Z^o_{2,n}(\tilde{p}_1;\alpha,\xi)=Z^o_{2,n}(\tilde{y}_1;\alpha,\xi)$ defined by Eq. (18) and Eq. (19) has the following properties:*
(a) It is a smooth, odd function of y_1 and symmetric about the origin.
(b) It is a convex function for $y_1\geq 0$ and a concave function for $y_1\leq 0$.
(c) It has two slant asymptotes for any $\xi\in(0,\infty)$ which take the general form below,

$$Z^\infty(y_1;\alpha,\xi)=\frac{\alpha(1-\beta_1\lambda_1)(1-\beta_2\lambda_2)}{\beta_1\lambda_2}\cdot\frac{y_1-\text{sign}(y_1)K(\xi)}{N+1}, \tag{23}$$

where the horizontal intercept $K(\xi)$ is inversely related to ξ but independent of α,

$$K(\xi)=\frac{\kappa(\xi)\sigma_u}{1-\beta_1\lambda_1}=\frac{\sigma_u^2\xi^{-1}}{\beta_1(1-\beta_1\lambda_1)}. \tag{24}$$

[12] For example, if λ_t and β_t are determined by the two-period Kyle model (see Proposition 1 of Huddart, Hughes, and Levine (2001)), one can verify that the equilibrium condition (22) holds even when $N\to\infty$.

(d) There exists a unique critical value, $\xi_c > 0$, endogenously determined by the equation:

$$1+\left(\frac{\sigma_u}{\beta_1\xi_c}\right)^2-\sqrt{\frac{2}{\pi}}\frac{\sigma_u}{\beta_1\xi_c}\frac{\exp(-\sigma_u^2/(2\beta_1^2\xi_c^2))}{\operatorname{erfc}(\sigma_u/(\sqrt{2}\beta_1\xi_c))}=\beta_1\lambda_1. \tag{25}$$

For $\xi \geq \xi_c$, $Z^o_{2,n}(y_1;\alpha,\xi)$ is an increasing function of y_1 and it has only one root at $y_1 = 0$. For $\xi < \xi_c$, $Z^o_{2,n}(y_1;\alpha,\xi)$ is a non-monotonic function of y_1 and it has three different roots. Note that ξ_c is independent of α and ξ_v; it depends on the ratio σ_u/β_1 and the product $\beta_1\lambda_1$.

Proof. See Appendix A.4. □

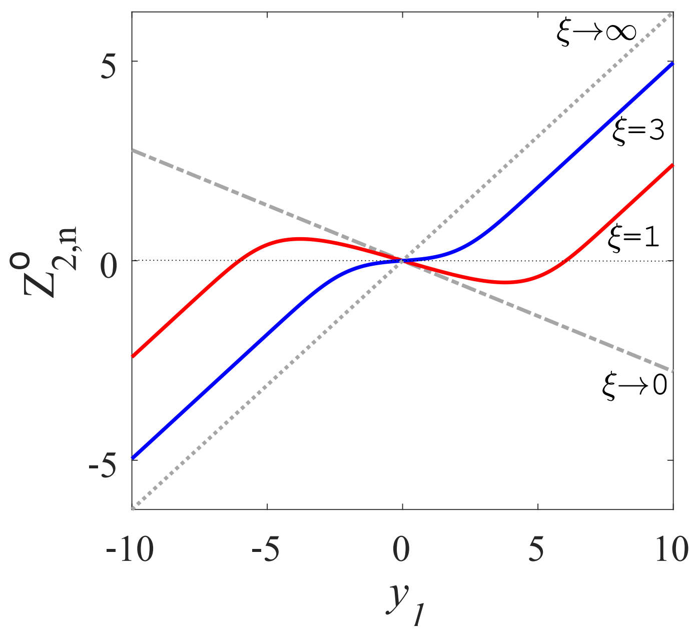

Figure 1. The subjective optimal strategy $Z^o_{2,n}(y_1;\alpha,\xi)$ in Eq. (18) for various values of ξ.

The main properties listed in Corollary 4.1 are reflected in Figure 1 where we plot the optimal strategy (18) for different values of ξ. An arbitrageur with the extreme prior $\xi \to 0$ believes that the asset value is unchanged (i.e., $\tilde{v} = 0$). This trader will attribute all the order flow y_1 to noise trading and trade against any price movements. In contrast, an arbitrageur with the extreme prior $\xi \to \infty$ believes that the first-period order flow was dominated by informed trading and will chase the price trend straightly. When $\xi < \xi_c$, this strategy is contrarian ("leaning against the wind") for small order flows but momentum for sufficiently large ones. Fatter tails in the prior (i.e., larger ξ) lead to a lower threshold for switching to momentum trading. When $\xi \geq \xi_c$, the strategy is always trend-following.

4.2 Robust strategy under uncertain priors

Model risks can be prominent in a fat-tailed trading environment where market meltdown may be triggered if some big trader has applied a wrong model. Institutional traders are often required to test their strategies across alternative scenarios. This pressure can motivate them to adopt strategies that sacrifice some optimality for robustness.

In our setup, how would arbitrageurs trade given their uncertain fat-tail prior $\mathcal{LG}(\alpha, \tilde{\xi})$? Figure 1 shows that they will face ambiguity about the profitable trading direction when they observe small and medium order flows. They may want to buy this asset under a high prior (blue line) but sell it under a low prior (red line). In case they use the wrong prior, they may trade on the wrong side and expose themselves to adverse fat-tail shocks. For robustness, they should not trade until there is little ambiguity about the trading direction.

Corollary 4.2. *When the range of uncertain Laplace prior satisfies $\xi_L < \xi_c \leq \xi_H$,*[13] *there is an equilibrium where arbitrageurs idle at $t = 1$ and their pure max-min strategy at $t = 2$ is $Z^o_{2,n}(y_1; \alpha, \xi_L)\mathbf{1}_{|y_1|>K_L}$, where $K_L(\xi_L)$ is the positive root of the equation: $Z^o_{2,n}(y_1; \alpha, \xi_L) = 0$. The solution of K_L is independent of α, according to the scaling property (21).*

Proof. See Appendix A.5. □

The nonlinear strategy $Z^o_{2,n}(y_1; \alpha, \xi_L)\mathbf{1}_{|y_1|>K_L}$ in Corollary 4.2 can be extremely biased because it is only determined by the lowest prior ξ_L and its threshold K_L can be arbitrarily large when ξ_L becomes arbitrarily small. This strategy may sacrifice too much optimality for robustness, making the equilibrium in Corollary 4.2 unappealing in reality.

Arbitrageurs may ponder a different equilibrium where they are not attached to the lowest prior ξ_L. This *debiased* equilibrium may be more attractive as it may balance robustness and optimality. In any possible equilibrium, the asymptotes of arbitrageurs' strategies at $t = 2$ cannot slope differently from Eq. (23); otherwise, for sufficiently large y_1, they would trade either more than the most optimistic strategy (following ξ_H) or less than the most pessimistic strategy (following ξ_L). Thus, in a symmetric and debiased equilibrium, arbitrageurs' strategies must have the same asymptotes in the form of (23), denoted $Z^\infty(y_1; \alpha, \xi_w)$ for some $\xi_w > \xi_L$. Arbitrageurs can average across multiple priors to achieve $\xi_w > \xi_L$. There is a formal argument for this debiasing mechanism: observing a sufficiently large order flow y_1 may convince arbitrageurs that y_1 contains a strong fat-tail signal, which may ease their concerns about trading on the wrong side and make them indifferent to model risks. Suppose they become ambiguity-neutral as $y_1 \to \pm\infty$. Then the asymptotes of their strategies must coincide

[13] If $\xi_c \leq \xi_L < \xi_H$, the max-min strategy is simply $Z^o_{2,n}(y_1; \alpha, \xi_L)$. If $\xi_L < \xi_H < \xi_c$, the max-min strategy becomes complicated and lacks empirical relevance. See Figure 11 and Appendix A.5.

with the asymptotes of their risk-neutral strategies which, after averaging across all possible priors, is given by $\mathrm{E}[Z^o_{2,n}(y_1;\alpha,\tilde{\xi})]$. For an arbitrary weight function $w\colon [\xi_L,\xi_H] \to [0,1)$ with $\int_{\xi_L}^{\xi_H} w(\xi)d\xi = 1$, Eq. (23) and Eq. (24) imply that ξ_w is the weighted harmonic mean of $\tilde{\xi}$:

$$\xi_w := \left(\int_{\xi_L}^{\xi_H} \xi^{-1} w(\xi) d\xi\right)^{-1}. \tag{26}$$

Depending on the weight function $w(\xi)$, the weighted average ξ_w can take any value over the interval (ξ_L,ξ_H) and correspondingly, the asymptotes $Z^\infty(y_1;\alpha,\xi_w)$ can shift freely as well.

Theorem 2. *There exists a symmetric, debiased equilibrium if the following conditions hold:*

(C1) arbitrageurs' admissible strategies converge to $Z^\infty(y_1;\alpha,\xi_w)$ with $K(\xi_w) < K_L$;

(C2) arbitrageurs' admissible strategies are convex for $y_1 \geq 0$ and concave for $y_1 \leq 0$;

(C3) the range of their uncertain priors satisfies $\xi_L < \xi_c \leq \xi_H$, where ξ_c solves Eq. (25).

In this equilibrium, arbitrageurs watch the market without any trading at $t = 1$. Their robust trading strategy at $t = 2$ is a soft-thresholding function of the order flow y_1 realized at $t = 1$:

$$\begin{aligned} Z_{2,n}(y_1;\alpha,K(\xi_w)) &= Z^\infty(y_1;\alpha,\xi_w)\mathbf{1}_{|y_1|>K(\xi_w)} = \frac{\alpha(1-\beta_1\lambda_1)(1-\beta_2\lambda_2)}{\beta_1\lambda_2(N+1)}\mathcal{S}(y_1;K(\xi_w)) \\ &= \frac{\alpha(1-\beta_1\lambda_1)(1-\beta_2\lambda_2)}{\beta_1\lambda_2(N+1)}\left[y_1 - \mathrm{sign}(y_1)K(\xi_w)\right]\mathbf{1}_{|y_1|>K(\xi_w)}. \end{aligned} \tag{27}$$

Proof. See Appendix A.6. The soft-thresholding operator $\mathcal{S}(y_1;K)$ is defined in Eq. (4). □

As discussed earlier, the first condition (C1) is endogenously required for the existence of a symmetric and debiased equilibrium. (C1) can be microfounded, for example, by assuming arbitrageurs are asymptotically indifferent to the model risk. If (C1) is absent, the condition (C2) alone will have no effect because this model economy will admit the same equilibrium stated in Corollary 4.2 which only requires the condition (C3). Only when (C1) holds, the condition (C2) can play a meaningful role in *regularizing* the max-min problem under (C1).[14] (C2) preserves the curvature property of optimal strategies (Corollary 4.1(b) and Figure 1). In practice, (C2) could be a desirable constraint to avoid overfitting with highly complex models (Kelly, Malamud, and Zhou (2022), Didisheim, Ke, Kelly, and Malamud (2023)).

Given (C1) and (C2), the admissible strategies must be enclosed by $Z^o_{2,n}(y_1;\alpha,\xi\to\infty)$, $Z^o_{2,n}(y_1;\alpha,\xi\to 0)$, and $Z^\infty(y_1;\alpha,\xi_w)$; see the shaded area in Figure 2. Any strategy that

[14]Without the curvature condition (C2), for any strategy that satisfies (C1), arbitrageurs can always find some deviations that trade more conservatively than this strategy. Such deviations are preferred under the max-min criterion but fail to support an equilibrium. Note that the deviations are possble because of the gap between $Z^\infty(y_1;\alpha,\xi_w)$ and $Z^\infty(y_1;\alpha,\xi_L)$, which is a direct implication of (C1).

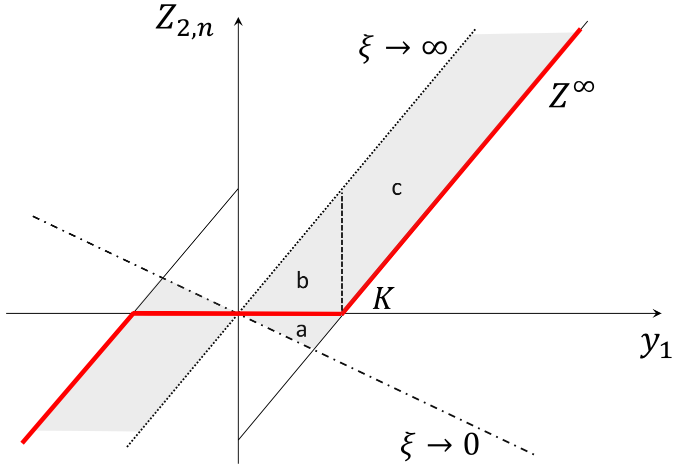

Figure 2. The robust trading strategy $Z_{2,n}(y_1; \alpha, K)$ under model risk as in Theorem 2.

goes out of the shaded area is either irrational or violating (C1) or (C2). We focus on the positive domain of y_1 and divide the shaded area into three regions. First, for $y_1 \in [0, K(\xi_w)]$, each arbitrageur will not sell against y_1, considering that she may lose money on average by doing so if the highest prior ξ_H is true. This rules out any decision point inside the triangle "a". Similarly, each arbitrageur will not purchase this asset, considering that she may lose money if the lowest prior ξ_L is true. This rules out any decision point inside the triangle "b". By the max-min criterion, each arbitrageur will not trade for $y_1 \in [0, K(\xi_w)]$. Next, for any $y_1 \in (K(\xi_w), \infty)$, each arbitrageur will not trade more than the amount of $Z^\infty(y_1; \alpha, \xi_w)$, because she understands that in the worst case she could lose more money by trading more. This argument rules out any decision point inside the open region "c". By symmetry, their equilibrium strategy is the red line in Figure 2, which is exactly characterized by Eq. (27).

This robust strategy $Z_{2,n}(y_1; \alpha, K(\xi_w))$ is a soft-thresholding function of the total order y_1. Its slope $\frac{\alpha(1-\beta_1\lambda_1)(1-\beta_2\lambda_2)}{\beta_1\lambda_2(N+1)}$ is independent of ξ_w, whereas its threshold $K(\xi_w) = \frac{\sigma_u^2 \xi_w^{-1}}{\beta_1(1-\beta_1\lambda_1)}$ is independent of α. The no-trade zone $[-K(\xi_w), K(\xi_w)]$ indicates infrequent trading activities at $t = 2$. Because arbitrageurs do not trade at $t = 1$ and only trade occasionally at $t = 2$, a range of small pricing errors can survive in this market. *Ex post*, an econometrician may find pervasive anomalies after analyzing the trading data in this model economy. She may question the rationality or capability of arbitrageurs. *Ex ante*, arbitrageurs have rationally assessed all the possible states and they let go many vague mispricings for robustness. There are no exogenous frictions that limit their trading ability, except the price impact costs, λ_1 and λ_2, which can be endogenized in a standard Kyle-type model. The major friction in our setup is the fat-tailed model risk which discourages cautious arbitrageurs from eradicating all possible pricing errors.

5 Main Results

5.1 LASSO as a Bayesian rational strategy

A strategy is a complete plan of action that a player chooses to take contingent on what circumstances might arise. The decision can depend on how the player predicts or estimates economic variables. When the prior is known and fixed, a Bayesian rational player should use the posterior mean estimate and optimize her expected utility. When the prior is unknown or uncertain, the player may have two options: either sacrifice rationality by altering the estimation method or preserve rationality by optimizing an alternative utility function. The same strategy may be derived from different approaches. In such cases, we cannot distinguish whether the player is taking the heuristic or the rational approach in decision making. Thus, we may have two theories *observationally equivalent* in that all of their empirical implications are identical but they express distinct theoretical perspectives.

Given the definition of LASSO estimates (Eq. (1) in Section 2), we call a plan of action a LASSO strategy if it can be effectively implemented with the LASSO estimate of variable(s). This definition takes no stance on the player's motives in developing the strategy and thus admits the possibility of *observational equivalence*. The strategy may be derived from either heuristic or rational arguments. It is usually easy and tempting to take the view that agents adhere to some heuristic approach.

Proposition 2. *Consider an otherwise identical environment except that arbitrageurs use the LASSO as a heuristic to estimate mispricings. Suppose their strategy space consists of all the linear functions* $Z_{2,n}^{\text{lasso}}(y_1) = a + b \cdot \widehat{\theta}^{\text{lasso}}(y_1)$*, where* $\widehat{\theta}^{\text{lasso}}$ *is a LASSO estimate of* $\tilde{\theta}$ *with an arbitrary threshold. If arbitrageurs are strategic and asymptotically rational, they will use a LASSO strategy observationally equivalent to the robust strategy in Theorem 2.*

Proof. By symmetry, we have $Z_{2,n}^{\text{lasso}}(y_1 = 0) = 0$ which requires $a = 0$. Clearly, the only strategy $Z_{2,n}^{\text{lasso}}(y_1) = b \cdot \widehat{\theta}^{\text{lasso}}(y_1)$ that satisfies the asymptotic rationality is the one that coincides with the REE asymptotes $Z^{\infty}(y_1; \alpha, \xi_w)$. This determines $Z_{2,n}^{\text{lasso}}(y_1) = Z_{2,n}(y_1; \alpha, K(\xi_w))$. □

Given this observation, we can call the robust strategy $Z_{2,n}(y_1; \alpha, K)$ an asymptotically rational LASSO strategy. In contrast to the heuristic argument in Proposition 2, we have presented a system of economic arguments to defend that this LASSO strategy can arise from a Bayesian rational equilibrium (Theorem 2). It remains to show how this strategy can be effectively implemented with the LASSO estimate of trading signals.

Theorem 3. *For any $\alpha \in (0,1]$, the robust strategy (27) in Theorem 2 is a LASSO strategy:*

$$Z_{2,n}(y_1;\alpha,K(\xi_w)) = \frac{\alpha(1-\beta_2\lambda_2)(\widehat{v}^{\text{lasso}} - \lambda_1 y_1)\mathbf{1}_{|y_1|>K(\xi_w)}}{\lambda_2(N+1)} = \frac{\alpha(1-\beta_2\lambda_2)}{\lambda_2(N+1)} \cdot \widehat{\theta}^{\text{lasso}}. \tag{28}$$

where the LASSO estimate $\widehat{v}^{\text{lasso}}$ is the solution to the LASSO objective function:

$$\widehat{v}^{\text{lasso}}(y_1;\xi_w) := \arg\min_v \left\{ \frac{1}{2}|y_1 - \beta_1 v|^2 + \frac{\sigma_u^2}{\xi_w}|v| \right\} = \frac{1}{\beta_1}\mathcal{S}(y_1;\kappa(\xi_w)\sigma_u) \tag{29}$$

and the LASSO estimate $\widehat{\theta}^{\text{lasso}}$ of the mispricing $\tilde{\theta} = \tilde{v} - p_1$ solves the LASSO objective:

$$\widehat{\theta}^{\text{lasso}}(y_1;\xi_w) := \arg\min_\theta \left\{ \frac{1}{2}\left| y_1 - \frac{\beta_1\theta}{1-\beta_1\lambda_1}\right|^2 + \frac{\sigma_u^2|\theta|}{(1-\beta_1\lambda_1)^2\xi_w} \right\} = \frac{1-\beta_1\lambda_1}{\beta_1}\mathcal{S}(y_1;K(\xi_w)). \tag{30}$$

The two LASSO thresholds satisfy $\kappa\sigma_u < K = \kappa\sigma_u/(1-\beta_1\lambda_1)$ since $\beta_1\lambda_1 \in (0,1)$.

Proof. See Appendix A.7. □

Theorem 3 confirms that the robust strategy (27) is always a LASSO strategy. This holds under an uncertain fat-tail prior $\mathcal{LG}(\alpha,\tilde{\xi})$ for any $\alpha > 0$. Theorem 2 provides the system of economic arguments for developing this robust strategy. Therefore, under fairly general conditions, we show that the LASSO can be a Bayesian rational strategy for agents who face prior uncertainty about the fat-tail shocks and who optimize the max-min expected utility.

This provides an economic rationale for us to use the LASSO. In contrast to Tibshirani's interpretation, our interpretation does not invoke the heuristic MAP estimate nor require a pure and fixed Laplace prior. In our setup, agents (i.e., arbitrageurs) are *Bayesian rational* because they use the posterior mean estimate to assess all the possible states and scenarios; the LASSO strategy they choose is also *sequentially rational* because each of them applies dynamic programming to solve a two-period objective function; the LASSO strategy also qualifies as an equilibrium strategy because each agent strategically considers the best responses of other agents. The prior is a general mixture distribution (9) with a raw kurtosis from 3 to 6.125, depending on the mixture weight α. Our theory holds for any $\alpha \in (0,1]$. This covers a wide range of fat-tailedness of the prior distribution.

What if arbitrageurs simply use the MAP estimate in this economic environment? This is related to the MAP-based interpretation (Tibshirani (1996)), which cannot produce the LASSO if the prior is not exactly Laplacian or if the prior is uncertain. In our setup, the MAP-based trading rule not only violates Bayesian rationality but also lacks sequential rationality. Being a non-equilibrium strategy, it can incur large losses.

Proposition 3. *When $\alpha = 1$, the robust LASSO strategy (27) is observationally equivalent to a heuristic, feedback trading strategy when arbitrageurs all adopt the MAP learning rule to estimate $\tilde{v}$ under a pure and fixed Laplace prior $\mathcal{L}(0, \xi = \xi_w)$. This can be written as*

$$Z_{2,n}(y_1; \alpha = 1, K(\xi_w)) = Z_{2,n}^{\text{map}}(y_1; \alpha = 1, \xi = \xi_w) = \frac{(1 - \beta_2\lambda_2)(\widehat{v}^{\text{map}} - \lambda_1 y_1)\mathbf{1}_{|y_1| > K(\xi_w)}}{\lambda_2(N+1)}, \quad (31)$$

where the MAP estimate $\widehat{v}^{\text{map}}$ coincides with the LASSO estimate $\widehat{v}^{\text{lasso}}$ defined by Eq. (29):

$$\widehat{v}^{\text{map}}(y_1; \alpha = 1, \xi = \xi_w) = \arg\max_v \frac{f(y_1|v)f(v; \alpha = 1, \xi_w)}{f(y_1)} = \frac{\mathcal{S}(y_1; \kappa(\xi_w)\sigma_u)}{\beta_1} = \widehat{v}^{\text{lasso}}(y_1; \xi_w). (32)$$

When $\alpha \neq 1$, the above coincidence fails. The MAP-based trading rule with the fixed prior $\mathcal{LG}(\alpha, \xi_w)$ differs from the robust LASSO strategy with the uncertain prior $\mathcal{LG}(\alpha, \tilde{\xi})$:

$$Z_{2,n}(y_1; \alpha, K(\xi_w)) \neq Z_{2,n}^{\text{map}}(y_1; \alpha, \xi_w), \qquad \textit{for any } \alpha \in (0, 1). \quad (33)$$

Proof. See Appendix A.8. □

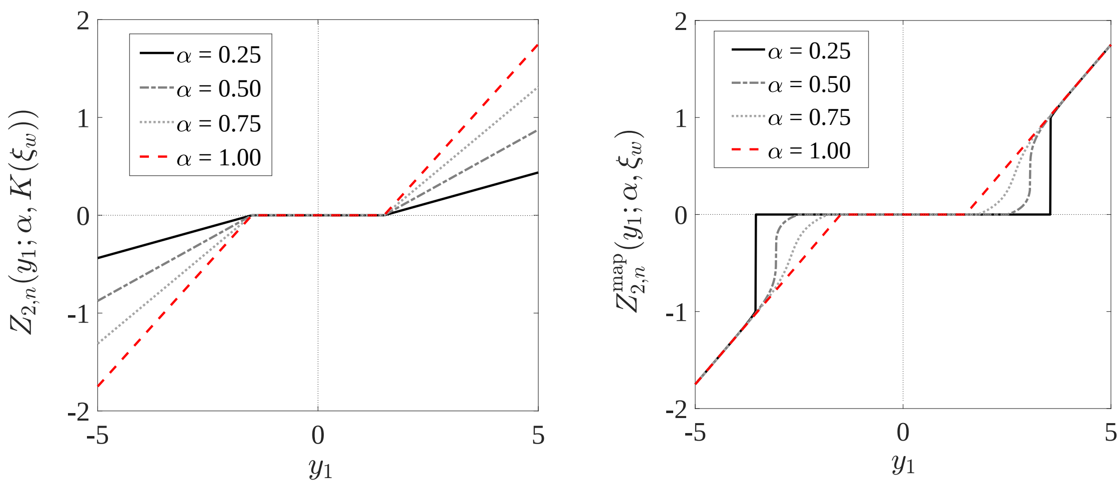

Figure 3. The robust LASSO strategy $Z_{2,n}(y_1; \alpha, K(\xi_w))$ and the MAP-based trading rule $Z_{2,n}^{\text{map}}(y_1; \alpha, \xi_w)$ for different values of α.

The MAP-based trading rule can produce a LASSO strategy only when traders have the pure fixed Laplace prior $\mathcal{LG}(\alpha = 1, \xi)$. For $\alpha \in (0, 1)$, we numerically solve the MAP strategy $Z_{2,n}^{\text{map}}$ by computing the posterior and finding its mode $\widehat{v}^{\text{map}}$. Figure 3 compare the profiles of two strategies. As a feature of Bayesian rationality which averages across all possibilities, the robust LASSO strategy satisfies the scaling $Z_{2,n}(y_1; \alpha, K(\xi_w)) = \alpha Z_{2,n}(y_1; \alpha = 1, K(\xi_w))$ and its threshold $K(\xi_w)$ is independent of α. In contrast, the MAP strategy lacks this scaling: its asymptotes are independent of α while its threshold increases with α.[15] The MAP strategy becomes discontinuous when α is sufficiently small. It has a wider inaction region but, once triggered, tends to respond at the maximum intensity. This *all-or-none* response just relies on a heuristic threshold (resulted from the MAP estimate) to decide whether the noisy order flow y_1 contains a Laplacian or Gaussian signal. When y_1 exceeds the threshold, the MAP strategy treats y_1 as if it contains the Laplacian signal for sure and thus believes $s = 1$ in Eq. (10). This is an example of data classification widely used in machine learning.[16] However, as demonstrated here, simple classifications may violate Bayesian rationality.

As a consequence of its overreaction for $\alpha \in (0, 1)$, the trading rule $Z^{\text{map}}_{2,n}(y_1)$ cannot meet sequential rationality. Implicitly, the MAP strategy assumes that no one would trade in the first period. In fact, a trader may profitably manipulate the price at $t = 1$, expecting that other feedback traders misinterpret the price disturbance as a signal and overreact to her trading. Therefore, the MAP-based trading rule is not an equilibrium strategy. It is not derived by rational agents who dynamically optimize their utility functions.

For the robust LASSO strategy $Z_{2,n}(y_1; \alpha, K)$, its inaction zone $[-K, K]$ does not play the role of signal classification but shrinks a range of ambiguous estimates to zero, consistent with the objective of robust optimization. It is also remarkable that all the Bayesian rational strategies solved in Section 4 are convex for $y_1 > 0$ and concave for $y_1 > 0$. Figure 3 shows that the MAP strategy violates this rational curvature property for $\alpha \in (0, 1)$.

Rational agents should not use the MAP-based alternatives. The MAP strategy may arise when some traders have been educated to apply the MAP method. It is unclear to what extend this method has been adopted by financial practitioners. The answer may well depend on their educational backgrounds. For example, a quant who was trained in the field of image analysis (e.g., Greig, Porteous, and Seheult (1989)) or speech recognition (e.g., Lim and Oppenheim (1979)) but lacks systematic training in finance or economics may apply the MAP method. When a substantial fraction of quants have used it in their trading strategies, the market may have systemic risk and disasters like the quant meltdown in 2007 may not be uncommon. As shown in Figure 4, the Bayesian-rational LASSO strategy strongly dominates the heuristic MAP-based strategy in terms of the total expected profits. This dominance holds for arbitrary values of N and ξ_w. The MAP-based strategy tends to incur significant losses when traders overestimate the scale of fat tails (i.e., $\xi_w > \xi_v$). The numerical example is suggestive. It shows the importance of bridging the gap between neoclassical economics and machine learning.

---

[15]We can show that $\lim_{y_1 \to \pm\infty} \frac{1}{y_1} Z^{\text{map}}_{2,n}(y_1; \alpha, \xi) = \frac{(1-\beta_1\lambda_1)(1-\beta_2\lambda_2)}{\beta_1\lambda_2(N+1)}$, which is independent of α.

[16]See, for example, the *naive Bayes classifier* (Domingos and Pazzani (1997)).

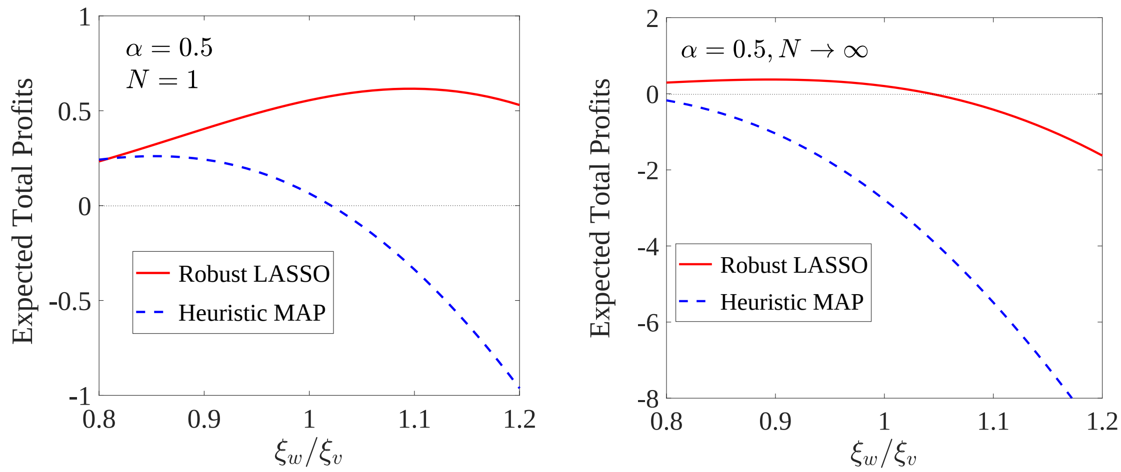

Figure 4. The expected total profits when arbitrageurs all follow the same robust LASSO strategy $Z_{2,n}(y_1; \alpha, K(\xi_w))$ versus those when they follow the MAP-based feedback trading rule $Z_{2,n}^{\text{map}}(y_1; \alpha, \xi_w)$ under the fixed prior $\mathcal{LG}(\alpha, \xi = \xi_w)$. We examine the case of $\alpha = 0.5$ and a wide range of values of ξ_w relative to the true prior value of ξ_v. The left panel reports the monopolistic result $N = 1$. The right panel is about the "competitive" limit $N \to \infty$.

5.2 Limits to arbitrage due to a "cartel" effect

Now we discuss the implications of the robust LASSO strategy in our model economy. The literature on limits of arbitrage (as reviewed by Gromb and Vayanos (2010)) has documented various market frictions which have a common feature as to limit arbitrageurs' *ability* to trade. Free from such frictions, our model is well suited to demonstrate a mechanism which only affects arbitrageurs' *willingness* to trade. In our setup, all the traders are risk neutral. They face no financial or trading constraints, except the price impact costs.[17] This setup highlights *model risk management* as the key friction. With uncertain fat-tail priors, this friction can lead to a robust strategy with a wide no-trade zone. There are two channels for this strategy to cause limited arbitrage. One is arbitrageurs' idleness in both the watching period ($t = 1$) and the speculation period ($t = 2$). The other channel is a subtle "cartel" effect due to arbitrageurs' conservative trading at $t = 2$.

To illustrate the first channel, we compare the unbiased robust strategy $Z_{2,n}(y_1; \alpha, K(\xi_v))$ (corresponding to the case of $\xi_w = \xi_v$) with the perfectly optimal strategy $Z_{2,n}^o(y_1; \alpha, \xi_v)$ (corresponding to the case that traders know the true prior ξ_v). By Theorem 3, the strategy $Z_{2,n}$ equivalently implements the LASSO estimate $\widehat{v}^{\text{lasso}}$,[18] whereas the optimal strategy $Z_{2,n}^o$

[17]In standard Kyle-type models, the price impact costs are not at all detrimental to market efficiency.
[18]By Proposition 3, $\widehat{v}^{\text{lasso}}$ coincides with the posterior mode estimate $\widehat{v}^{\text{map}}$ only when $\alpha = 1$.

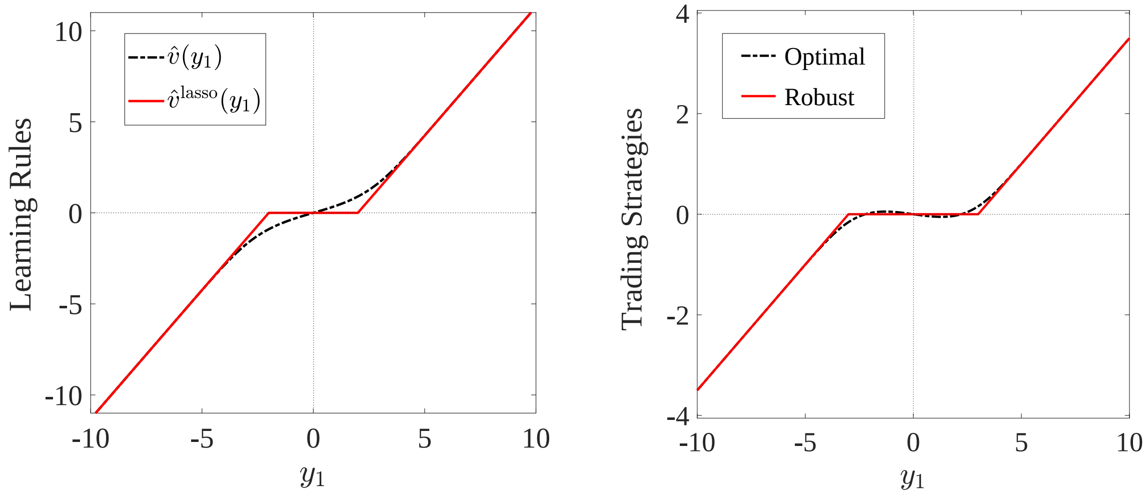

Figure 5. Left: the posterior mean estimate $\widehat{v}(y_1)$ and the LASSO estimate $\widehat{v}^{\text{lasso}}(y_1)$ under a common prior $\mathcal{LG}(\alpha, \xi_v)$. Right: the perfectly optimal strategy $Z^o_{2,n}(y_1; \alpha, \xi_v)$ and the asymptotically unbiased robust strategy $Z_{2,n}(y_1; \alpha, K(\xi_v))$.

is directly driven by the posterior mean $\widehat{v}$. As the left panel of Figure 5 shows, the posterior mean estimate $\widehat{v}$ is a smooth nonlinear function, while the LASSO estimate $\widehat{v}^{\text{lasso}}$ is a soft-thresholding function which is zero for $y_1 \in [-\kappa\sigma_u, \kappa\sigma_u]$ and is linear outside that region. One can also see that $\widehat{v} > \widehat{v}^{\text{lasso}}$ when $y_1 > 0$ and $\widehat{v} < \widehat{v}^{\text{lasso}}$ when $y_1 < 0$. The right panel of Figure 5 compares these two strategies. Again, the optimal strategy $Z^o_{2,n}$ is a nonlinear function with asymptotic linearity, whereas the robust strategy $Z_{2,n}$ is a soft-thresholding function with an inaction zone $[-K, K]$. It can be verified that $Z_{2,n}(y_1; \alpha, K(\xi_v)) < Z^o_{2,n}(y_1; \alpha, \xi_v)$ for $|y_1| > K$. This follows from robust control which tends to produce conservative responses.

The robust LASSO strategy only reacts to large events and deliberately ignores small ones. This feature is similar to many phenomena in behavioral economics, including limited attention, status quo bias, anchoring and adjustment, among others; see Barberis and Thaler (2003) and Gabaix (2014). Despite such similarities, the LASSO strategy (27) is the rational choice by traders in our setting. They leave money on the table because the fat-tailed model risk discourages them from betting on directionally ambiguous pricing errors.

To elucidate the second channel, we need to compare the profitability of the robust strategy $Z_{2,n}(y_1; \alpha, K(\xi_w))$ with that of the optimal strategy $Z^o_{2,n}(y_1; \alpha, \xi_w)$.[19] The comparison is on a fair ground when these strategies can converge as $|y_1| \to \infty$. Figure 6 shows the total trading profits earned by all arbitrageurs when they follow the same strategy. The left panel

[19] This is a benchmark strategy as if traders ignore model risk and follow the "optimal" strategy using the averaged prior ξ_w. Another benchmark is the rational-expectations strategy, $\mathrm{E}_w[Z^o_{2,n}(y_1; \alpha, \tilde{\xi})]$, which is less tractable but extremely close to $Z^o_{2,n}(y_1; \alpha, \xi_w)$. We have $\mathrm{E}_w[Z^o_{2,n}(y_1; \alpha, \tilde{\xi})] \to Z^o_{2,n}(y_1; \alpha, \xi_w)$ as $|y_1| \to \infty$.

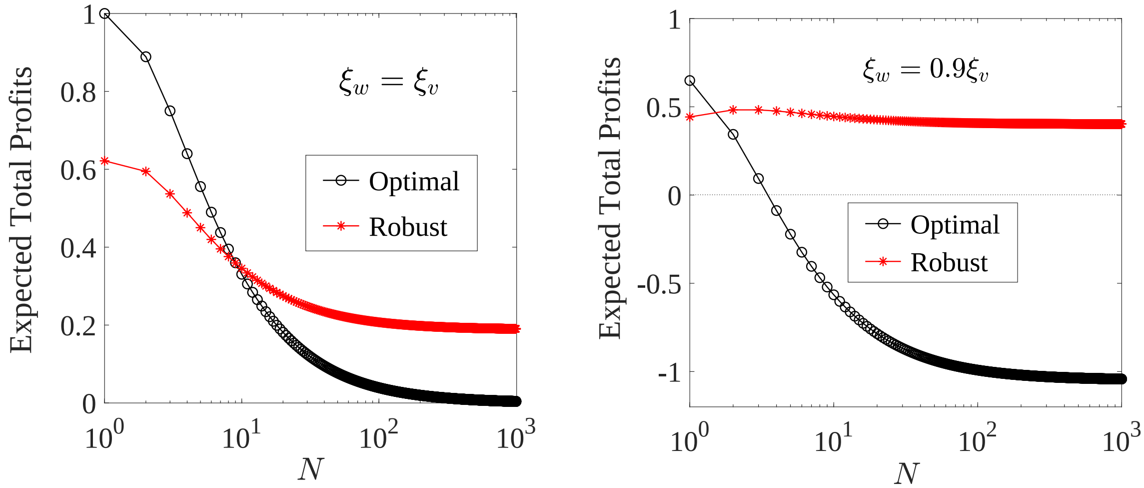

Figure 6. The expected total trading profits of all arbitrageurs when they follow the robust strategy $Z_{2,n}(y_1; \alpha, K(\xi_w))$ versus when they follow the optimal strategy $Z^o_{2,n}(y_1; \alpha, \xi_w)$.

shows the case of $\xi_w = \xi_v$, i.e., arbitrageurs are asymptotically unbiased. With the optimal strategy, they enjoy oligopoly profits for small values of N, while their total profits decay rapidly toward zero as N increases to infinity. In contrast, arbitrageurs' total profits decay slowly when they follow the robust LASSO strategy. For $N \leq 10$, this strategy allows them to capture the majority of what they would earn using the optimal strategy. For $N > 10$, their robust trading profits even surpass the optimal trading profits. In the limit $N \to \infty$, their total profit converges to a level at about 20% of the maximal monopoly profit (normalized to one). This result may be surprising given that the difference between these two strategies vanishes everywhere: $\lim_{N\to\infty} |Z_{2,n}(y_1; \alpha, K(\xi_w)) - Z^o_{2,n}(y_1; \alpha, \xi_w)| = 0$.

The right panel of Figure 6 shows the case when $\xi_w = 0.9\xi_v$. With the robust strategy, arbitrageurs' profits are significant and almost flat as N increases. In contrast, the optimal strategy leads to increasing losses as N increases to infinity. Its performance is much more sensitive to the competition. From extensive numerical experiments, we find that the robust strategy can often outperform the optimal strategy when they are both biased ($\xi_w \neq \xi_v$). Thus, the optimal strategy is unattractive under estimation bias and trading competition.

Theorem 4. *In the symmetric equilibrium of Theorem 1, the expectation of arbitrageurs' aggregate profit when they all follow the subjective optimal strategy $Z^o_{2,n}(y_1; \alpha, \xi_w)$ is*

$$\mathrm{E}\left[\sum_{n=1}^{N}(\tilde{v} - \tilde{p}_2)Z^o_{2,n}\right] = \frac{N(1-\beta_2\lambda_2)^2}{(N+1)\lambda_2}\mathrm{E}\left[\widehat{\theta}(\tilde{y}_1; \alpha, \xi_v)\widehat{\theta}(\tilde{y}_1; \alpha, \xi_w) - \frac{N}{N+1}\widehat{\theta}(\tilde{y}_1; \alpha, \xi_w)^2\right], \quad (34)$$

where $\widehat{\theta}(y_1;\alpha,\xi_w) := \mathrm{E}[\tilde{v}-\tilde{p}_1|y_1;\alpha,\xi_w] = \widehat{v}(y_1;\alpha,\xi_w) - \lambda_1 y_1$ *and* $\widehat{v}(y_1)$ *is given by Eq. (19). When* $\xi_w = \xi_v$ *and* $N = 1$*, we obtain the maximal monopoly profit,* $\frac{(1-\beta_2\lambda_2)^2}{4\lambda_2}\mathrm{E}\left[\widehat{\theta}(\tilde{y}_1;\alpha,\xi_w)^2\right]$*. When* $\xi_w = \xi_v$ *and* $N \to \infty$*, arbitrageurs will compete away their aggregate trading profit.*

In the symmetric equilibrium of Theorem 2, the expectation of arbitrageurs' aggregate trading profit when they all follow the robust strategy $Z_{2,n}(y_1;\alpha,K(\xi_w))$ *is given by*

$$\mathrm{E}\left[\sum_{n=1}^{N}(\tilde{v}-\tilde{p}_2)Z_{2,n}\right] = \frac{N\alpha^2(1-\beta_2\lambda_2)^2}{(N+1)\lambda_2}\mathrm{E}\left[(\widehat{v}(\tilde{y}_1;\xi_v) - \widehat{v}^{\text{lasso}}(\tilde{y}_1;\xi_w))\widehat{\theta}^{\text{lasso}} + \frac{(\widehat{\theta}^{\text{lasso}})^2}{N+1}\right]. \quad (35)$$

When $0 < \xi_w \le \xi_v$*, the expected aggregate profit is always positive and has a positive limit:*

$$\lim_{N\to\infty}\mathrm{E}\left[\sum_{n=1}^{N}(\tilde{v}-\tilde{p}_2)Z_{2,n}\right] = \frac{\alpha^2(1-\beta_2\lambda_2)^2}{\lambda_2}\mathrm{E}[(\widehat{v}(\tilde{y}_1;\xi_v) - \widehat{v}^{\text{lasso}}(\tilde{y}_1;\xi_w))\widehat{\theta}^{\text{lasso}}(\tilde{y}_1;\xi_w)] > 0, \quad (36)$$

because $\widehat{v}(\tilde{y}_1;\xi_v) - \widehat{v}^{\text{lasso}}(\tilde{y}_1;\xi_w)$ *and* $\widehat{\theta}^{\text{lasso}}(\tilde{y}_1;\xi_w)$ *take the same signs when* $|\tilde{y}_1| > K(\xi_w)$*.*

Proof. See Appendix A.9. Here, we have used $\widehat{v}(\tilde{y}_1;\xi_v)$ to denote $\widehat{v}(\tilde{y}_1;\alpha=1,\xi_v)$. □

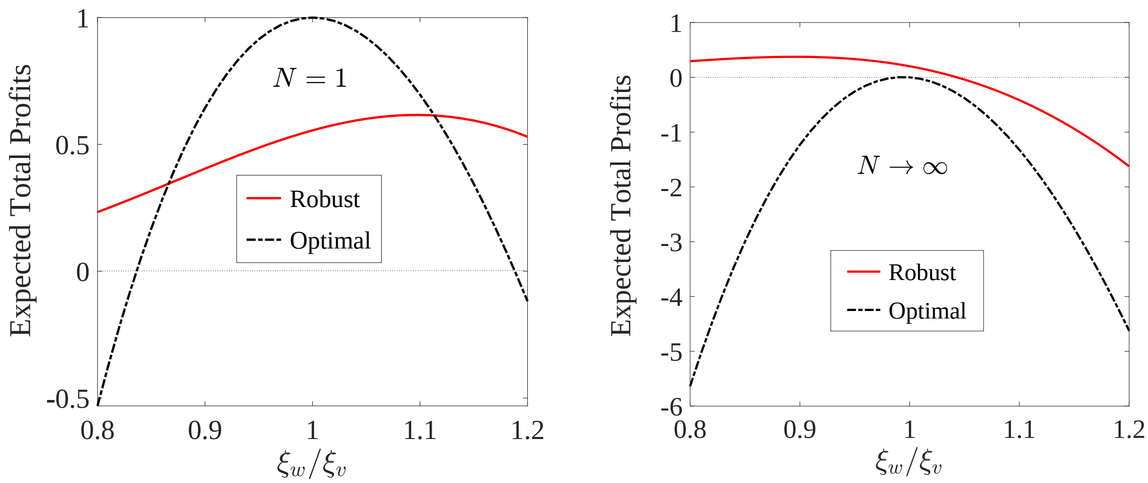

Figure 7. The profitability of the robust strategy $Z_{2,n}(\tilde{y}_1;\alpha,K(\xi_w))$ versus that of the optimal strategy $Z^o_{2,n}(\tilde{y}_1;\alpha,\xi_w)$ for a range of values of ξ_w relative to the true prior ξ_v.

Figure 7 compares the profitability of these two strategies for different values of ξ_w in two extreme cases. In the monopoly case ($N=1$), a single arbitrageur earns the maximal profit if she use the optimal strategy when $\xi_w = \xi_v$. Her expected profit is, however, sensitive to the bias and becomes negative when $|\xi_w - \xi_v|/\xi_v \gtrsim 20\%$. In contrast, the performance of the robust strategy is less sensitive to the bias and remains positive for a much wider range of ξ_w. By following the price trends the robust LASSO strategy is more likely to trade on the right side, whereas the optimal strategy can bet on the wrong side much more frequently. In the competitive case ($N \to \infty$), the total arbitrage profit under the optimal strategy is almost always negative unless $\xi_w \approx \xi_v$. This contrasts with the robust strategy which maintains positive profitability for $\xi_w \leq \xi_v$. While both strategies may lose money for $\xi_w > \xi_v$, the robust strategy lose much less. Our results share some similarities with the findings of Zhu and Zhou (2009) who report that the technical trading rules can be robust to model specification and tend to substantially outperform the seemingly optimal trading strategies under model uncertainty.

How can an infinite number of traders make significant profits when they follow the same strategy? Figure 8 (left) plots $[\widehat{v}(y_1;\xi_v) - \widehat{v}^{\text{lasso}}(y_1;\xi_w)] \cdot \widehat{\theta}^{\text{lasso}}(y_1;\xi_w)$ as a function of the input y_1 when $\xi_w = \xi_v$. This product term drives the positive limit of Eq. (36). It shows two sharp peaks in the outskirts of no-trade zone. Recall that the LASSO estimate shrinks the mean prediction (Figure 5): $|\widehat{v}^{\text{lasso}}(y_1)| < |\widehat{v}(y_1)|$. As this shrinkage is independent of N, it can actually benefit an arbitrary number of traders such that the entire group of them can buy (or sell) the asset on average below (or above) the fair price; see the right panel of Figure 8. Consequently, this *statistical arbitrage* remains profitable even when $N \to \infty$.

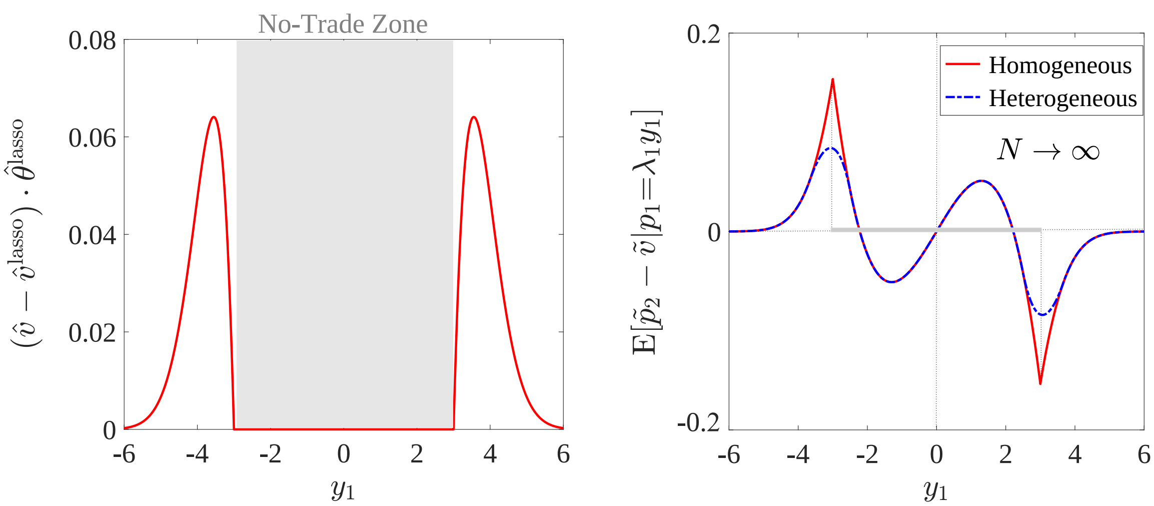

Figure 8. Left: $[\widehat{v}(y_1;\xi_v) - \widehat{v}^{\text{lasso}}(y_1;\xi_w)] \cdot \widehat{\theta}^{\text{lasso}}(y_1;\xi_w)$ as a function of y_1 when $\xi_w = \xi_v$. Right: $\mathrm{E}[\tilde{p}_2 - \tilde{v}|p_1 = \lambda_1 y_1]$ versus y_1, where the shaded segment indicates the no-trade zone. The red solid line is when arbitrageurs follow the robust strategy with the same threshold $K(\xi_v)$. The blue dashed line is when they use heterogeneous thresholds, denoted $K(\xi_{w,n})$. We impose $\frac{1}{N}\sum_{n=1}^{N} \xi_{w,n}^{-1} = \xi_v^{-1}$ to facilitate a fair comparison.

The non-vanishing profit earned by arbitrageurs can be attributed to their conservative trading outside the no-trade zone. When $\xi_w \leq \xi_v$, the robust LASSO strategy always trades less than the unbiased optimal strategy: $|Z_{2,n}(y_1; \alpha, K(\xi_w))| < |Z^o_{2,n}(y_1; \alpha, \xi_v)|$. This under-trading mitigates their competition, allowing traders to accumulate extra market power that facilitates a "cartel" to protect their profits. Eq. (28) allows us to rewrite Eq. (36) as

$$\frac{\alpha^2(1-\beta_2\lambda_2)^2}{\lambda_2}\mathrm{E}[(\widehat{v}-\widehat{v}^{\text{lasso}})\widehat{\theta}^{\text{lasso}}] = 2\alpha(1-\beta_2\lambda_2)\mathrm{E}[(\tilde{v}-\widehat{v}^{\text{lasso}})\cdot Z_{2,n}(\tilde{y}_1; \alpha, \xi_w, N=1)], \quad (37)$$

similar to the monopoly case that an arbitrageur can pay a cheaper price $\widehat{v}^{\text{lasso}}$ to receive $\tilde{v}$. This seemingly collusive outcome is not due to any trading frictions or financial constraints. Like tacit collusion, it requires no communication device or explicit agreements. The "cartel" is facilitated by traders' strategic exercise of robust optimization. By rewarding arbitrageurs, this novel effect serves as another channel that inevitably impedes market efficiency.

Price efficiency can be fully restored at $t = 2$ if the economy hosts an infinite number of arbitrageurs and all of them follow the unbiased optimal strategy. Nonetheless, if arbitrageurs are constrained by model risks and all adopt the robust LASSO strategy, then there will be persistent pricing errors in the neighborhoods of $p_1 = \pm\lambda_1 K$, as shown in the right panel of Figure 8. In general, whenever a mass of them are constrained by model risks, they will trade conservatively, amass extra market power, and sustain inefficient prices.

Proposition 4. *Suppose the economy hosts an infinite number of risk-neutral arbitrageurs ($N \to \infty$). If there is ever a finite measure $\phi \in (0, 1]$ of them constrained by model risk as in Theorem 2, then the asset price $\tilde{p}_2$ is inefficient for almost all realizations of $\tilde{p}_1 = p_0 + \lambda_1\tilde{y}_1$:*

$$\mathrm{E}[\tilde{p}_2 - \tilde{v}|\tilde{p}_1] \to \alpha(1-\beta_2\lambda_2)\left\{[\widehat{v}(\tilde{y}_1; \xi_w) - \widehat{v}(\tilde{y}_1; \xi_v)] + \phi\left[\widehat{\theta}^{\text{lasso}}(\tilde{y}_1; \xi_w) - \widehat{\theta}(\tilde{y}_1; \xi_w)\right]\right\} \neq 0. \quad (38)$$

Proof. See Appendix A.10. Note that we have used $\widehat{\theta}(\tilde{y}_1; \xi_w)$ to stand for $\widehat{\theta}(\tilde{y}_1; \alpha = 1, \xi_w)$. □

Eq. (38) shows two sources for inefficient prices at $t = 2$. One is the estimation bias (i.e., $\xi_w \neq \xi_v$) which is applicable to all traders. The other is the under-trading by a fraction of traders who exercise robust control. The price is inefficient almost everywhere, unless all the arbitrageurs know the true prior (i.e., $\xi_w = \xi_v$ and $\phi = 0$). Bossaerts et al. (2010) and Ahn et al. (2014) document considerable heterogeneity in risk and ambiguity aversion and find that a fraction of individuals' behavior is consistent with the standard expected utility. Proposition 4 shows that our result is robust to investors' heterogeneous preferences.

This implication also holds when arbitrageurs use the LASSO strategy (27) with heterogeneous thresholds, denoted $K(\xi_{w,n})$ for $n = 1, ..., N$. The aggregate robust trading with heterogeneous thresholds becomes a smoother function of y_1 (Figure 16 in Appendix A.10). The right panel of Figure 8 shows that heterogeneous thresholds can partially smooth out the "cartel" effect but cannot restore price efficiency.

6 Extensions and Applications

6.1 Uncertainty about the frequency of fat-tail shocks

Our setup can be easily extended to the case that traders are uncertain about both prior parameters, α and ξ_v. In other words, they face model risk about the frequency and magnitude of mispricings simultaneously. Under the Gaussian-Laplacian mixture distribution (9), the *fat-tailedness* (as measured by the raw kurtosis) of $\tilde{v}$ is given by $3+3\frac{(4+\alpha)(1-\alpha)}{(2-\alpha)^2} \in (3, \frac{49}{8}]$, which is a simple function of α. So we are considering a general situation where traders face uncertainty about not only the stochastic volatility but also the fat-tailedness of stock value.

Proposition 5. *Suppose arbitrageurs are uncertain about both α and ξ_v, with common priors denoted as $\mathcal{LG}(\tilde{\alpha}, \tilde{\xi})$ where $\tilde{\alpha} \in [\alpha_L, \alpha_H]$ and as before $\tilde{\xi} \in [\xi_L, \xi_H]$. If they are asymptotically ambiguity-neutral as (C1) in Theorem 2 and if the other two conditions (C2) and (C3) also hold, then there exists a symmetric debiased equilibrium where arbitrageurs watch the market without trading at $t = 1$ and they follow a robust LASSO strategy at $t = 2$:*

$$Z_{2,n}(y_1; \overline{\alpha}, K(\xi_w)) = \frac{\overline{\alpha}(1-\beta_1\lambda_1)(1-\beta_2\lambda_2)}{\beta_1\lambda_2(N+1)}\mathcal{S}(y_1; K) = \frac{\overline{\alpha}(1-\beta_2\lambda_2)}{\lambda_2(N+1)} \cdot \widehat{\theta}^{\text{lasso}}(y_1; K), \quad (39)$$

where $\overline{\alpha} := \mathrm{E}[\tilde{\alpha}]$ is the prior mean of $\tilde{\alpha}$ and $\widehat{\theta}^{\text{lasso}}(y_1; K)$ is still defined by Eq. (30).

Proof. Eq. (39) follows from Theorem 2 and the scaling property that α can be factored out from the fixed-prior solution (27). The more rigorous proof resembles Appendix A.6. □

Recall from Figure 3 that the robust LASSO strategy has the same threshold $K(\xi_w)$ for any positive values of α. With Bayesian learning, traders average across all possible priors. They respond with $\overline{\alpha} = \mathrm{E}[\tilde{\alpha}]$ in Eq. (39), since $\mathrm{E}[Z_{2,n}(y_1; \tilde{\alpha}, K)] = \mathrm{E}[\tilde{\alpha}]Z_{2,n}(y_1; \alpha = 1, K)$. This also explains why our original setup only focuses on the prior uncertainty about ξ_v.

Thus, our derivation of the LASSO strategy holds in the general situation where agents have formed a mixture prior on the estimated parameter but know little about the frequency or the scale of its fat-tail component. This general applicability is in sharp contrast with the MAP-based interpretation which requires a pure and fixed Laplace prior.

6.2 LASSO for a long-short portfolio of mispriced stocks

The result below rationalizes the application of a vector form of the LASSO algorithm. It is a technical strategy that forecasts and exploits multiple misvalued stocks simultaneously.

Proposition 6. *Suppose arbitrageurs anticipate multiple independent assets to be mispriced. Each asset, indexed by $j = 1, ..., J$, is traded in a two-period environment with linear price movements as in (11) and aggregate order flows as in (15). Arbitrageurs have uncertain priors on each asset, denoted $\tilde{v}_j \sim \mathcal{LG}(\tilde{\alpha}_j, \tilde{\xi}_j)$. If (C1)-(C3) in Theorem 2 hold, there is an equilibrium where arbitrageurs choose to idle at $t = 1$ and follow a LASSO strategy at $t = 2$:*

$$\mathbf{Z}_{2,n}(\mathbf{p}_1) := \left\{ \frac{\overline{\alpha}_j(1 - \beta_{2,j}\lambda_{2,j})}{\lambda_{2,j}(N+1)} \cdot \widehat{\theta}_j^{\text{lasso}}(p_{1,j}; \xi_{w,j}) : j = 1, ..., J \right\}, \tag{40}$$

where $\overline{\alpha} = \mathrm{E}^j[\tilde{\alpha}_j]$ is the prior mean frequency and $\xi_{w,j}$ is the weighted harmonic mean of $\tilde{\xi}_j$. The vector of LASSO estimates, $\widehat{\mathbf{\Theta}}^{\text{lasso}} = \{\widehat{\theta}_j^{\text{lasso}}(p_{1,j}; \xi_{w,j}) : j = 1, ..., J\}$, is defined by

$$\widehat{\mathbf{\Theta}}^{\text{lasso}} := \underset{\{\theta_1,...,\theta_J\}}{\arg\min} \sum_{j=1}^{J} \left\{ \frac{1}{2} \left| p_{1,j} - \frac{\beta_{1,j}\lambda_{1,j}}{1 - \beta_{1,j}\lambda_{1,j}} \theta_j \right|^2 + \left(\frac{\lambda_{1,j}\sigma_{u,j}}{1 - \beta_{1,j}\lambda_{1,j}} \right)^2 \frac{|\theta_j|}{\xi_{w,j}} \right\}. \tag{41}$$

Proof. See Appendix A.11. This is based on Theorem 2, Theorem 3, and Proposition 5. □

Proposition 6 can give rational interpretations of certain technical or algorithmic trading rules. These are triggered to trade whenever stock prices hit across predefined price levels (Lo, Mamaysky, and Wang (2000); Zhu and Zhou (2009); Han, Liu, Zhou, and Zhu (2021)). At first glance, such mechanical trading plans are at odds with Bayesian rationality. Proposition 6 suggests that simple trading rules for constructing a long-short portfolio may well be the solution of sophisticated risk management. This viewpoint agrees with Zhu and Zhou (2009) who show that the widely used moving average trading rule can add value to asset allocation under uncertainty about the return predictability or about the date-generating process of stock prices.

Proposition 6 can also provide a rational view on those feedback traders who extrapolate price trends and who are often stereotyped in the behavioral literature (DeLong et al. (1990), Barberis et al. (2015, 2018)). The multi-asset LASSO strategy (40) shows a similar extrapolative pattern. Its two momentum "arms" are ready to long and short equities when their prices move beyond the endogenous no-trade zones. This momentum strategy responds to most recent winners and losers.

6.3 Sparse predictors in the cross-section of stock returns

When a market lacks weak-form efficiency, we may predict stock returns based on historical stock prices and trading volumes. This is not restricted to the prediction of a stock by using its own trading data (as in Theorem 2 or Proposition 6) but also permits cross-asset return predictions (Chinco et al. (2019) and Kelly et al. (2023)). Here, we briefly extend our theory to understand their empirical quest. Our discussion may provide some intuition about the return predictability and the usefulness of LASSO in their study. We also propose that the LASSO may be used to filter predictive stock returns for robustness.

Consider a two-period market where a large number of stocks are traded. Some stocks have fat tails which may be correlated with the fat tails of other stocks. Like Chinco et al. (2019), we consider the task to predict the next-minute stock returns using the entire cross-section return data over a short-term window, denoted as $\{r_{1,j} : j = 1, ..., J\}$. To be specific, we attempt to forecast stock k's next-minute return, $\tilde{r}_{2,k}$. There are perhaps a small number of stocks that have predictive power for $\tilde{r}_{2,k}$. Let $\mathbf{S}_k \coloneqq \{j : \tilde{r}_{1,j} \text{ is informative about } \tilde{r}_{2,k}\}$ be the subset of such stocks (predictors). For simplicity, suppose we have learned about the subset $\mathbf{S}_k$. This is actually achieved in Chinco et al. (2019) by applying the LASSO to select the subset of predictors. Suppose the economy operates in a way that the mispricing of stock k is approximately a linear combination of the mispricings of stocks in the subset $\mathbf{S}_k$:

$$\tilde{\theta}_k \approx \sum_{j \in \mathbf{S}_k} \Omega_{k,j} \tilde{\theta}_j = \sum_{j \in \mathbf{S}_k} \Omega_{k,j} (\tilde{v}_j - p_{1,j}), \tag{42}$$

where $\Omega_{k,j}$ reflects the degree of predictive power of stock j on stock k. Eq. (42) serves as an assumption. The microfoundation is beyond the scope of this paper. Correlated fat-tail mispricings might arise from correlated news shocks or from asynchronous trading on dispersed private signals drawn from a multivariate fat-tail distribution.

As before, our prior belief for each stock $j \in \mathbf{S}_k$ is $\tilde{v}_j \sim \mathcal{LG}(\alpha_j, \tilde{\xi}_j)$ where $\tilde{\xi}_j$ represents our prior uncertainty about the scale of fat-tail shocks. If the conditions for Theorem 2 hold for each stock, then we can apply the LASSO estimate for each $\tilde{\theta}_j$ (by Eq. (30)) and add them together to form a robust estimate of the pricing error $\tilde{\theta}_k$:

$$\widehat{\theta}_k^{\text{rob}} \coloneqq \sum_{j \in \mathbf{S}_k} \Omega_{k,j} \widehat{\theta}_j^{\text{lasso}} = \sum_{j \in \mathbf{S}_k} \Omega_{k,j} \alpha_j \frac{1 - \beta_{1,j}\lambda_{1,j}}{\beta_{1,j}} \left[y_{1,j} - \operatorname{sign}(y_{1,j}) K_j(\xi_{w,j}) \right] \mathbf{1}_{|y_{1,j}| > K_j(\xi_{w,j})}. \tag{43}$$

In general, $\widehat{\theta}_k^{\text{rob}}$ is not exactly the LASSO estimate of $\tilde{\theta}_k$, because the threshold K_j can be different across predictors, preventing $\widehat{\theta}_k^{\text{rob}}$ from being a simple soft-thresholding function.

If we normalize the initial stock price to be one (i.e., $p_{0,j} = 1$) for each stock $j \in \mathbf{S}_k$, then the lagged return for stock j can be written as $r_{1,j} = (p_{1,j} - p_{0,j})/p_{0,j} = p_{1,j} - 1 = \lambda_{1,j} y_{1,j}$. Suppose the short-term cross-sectional return predictability has not been exploited by any traders (e.g., prior to the publication of Chinco et al. (2019)). Then, the order flows for each stock can be described by Eq. (12). For stock k, the next-minute return is

$$\tilde{r}_{2,k} := \frac{\tilde{p}_{2,k} - p_{1,k}}{p_{1,k}} = \frac{\lambda_{2,k}\tilde{y}_{2,k}}{p_{1,k}} = \frac{\lambda_{2,k}(\beta_{2,k}\tilde{\theta}_k + \tilde{u}_k)}{1 + r_{1,k}} \approx \frac{\lambda_{2,k}\beta_{2,k}\sum_{j\in\mathbf{S}_k}\Omega_{k,j}\tilde{\theta}_j + \lambda_{2,k}\tilde{u}_k}{1 + r_{1,k}}. \quad (44)$$

The posterior mean estimate of $\tilde{r}_{2,k}$ is

$$\widehat{r}_{2,k} := \mathrm{E}\left[\tilde{r}_{2,k} | \{r_{1,j} : j \in \mathbf{S}_k\}\right] = \frac{\lambda_{2,k}\beta_{2,k}}{1 + r_{1,k}} \sum_{j\in\mathbf{S}_k} \Omega_{k,j} \int w_j(\xi) \mathrm{E}\left[\tilde{\theta}_j | r_{1,j}; \alpha_j, \tilde{\xi}_j = \xi\right] d\xi, \quad (45)$$

which is a complicated nonlinear function of the vector $\{r_{1,j} : j \in \mathbf{S}_k\}$. To make a robust estimate of $\tilde{r}_{2,k}$, we can replace the above posterior means with the LASSO estimates:

$$\widehat{r}^{\text{rob}}_{2,k} := \frac{\lambda_{2,k}\beta_{2,k}}{1 + r_{1,k}} \sum_{j\in\mathbf{S}_k} \Omega_{k,j} \widehat{\theta}^{\text{lasso}}_j(r_j; \alpha_j, K_j(\xi_{w,j})) = a_k + \sum_{j\in\mathbf{S}_k, |r_{1,j}| > \lambda_{1,j}K_j} b_{k,j} r_{1,j}. \quad (46)$$

This is a simple linear function of the lagged stock returns (predictors), where

$$a_k := - \sum_{j\in\mathbf{S}_k, |r_{1,j}| > \lambda_{1,j}K_j} b_{k,j}\lambda_{1,j}K_j \mathrm{sign}(r_{1,j}) \quad \text{and} \quad b_{k,j} := \frac{\lambda_{2,k}\beta_{2,k}}{1 + r_{1,k}} \Omega_{k,j}\alpha_j \frac{1 - \beta_{1,j}\lambda_{1,j}}{\lambda_{1,j}\beta_{1,j}}. \quad (47)$$

Eq. (46) expresses both sparsity and linearity of predictive signals in the cross-section of stock returns. The sparsity comes from two selections, $j \in \mathbf{S}_k$ and $|r_{1,j}| > \lambda_{1,j}K_j$. This may help understand the empirical findings of Chinco et al. (2019) who have focused on the first selection $\mathbf{S}_k$ and use the LASSO to pick up the nontrivial coefficients $b_{k,j}$ of all candidate predictors $\{r_{1,j} : j = 1, ..., J\}$ where J is a large number. Their variable selection ($j \in \mathbf{S}_k$) is different from our filtering ($|r_{1,j}| > \lambda_{1,j}K_j$) on predictors, although both have involved the LASSO estimates. Our argument of robust optimization will effectively impose the filtering threshold $\lambda_{1,j}K_j$ on the observed predictive stock returns. Here, $\lambda_{1,j}$ reflects the transaction cost per unit of order flow, while K_j is the endogenous threshold that appears in the LASSO estimation of the residual signal $\tilde{\theta}_j$.[20] The filtering effect in our theory can also "sparsify" the selection of return predictors. Therefore, we propose an application of LASSO to directly filter the cross-section of stock returns. This procedure may improve the robustness and even the performance in the prediction task of Chinco et al. (2019).

[20] Chinco et al. (2019) have taken into account the bid-ask spread in their LASSO-implied strategy. This spread cost may partially capture the filtering effect of $\lambda_{1,j}K_j$ derived from our robust optimization.

6.4 Ridge regression based on a Gaussian mixture prior

In the statistical interpretation of LASSO, the l_1 penalty term is inserted into the objective function through the MAP estimation under a pure Laplace prior. In our interpretation, the effect of l_1 penalty is imposed effectively by the max-min decision criteria and the curvature condition (C2) in Theorem 2. (C2) helps regularize the robust optimization problem by requiring admissible strategies to preserve the curvature of Bayesian optimal strategies.

The key assumption in our previous development is the fat-tailed prior distribution. If we replace this with a Gaussian mixture prior and keep everything else equal, we can show that the max-min robust strategy is the *ridge regression*. In contrast to the LASSO, the ridge regression contains an l_2 penalty in its objective function and thus uniformly shrinks all coefficients without sending any of them to zero.

Proposition 7. *If, instead, arbitrageurs have a mixture Gaussian prior about the asset liquidation value, then their equilibrium strategy, either optimal or robust, is observationally equivalent to the ridge regression which is a linear function of the input order flow* y_1 *without any finite inaction region.*

Proof. See Appendix A.12. □

The conditions (C1)-(C3) in Theorem 2 are not needed in Proposition 7. These conditions are needed to ensure the existence of a *debiased* equilibrium when traders have uncertain fat-tail priors. If we remove the assumption of fat tails, we are back to the Gaussian world where inference problems are often linear and the mean estimate coincides with the mode estimate. As a control, Proposition 7 shows that the fat-tail distribution is the key assumption in our interpretation of the LASSO.

Ridge regression is another basic and popular machine learning method. Its l_2 penalty has been integrated in other techniques. For example, the *elastic net* is a regularized regression that linearly combines the l_1 and l_2 penalties of the LASSO and ridge methods (Hastie et al. (2009)). While Proposition 7 seems to rationalize the ridge regression, we choose to focus on the LASSO regression in this paper. Our purpose is to highlight the importance of Bayesian rational thinking before we apply machine learning methods to economic problems. Of course, we do not expect every machine learning algorithm to have a rational interpretation or version. It is left for future work to understand the economic rationales (if any) underneath many other techniques.

7 Conclusion

Machine learning seems an inevitable trend in the era of big data. Are machine learning methods heuristic approximations or rational choices for economic agents? As a starting point, this paper provides a Bayesian-rational interpretation for one of the most widely used machine learning method, the LASSO algorithm. Unlike the interpretation by Tibshirani (1996), our rationalization of LASSO does not invoke the heuristic MAP (i.e., the posterior mode) estimation or hinge on the restrictive assumption of a pure fixed Laplace prior. In our setup, agents (i.e., arbitrageurs) consistently use the Bayesian-rational learning (i.e., the posterior mean estimate) to evaluate all possible states. They have sequential rationality by using dynamic programming to solve their well-defined max-min expected utility function. Under general fat-tailed model risks, their robust strategy is a LASSO strategy. We also show that this robust LASSO strategy can reduce traders' competition even when the number of them becomes infinite. This induces a seemingly collusive outcome as traders' total profit does not vanish even in the competitive limit. This is a novel mechanism for limited arbitrage.

This paper provides a showcase for a popular machine learning method within the neoclassic framework of financial economics. It calls for more interdisciplinary studies to better understand the economic rationales and implications of other machine learning techniques.

A Appendix

A.1 Example of Microfoundation

Here, we discuss one microfoundation for the trading environment which can serve as the background of our model setup. Consider a two-period model of Kyle (1985) with multiple ($M \geq 1$) informed traders as extended by Holden and Subrahmanyam (1992). Suppose all market participants (informed traders and market makers) hold the common knowledge that the stock liquidation value is normally distributed. Then there exists a unique subgame perfect linear equilibrium, based on the general procedures of Proposition 1 in Holden and Subrahmanyam (1992). We use the same notation $\Sigma_t := \text{Var}(\tilde{v}_0|\tilde{p}_t, \tilde{p}_{t-1}, ...\tilde{p}_0)$ for $t \in \{0, 1, 2\}$, which is the posterior variance of $\tilde{v}_0$ conditional on the price history up to time t. There are competitive market makers who will accommodate the following aggregate order flows:

$$\tilde{y}_1 = \sum_{m=1}^{M} \tilde{x}_{m,1} + \tilde{u}_1 = \beta_1(\tilde{v} - p_0) + \tilde{u}_1, \tag{A1}$$

$$\tilde{y}_2 = \sum_{m=1}^{M} \tilde{x}_{m,2} + \tilde{u}_2 = \beta_2(\tilde{v} - \tilde{p}_1) + \tilde{u}_2, \tag{A2}$$

where $\tilde{x}_{m,t}$ denotes the order placed by the m-th informed trader at time t and β_t represents the aggregate informed trading intensity at time t. In the conjectured linear equilibrium, market makers follow the linear pricing strategies below:

$$\tilde{p}_1 = p_0 + \lambda_1 \left(\sum_{m=1}^{M} \tilde{x}_{m,1} + \tilde{u}_1\right) = p_0 + \lambda_1 \tilde{y}_1, \tag{A3}$$

$$\tilde{p}_2 = \tilde{p}_1 + \lambda_2 \left(\sum_{m=1}^{M} \tilde{x}_{m,2} + \tilde{u}_2\right) = \tilde{p}_1 + \lambda_2 \tilde{y}_2, \tag{A4}$$

where the pricing coefficients (Kyle lambdas) are given by

$$\lambda_1 = \frac{\beta_1 \Sigma_1}{\sigma_u^2}, \qquad \lambda_2 = \frac{\beta_2 \Sigma_2}{\gamma \sigma_u^2}, \tag{A5}$$

As a boundary condition, it is easy to derive the total intensity of informed trading at $t = 2$:

$$\beta_2 = \frac{M}{\lambda_2(M+1)}. \tag{A6}$$

By backward induction, we can further derive the total intensity of informed trading at $t = 1$,

$$\beta_1 = \frac{\delta M(M+1)^2 - 2M}{\lambda_1[\delta(M+1)^3 - 2M]}, \tag{A7}$$

where $\delta := \lambda_2/\lambda_1$ is the ratio of Kyle lambdas. These results imply that

$$1-\beta_1\lambda_1 = \frac{\delta(M+1)^2}{\delta(M+1)^3-2M}, \qquad 1-\beta_2\lambda_2 = \frac{1}{M+1}.$$

The optimal trading strategy for each informed trader is

$$\tilde{x}_{m,1} = \frac{\beta_1}{M}(\tilde{v}-p_0) = \frac{\delta(M+1)^2-2}{\delta(M+1)^3-2M}\cdot\frac{\tilde{v}-p_0}{\lambda_1}, \tag{A8}$$

$$\tilde{x}_{m,2} = \frac{\beta_2}{M}(\tilde{v}-p_1) = \frac{\tilde{v}-\tilde{p}_1}{\lambda_2(M+1)}. \tag{A9}$$

The Bayesian update of the posterior variance about $\tilde{v}$ follows

$$\Sigma_2 = (1-\beta_2\lambda_2)\Sigma_1 = \frac{\Sigma_1}{M+1}, \qquad \Sigma_1 = (1-\beta_1\lambda_1)\Sigma_0 = \frac{\delta(M+1)^2\sigma_v^2}{\delta(M+1)^3-2M}. \tag{A10}$$

Combining (A5), (A6), (A7), (A8), (A9), and (A10), we find the expressions of λ_1 and λ_2:

$$\lambda_1 = \frac{\sqrt{\delta M(M+1)^2(\delta(M+1)^2-2)}}{\delta(M+1)^3-2M}\cdot\frac{\sigma_v}{\sigma_u}, \tag{A11}$$

$$\lambda_2 = \delta\lambda_1 = \sqrt{\frac{\delta M/\gamma}{\delta(M+1)^3-2M}}\cdot\frac{\sigma_v}{\sigma_u}. \tag{A12}$$

In order to have $\lambda_1 > 0$ and $\lambda_2 > 0$, it is equivalent to impose $\delta(M+1)^2 > 2$. This will guarantee that the denominators in (A11) and (A12) are strictly positive, $\delta(M+1)^3-2M > 0$. These allow us to rewrite the ratio of Kyle lambdas which is equal to δ by definition:

$$\frac{\lambda_2^2}{\lambda_1^2} = \frac{\delta(M+1)^3-2M}{\delta\gamma(M+1)^4-2\gamma(M+1)^2} = \delta^2, \tag{A13}$$

By Eq. (A13), the equilibrium ratio $\delta = \delta(M,\gamma)$ must satisfy the cubic equation:

$$(M+1)^4\gamma\delta^3 - 2(M+1)^2\gamma\delta^2 - (M+1)^3\delta + 2M = 0, \quad \text{s.t.} \quad \delta(M+1)^2 > 2. \tag{A14}$$

Huddart et al. (2001) studied the two period Kyle model with a single informed trader ($M=1$) and constant noise trading volatility ($\gamma = 1$). They obtain the cubic equation $8\delta^3 - 4\delta^2 - 4\delta + 1 = 0$ which coincides with (A14) if we set $M=1$ and $\gamma=1$. The economic solution is the largest root $\delta \approx 0.901$. Similarly, when $M > 1$, there is a unique solution that meets the second order condition and the requirement $\lambda_1 > 0$ and $\lambda_2 > 0$.

This microfoundation is used in all our numerical examples.

A.2 Proof of Theorem 1

Since arbitrageurs' prior is non-directional, we conjecture first and verify later that they do not trade in the first period, $Z_{1,n} = 0$, for $n = 1, ..., N$. Under this conjecture, we can solve their optimal strategy at $t = 2$. Arbitrageurs conjecture the market-clearing price as

$$\tilde{p}_2 = \tilde{p}_1 + \lambda_2 \left[\beta_2(\tilde{v} - \tilde{p}_1) + \sum\nolimits_{n=1}^{N} Z_{2,n}(\tilde{p}_1) + \tilde{u}_2 \right]. \tag{A15}$$

They estimate $\tilde{v}$ conditional on the observed order flow y_1 and the Laplacian-Gaussian prior $\tilde{v} \sim \mathcal{LG}(\alpha, \tilde{\xi})$. For any fixed prior $\tilde{\xi} = \xi \in (0, \infty)$, the n-th trader solves her optimal strategy,

$$Z^o_{2,n}(p_1; \alpha, \xi) = \arg\max_{z_{2,n}} \mathrm{E}^{\mathcal{A}} \left[(\tilde{v} - \tilde{p}_2) z_{2,n} | p_1 \right], \tag{A16}$$

where $\mathcal{A}$ denotes the prior belief $\tilde{v} \sim \mathcal{LG}(\alpha, \xi)$. We can use $Z^o_{2,-n} = \sum_{n' \neq n} Z^o_{2,n'}$ to denote the aggregate trading by all arbitrageurs except the n-th one. The first order condition is

$$\mathrm{E}^{\mathcal{A}}[\tilde{v}|p_1] - p_1 = \lambda_2 \left(\beta_2 \mathrm{E}^{\mathcal{A}}[\tilde{v}|p_1] - \beta_2 p_1 + 2 z_{2,n} + \mathrm{E}^{\mathcal{A}}[Z^o_{2,-n}|p_1] \right). \tag{A17}$$

Let $\widehat{v}(p_1; \alpha, \xi) = \mathrm{E}^{\mathcal{A}}[\tilde{v}|p_1]$ be the posterior mean estimate of $\tilde{v}$. Then the strategy solution is

$$Z^o_{2,n}(p_1; \alpha, \xi) = \frac{(1 - \beta_2\lambda_2)(\widehat{v} - p_1)}{2\lambda_2} - \frac{\mathrm{E}^{\mathcal{A}}[Z^o_{2,-n}(p_1)|p_1]}{2} \tag{A18}$$

The n-th arbitrageur conjectures that every other arbitrageur solves the same problem and trades $Z^o_{2,n'} = \eta \cdot (\widehat{v} - p_1)$ for any $n' \neq n$, with a coefficient η to be determined. This suggests

$$Z^o_{2,n}(p_1; \alpha, \xi) = \frac{\widehat{v} - p_1}{2\lambda_2} \left[1 - \beta_2\lambda_2 - \eta\lambda_2(N - 1) \right]. \tag{A19}$$

As arbitrageurs make the same conjecture in a symmetric equilibrium, they can find that $\eta = \frac{1 - \beta_2\lambda_2 - \eta\lambda_2(N-1)}{2\lambda_2}$, which has a unique solution

$$\eta = \frac{1 - \beta_2\lambda_2}{\lambda_2(N + 1)} > 0. \tag{A20}$$

Since $p_1 = p_0 + \lambda_1 y_1$, their optimal strategy at $t = 2$ under the fixed prior $\mathcal{LG}(\alpha, \xi)$ is

$$Z^o_{2,n}(y_1; \alpha, \xi) = \frac{1 - \beta_2\lambda_2}{\lambda_2(N + 1)} (\widehat{v}(p_1; \alpha, \xi) - p_1) = (1 - \beta_2\lambda_2) \frac{\widehat{\theta}(y_1; \alpha, \xi)}{\lambda_2(N + 1)}, \quad \text{for} \;\; n = 1, ..., N. \tag{A21}$$

Let $\tilde{x}_1 := \beta_1(\tilde{v} - p_0)$ denote the total informed order flow at $t = 1$. Then, arbitrageurs' prior about $\tilde{x}_1$ is another Laplacian-Gaussian mixture distribution:

$$f(x_1) = \frac{\alpha}{2\beta_1\xi}\exp\left(-\frac{|x_1|}{\beta_1\xi}\right) + \frac{1-\alpha}{\beta_1\sqrt{2\pi\sigma_v^2}}\exp\left(-\frac{x_1^2}{2(\beta_1\sigma_v)^2}\right). \tag{A22}$$

By Bayes' rule, the posterior probability of x_1 conditional on y_1 is found to be

$$\begin{aligned} & f(x_1|y_1) = f(y_1,x_1)/f(y_1) = f(y_1|x_1)f(x_1)/f(y_1) \\ = \; & \frac{\alpha/(2\beta_1\xi)}{\sqrt{2\pi\sigma_u^2}f(y_1)}\exp\left[-\frac{(y_1-x_1)^2}{2\sigma_u^2} - \frac{|x_1|}{\beta_1\xi}\right] + \frac{(1-\alpha)/\beta_1}{2\pi\sigma_u\sigma_v f(y_1)}\exp\left[-\frac{(y_1-x_1)^2}{2\sigma_u^2} - \frac{x_1^2}{2(\beta_1\sigma_v)^2}\right]. \end{aligned}$$

By direct integration, the probability density function of $\tilde{y}_1 = \beta_1(\tilde{v}-p_0) + \tilde{u}_1 = \tilde{x}_1 + \tilde{u}_1$ is

$$\begin{aligned} f(y_1) \; = \; & \frac{\alpha}{4\beta_1\xi}\exp\left(\frac{\sigma_u^2}{2(\beta_1\xi)^2}\right)\left[e^{-\frac{y_1}{\beta_1\xi}}\operatorname{erfc}\left(\frac{\sigma_u^2/(\beta_1\xi)-y_1}{\sqrt{2}\sigma_u}\right) + e^{\frac{y_1}{\beta_1\xi}}\operatorname{erfc}\left(\frac{\sigma_u^2/(\beta_1\xi)+y_1}{\sqrt{2}\sigma_u}\right)\right] \\ + \; & \frac{1-\alpha}{\sqrt{2\pi(\sigma_u^2+(\beta_1\sigma_v)^2)}}\exp\left[-\frac{y_1^2}{2(\sigma_u^2+(\beta_1\sigma_v)^2)}\right]. \end{aligned} \tag{A23}$$

Define two dimensionless parameters: $\kappa(\xi) \coloneqq \sigma_u/(\beta_1\xi)$ and $\varrho \coloneqq (\beta_1\sigma_v/\sigma_u)^2$. By rescaling the order flows as $y_1 = y\sigma_u$, we can express $f(y_1 = y\sigma_u)$ in a dimensionless form

$$f(y) = \frac{\alpha\kappa e^{\frac{\kappa^2}{2}}}{4}\left[e^{-\kappa y}\operatorname{erfc}\left(\frac{\kappa-y}{\sqrt{2}}\right) + e^{\kappa y}\operatorname{erfc}\left(\frac{\kappa+y}{\sqrt{2}}\right)\right] + \frac{1-\alpha}{\sqrt{2\pi(1+\varrho)}}\exp\left[-\frac{y^2}{2(1+\varrho)}\right],$$

which is a symmetric function and decays exponentially at large $|y|$. Bayes' rule implies that

$$\mathrm{E}^{\mathcal{A}}[\tilde{x}_1 = x\sigma_u | y_1 = y\sigma_u, \xi] = \sigma_u\int_{-\infty}^{\infty} x f(x|y)dx = \sigma_u\int_{-\infty}^{\infty} x f(y|x)f(x)/f(y)dx. \tag{A24}$$

Given all the above results, we can derive the posterior expectation of $\tilde{v}$:

$$\begin{aligned} & \widehat{v}(y) = \mathrm{E}^{\mathcal{A}}[\tilde{v}|y] = \alpha\mathrm{E}^{\mathcal{A}}[\tilde{v}|y, s=1] + (1-\alpha)\mathrm{E}^{\mathcal{A}}[\tilde{v}|y, s=0] \\ = \; & \frac{\alpha\xi\kappa(y-\kappa)\operatorname{erfc}\left(\frac{\kappa-y}{\sqrt{2}}\right)}{\operatorname{erfc}\left(\frac{\kappa-y}{\sqrt{2}}\right) + e^{2\kappa y}\operatorname{erfc}\left(\frac{\kappa+y}{\sqrt{2}}\right)} + \frac{\alpha\xi\kappa(y+\kappa)\operatorname{erfc}\left(\frac{\kappa+y}{\sqrt{2}}\right)}{\operatorname{erfc}\left(\frac{\kappa+y}{\sqrt{2}}\right) + e^{-2\kappa y}\operatorname{erfc}\left(\frac{\kappa-y}{\sqrt{2}}\right)} + (1-\alpha)\lambda_1 y\sigma_u. \end{aligned} \tag{A25}$$

One can verify that $\widehat{v}(y)$ is an increasing function of $y \coloneqq p_1/(\lambda_1\sigma_u)$ and its shape depends on two dimensionless parameters, α and $\kappa(\xi)$. Asymptotically, $\widehat{v}$ becomes linear in y_1:

$$\widehat{v}(y_1) \to \alpha[y_1 - \operatorname{sign}(y_1)\kappa\sigma_u]/\beta_1 + (1-\alpha)\lambda_1 y_1, \qquad \text{as} \quad |y_1| \to \infty. \tag{A26}$$

A.3 Proof of Proposition 1

Now we examine the equilibrium existence condition. If the n-th trader does not deviate from the conjectured no-trade strategy at $t = 1$, her optimal strategy should be given by the original result $Z^o_{2,n}(p_1;\alpha,\xi) = \eta \cdot [\widehat{v}(p_1;\alpha,\xi) - p_1]$, where $p_1 = p_0 + \lambda_1 y_1 = \lambda_1(\beta_1 v + u_1)$. To verify that no arbitrageur would trade at $t = 1$, we have to examine the condition (16). Suppose the n-th trader deviates from the no-trade strategy by placing an order $Z'_{1,n}{=}z_1 \neq 0$. Then the actual total order flow at $t = 1$ will be $\tilde{y}'_1 = \beta_1\tilde{v} + \tilde{z}_1 + \tilde{u}_1$, instead of $\tilde{y}_1 = \beta_1\tilde{v} + \tilde{u}_1$ in the conjectured equilibrium. Taking as given $Z^o_{2,n'}(y'_1;\alpha,\xi) = \eta \cdot [\widehat{v}(y'_1;\alpha,\xi) - \lambda_1 y'_1]$ for any other trader $n' \neq n$, the n-th arbitrageur can solve her best response at $t = 2$ conditional on her knowledge of y_1 and z_1. It is found to be

$$
\begin{aligned}
Z'_{2,n}(y'_1;\alpha,\xi) &= \frac{\widehat{v}(y_1;\alpha,\xi) - \lambda_1 y'_1}{2\lambda_2} - \frac{\mathrm{E}^{\mathcal{A}}[\beta_2(\tilde{v} - \lambda_1 y'_1)|y_1,z_1] + \mathrm{E}^{\mathcal{A}}[Z^o_{2,-n}(y'_1;\alpha,\xi)|y_1,z_1]}{2} \\
&= \frac{1-\beta_2\lambda_2}{\lambda_2(N+1)}\left([\widehat{v}(y_1;\alpha,\xi) - \lambda_1 y'_1] + \frac{N-1}{2}[\widehat{v}(y_1;\alpha,\xi) - \widehat{v}(y'_1;\alpha,\xi)]\right). \quad \text{(A27)}
\end{aligned}
$$

We add a few useful notations:

$$
\begin{aligned}
&\Delta\widehat{v} \coloneqq \widehat{v}(\tilde{y}'_1;\alpha,\xi) - \widehat{v}(\tilde{y}_1;\alpha,\xi), \quad \Delta P_1 \coloneqq \lambda_1(\tilde{y}'_1 - \tilde{y}_1) = \lambda_1 z_1, \quad \Delta P_2 \coloneqq \tilde{p}_2(\mathbf{Z}') - \tilde{p}_2(\mathbf{Z}) \\
&\Delta Z \coloneqq Z'_{2,n}(y'_1;\alpha,\xi) - Z^o_{2,n}(y_1;\alpha,\xi) = -\frac{(1-\beta_2\lambda_2)\lambda_1 z_1}{\lambda_2(N+1)} - \frac{(1-\beta_2\lambda_2)(N-1)}{2(N+1)\lambda_2}\Delta\widehat{v}.
\end{aligned}
$$

Note that $\mathbf{Z}' \coloneqq [\langle 0, Z^o_{2,1}\rangle, ...\langle Z'_{1,n}, Z'_{2,n}\rangle, ...\langle 0, Z^o_{2,N}\rangle]$ differs from $\mathbf{Z} \coloneqq [\langle 0, Z^o_{2,1}\rangle, ...\langle 0, Z^o_{2,N}\rangle]$ only in the n-th element $(\mathbf{Z}')_n = \langle Z'_{1,n}, Z'_{2,n}\rangle$. We can derive the following result:

$$
\Delta P_2 = \lambda_1 z + \lambda_2[\Delta Z + \beta_2(\tilde{v} - \lambda_1\tilde{y}'_1) - \beta_2(\tilde{v} - \lambda_1\tilde{y}_1) + Z^o_{2,-n}(\tilde{y}'_1) - Z^o_{2,-n}(\tilde{y}_1)] = -\lambda_2\Delta Z.
$$

Note that $\mathrm{E}^{\mathcal{A}}[\tilde{y}_1 z_1] = 0$ and $\mathrm{E}^{\mathcal{A}}[\widehat{v}(\tilde{y}_1;\xi)z_1] = 0$ because $\tilde{y}_1 = X_1(\tilde{v}) + \tilde{u}_1$ is symmetrically distributed and $\widehat{v}(-y_1;\xi) = -\widehat{v}(y_1;\xi)$. The extra payoff from this unilateral deviation is

$$
\begin{aligned}
\Delta\Pi^d_{z,n} &= \mathrm{E}^{\mathcal{A}}[(\tilde{v} - \tilde{p}_2(\mathbf{Z}'))Z'_{2,n} + (\tilde{v} - \tilde{p}_1(\mathbf{Z}'))z_1 - (\tilde{v} - \tilde{p}_2(\mathbf{Z}))Z^o_{2,n}] \\
&= -\lambda_1 z_1^2 + \mathrm{E}^{\mathcal{A}}[\mathrm{E}^{\mathcal{A}}[(\tilde{v} - \tilde{p}_2(\mathbf{Z}) + \lambda_2 Z'_{2,n})\Delta Z|\tilde{y}_1]] \\
&= -\lambda_1 z_1^2 + \mathrm{E}^{\mathcal{A}}\left[\left(\frac{\widehat{v}(\tilde{y}_1;\alpha,\xi) - \lambda_1\tilde{y}_1}{N+1} + \lambda_2\Delta Z\right)\cdot\Delta Z\right] \\
&= -\lambda_1 z_1^2 + (1-\beta_2\lambda_2)^2\frac{\mathrm{E}^{\mathcal{A}}\left[\left(\lambda_1 z_1 + \frac{N-1}{2}\Delta\widehat{v}\right)^2\right]}{(N+1)^2\lambda_2} \\
&\quad -(1-\beta_2\lambda_2)\frac{(N-1)\mathrm{E}^{\mathcal{A}}\left[(\widehat{v}(\tilde{y}_1;\alpha,\xi) - \lambda_1\tilde{y}_1)\Delta\widehat{v}\right]}{2(N+1)^2\lambda_2}. \quad \text{(A28)}
\end{aligned}
$$

It is useful to prove that $\mathrm{E}^{\mathcal{A}}\left[(\widehat{v}(\tilde{y}_1;\alpha,\xi)-\lambda_1\tilde{y}_1)(\Delta\widehat{v}-[\alpha/\beta_1+(1-\alpha)\lambda_1]z_1)\right]\geq 0$. First, we consider $z_1\geq 0$, under which $\Delta\widehat{v}(y_1,z_1)\leq\lim_{|y_1|\to\infty}\Delta\widehat{v}(y_1,z_1)=[\alpha/\beta_1+(1-\alpha)\lambda_1]z_1$ and the equality holds at infinity. This follows from Eq. (A26) and that $\widehat{v}$ is always convex for $y_1\geq 0$ and concave for $y_1\leq 0$. Second, given $\widehat{v}(-y_1;\alpha,\xi)=-\widehat{v}(y_1;\alpha,\xi)$, it follows that $\Delta\widehat{v}(y_1-\frac{z_1}{2},z_1)=\widehat{v}(y_1+\frac{z_1}{2};\alpha,\xi)-\widehat{v}(y_1-\frac{z_1}{2};\alpha,\xi)$ is an even function of y_1 and its minimum is achieved at $y_1=0$. This means $\Delta\widehat{v}(y_1,z_1)$ is shifted to the left and its minimum locates at $y_1=-\frac{z_1}{2}$. Hence, $\Delta\widehat{v}(y_1,z_1)-[\alpha/\beta_1+(1-\alpha)\lambda_1]z_1$ is non-positive and its minimum locates at $y_1=-\frac{z_1}{2}\leq 0$. Given the distributional symmetry of $\tilde{y}_1$ and the fact that $\widehat{v}(y_1;\alpha,\xi)-\lambda_1 y_1$ is an odd function, it follows that $\mathrm{E}^{\mathcal{A}}\left[(\widehat{v}(\tilde{y}_1;\alpha,\xi)-\lambda_1\tilde{y}_1)(\Delta\widehat{v}-[\alpha/\beta_1+(1-\alpha)\lambda_1]z_1)\right]\geq 0$. The same inequality holds for $z_1\leq 0$ under which $\Delta\widehat{v}(y_1,z_1)-[\alpha/\beta_1+(1-\alpha)\lambda_1]z_1\geq 0$, with its maximum located at $y_1=-\frac{z_1}{2}\geq 0$. Based on the obvious result (by symmetry) that $\mathrm{E}^{\mathcal{A}}\left[(\widehat{v}(\tilde{y}_1;\alpha,\xi)-\lambda_1\tilde{y}_1)z_1\right]=0$, the previous inequality implies that

$$
\begin{aligned}
\Delta\Pi^d_{z,n} &\leq -\lambda_1 z_1^2+(1-\beta_2\lambda_2)^2\frac{\mathrm{E}^{\mathcal{A}}\left[\left(\lambda_1 z_1+\frac{N-1}{2}\Delta\widehat{v}\right)^2\right]}{(N+1)^2\lambda_2} \\
&\quad -(1-\beta_2\lambda_2)\frac{(N-1)[\alpha/\beta_1+(1-\alpha)\lambda_1]}{2(N+1)^2\lambda_2}\mathrm{E}^{\mathcal{A}}\left[(\widehat{v}(\tilde{y}_1;\alpha,\xi)-\lambda_1\tilde{y}_1)z_1\right] \\
&= -\lambda_1 z_1^2+(1-\beta_2\lambda_2)^2\frac{\mathrm{E}^{\mathcal{A}}\left[(2\lambda_1 z_1+(N-1)\Delta\widehat{v})^2\right]}{4(N+1)^2\lambda_2}. \qquad \text{(A29)}
\end{aligned}
$$

As $0\leq\Delta\widehat{v}(y_1,z_1)\leq(\alpha/\beta_1+(1-\alpha)\lambda_1)z_1$ for $z_1\geq 0$ and $(\alpha/\beta_1+(1-\alpha)\lambda_1)z_1\leq\Delta\widehat{v}(y_1,z_1)\leq 0$ for $z_1\leq 0$, the last expression of (A29) has an upper bound which is achieved when $\Delta\widehat{v}=(\alpha/\beta_1+(1-\alpha)\lambda_1)z_1$. This yields

$$
\Delta\Pi^d_{z,n}(z_1;\alpha,\xi)\leq -\lambda_1 z_1^2+(1-\beta_2\lambda_2)^2\lambda_1 z_1^2\frac{\left[2+(N-1)(\alpha/(\beta_1\lambda_1)+1-\alpha)\right]^2}{4\delta(N+1)^2} \qquad \text{(A30)}
$$

Intuitively, the maximal benefit for the n-th trader (who unilaterally deviates) is achieved when arbitrageurs have the same extreme prior $\xi\to\infty$. In this limit, their reactions to the past order flows become the strongest and exactly linear:

$$
Z^o_{2,n}(y_1;\alpha,\xi\to\infty)=\frac{\alpha(1-\beta_1\lambda_1)(1-\beta_2\lambda_2)}{\beta_1\lambda_2}\frac{y_1}{N+1}. \qquad \text{(A31)}
$$

It is easy to verify that $\lim_{\xi\to\infty}\Delta\widehat{v}(\tilde{y}_1,z_1;\alpha,\xi)=\alpha z_1/\beta_1+(1-\alpha)\lambda_1 z_1$ such that

$$
\lim_{\xi\to\infty}\Delta\Pi^d_{z,n}(z_1;\alpha,\xi)=-\lambda_1 z_1^2+(1-\beta_2\lambda_2)^2\lambda_1 z_1^2\frac{[2+(N-1)(\alpha/(\beta_1\lambda_1)+1-\alpha)]^2}{4\delta(N+1)^2} \qquad \text{(A32)}
$$

which is exactly the right-hand side of (A30). Thus, $\Delta\Pi^d_{z,n} < 0$ holds if the above coefficient in front of z_1^2 is negative. This leads to the equilibrium existence condition (22), i.e.,

$$1 + \frac{\alpha(1-\beta_1\lambda_1)}{\beta_1\lambda_1} \cdot \frac{N-1}{N+1} < \frac{2\sqrt{\lambda_2/\lambda_1}}{1-\beta_2\lambda_2}. \tag{A33}$$

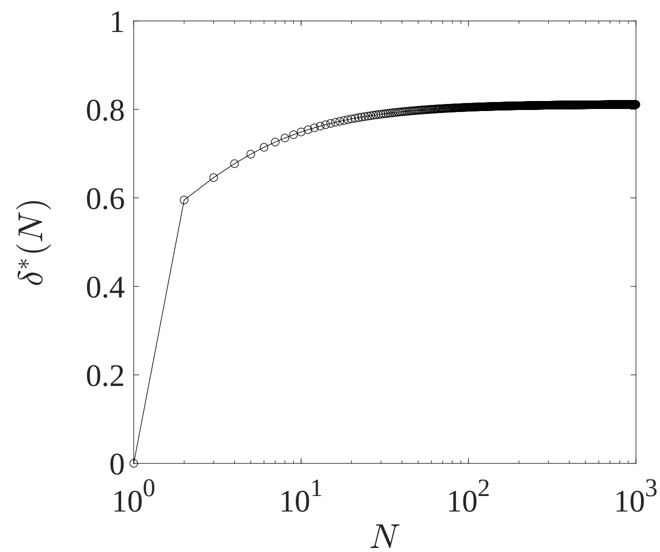

Figure 9. The critical value $\delta^*(N)$ as a function of the number of arbitrageurs N. Here, λ_t and β_t are determined in a two-period Kyle model ($M = 1$).

Consider the special case $M = 1$ in the example of microfoundation in Appendix A.1. This corresponds to a two-period a Kyle model with a single informed trader. When $\alpha = 1$, the inequality (22) lead to the condition $\delta > \delta^*(N)$ where $\delta := \lambda_2/\lambda_1$ is the ratio of Kyle lambdas and $\delta^*(N)$ is the largest root of the nonlinear equation:

$$1 + \left(\frac{N-1}{N+1}\right)\frac{2\delta}{2\delta-1} = 4\sqrt{\delta}. \tag{A34}$$

We find that $\delta^*(N{=}2) \approx 0.5951$, $\delta^*(N{=}3) \approx 0.6458$, and $\delta^*(N{=}10) \approx 0.7489$. Moreover, $\lim_{N\to\infty}\delta^*(N) \approx 0.8117$, which is the largest root of the equation: $64\delta^3 - 80\delta^2 + 24\delta - 1 = 0$. The critical values of $\delta^*(N)$ are plotted in Figure 9. The equilibrium ratio δ can vary with the ratio of noise trading volatilites (γ). In the liquidity regime $\delta > \delta^*(N)$, it is unprofitable for any arbitrageur to trade in the first period, that is, $\Delta\Pi^d_{z,n}(z_1;\xi) < 0$ for $z_1 \neq 0$. This confirms our conjecture that arbitrageurs will not trade at $t = 1$. When $\delta > \delta_\infty \approx 0.8117$, the equilibrium can host an infinite number of arbitrageurs.

A.4 Proof of Corollary 4.1

(a) follows from the property that $\widehat{v}(y_1;\alpha,\xi)$ is an odd function: $\widehat{v}(-y_1) = -\widehat{v}(y_1)$.
(b) follows from the property that $\widehat{v}(y_1;\alpha,\xi)$ is convex for $y_1 \geq 0$ and concave for $y_1 \leq 0$.
(c) follows from the result that $\lim_{|y_1|\to\infty} \widehat{v}(y_1) = (\alpha/\beta_1)[y_1 - \text{sign}(y_1)\kappa\sigma_u] + (1-\alpha)\lambda_1 y_1$.

(d) follows from the condition $\frac{\partial}{\partial y_1} Z^o_{2,n}(y_1;\alpha,\xi=\xi_c)\big|_{y_1=0} = 0$, which is equivalent to

$$1 + \kappa(\xi_c)^2 - \frac{\kappa(\xi_c)e^{-\kappa(\xi_c)^2/2}}{\text{erfc}(\kappa(\xi_c)/\sqrt{2})}\sqrt{\frac{2}{\pi}} = \lambda_1\beta_1 \tag{A35}$$

The left-hand side of (A35) is a monotonic function of κ, which decreases from 1 to 0 when κ increases from 0 to ∞. The right-hand side, $\lambda_1\beta_1$, is a constant that takes a value between 0 and 1. Hence, Eq. (A35) admits a unique positive solution $\kappa^* > 0$. Since $\kappa \coloneqq \sigma_u/(\beta_1\xi)$, we can find the unique solution (Figure 10), $\xi_c = \sigma_u/(\beta_1\kappa^*)$, which depends on σ_u and σ_v.

From the curvature property of $Z^o_{2,n}(y_1;\alpha,\xi)$, we have that $\frac{\partial^2 Z^o_{2,n}}{\partial y_1^2} > 0$ for $y_1 \geq 0$ and $\frac{\partial^2 Z^o_{2,n}}{\partial y_1^2} < 0$ for $y_1 \leq 0$. Therefore, when $\xi \geq \xi_c$, we have $\frac{\partial Z^o_{2,n}(y_1;\alpha,\xi)}{\partial y_1} \geq 0$, showing that $Z^o_{2,n}(y_1;\alpha,\xi)$ is an increasing function of y_1 which has only one root at $y_1 = 0$. In contrast, $Z^o_{2,n}(y_1;\alpha,\xi)$ becomes a non-monotonic function when $\xi < \xi_c$ and has three different roots.

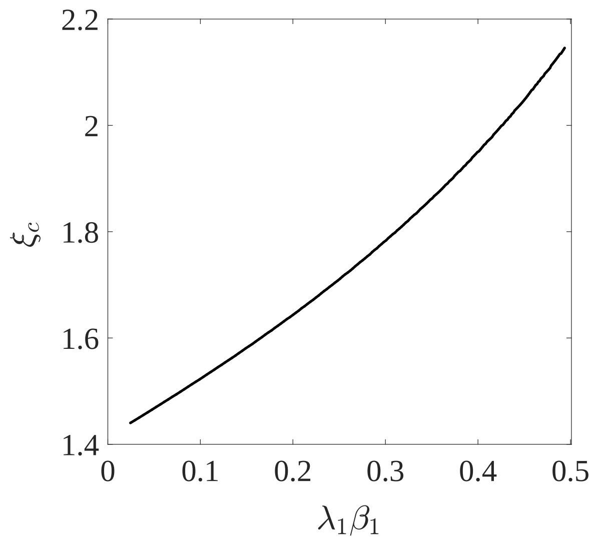

Figure 10. The critical value ξ_c as a function of the value of $\lambda_1\beta_1$, where λ_1 and β_1 are determined in a two-period Kyle model.

A.5 Proof of Corollary 4.2

First, when $\xi_L < \xi_c \le \xi_H$, arbitrageurs' best response at $t = 2$ is to stop trading on any $y_1 \in [-K_L, K_L]$ where K_L is the positive root of the equation $Z^o_{2,n}(y_1; \alpha, \xi_L) = 0$. If other arbitrageurs do not trade on $y_1 \in [-K_L, K_L]$, any individual arbitrageur would not deviate because buying or selling this asset could lose money in case when the true prior turns out to be on the opposite side of such trading. For any $|y_1| > K_L$, the max-min strategy simply follows $Z^o_{2,n}(y_1; \alpha, \xi_L)$. This is because any unilateral deviation from this most conservative strategy may lose money in case that the reality happens to be the lowest prior ξ_L. Thus, no one would trade more than the most conservative strategy $Z^o_{2,n}(y_1; \alpha, \xi_L)\mathbf{1}_{|y_1|>K_L}$.

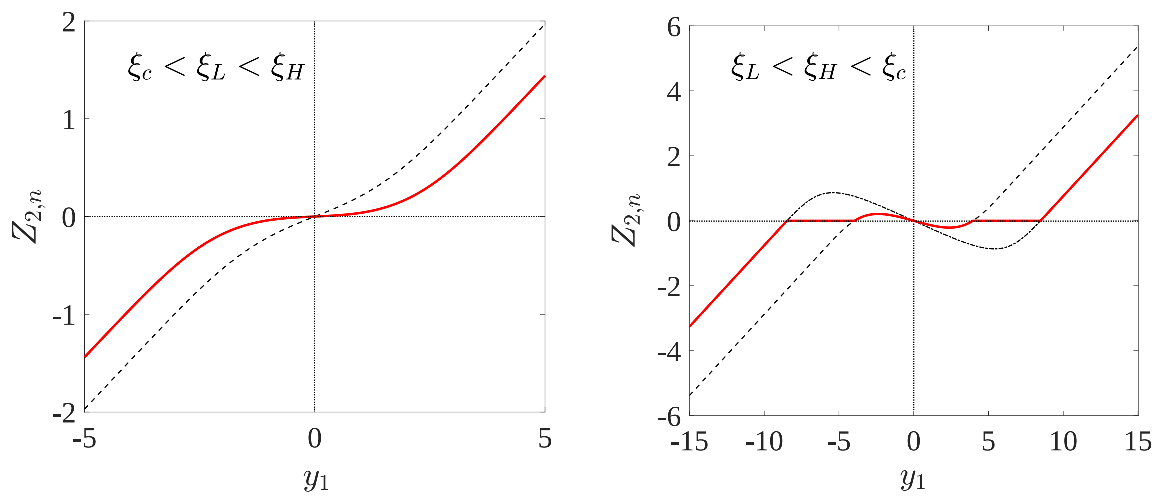

Figure 11. The equilibrium max-min strategies when $\xi_c < \xi_L < \xi_H$ and $\xi_L < \xi_H < \xi_c$.

Second, when $\xi_c \le \xi_L < \xi_H$, the two extreme strategies $Z^o_{2,n}(y_1; \alpha, \xi_H)$ and $Z^o_{2,n}(y_1; \alpha, \xi_L)$ agree on the trading direction for all realized y_1. Thus, the max-min strategy is $Z^o_{2,n}(y_1; \alpha, \xi_L)$, as shown by the red solid line in the left panel of Figure 11.

Last, when $\xi_L < \xi_H < \xi_c$, both $Z^o_{2,n}(y_1; \alpha, \xi_L)$ and $Z^o_{2,n}(y_1; \alpha, \xi_H)$ are non-monotonic, and each has three roots. The max-min strategy is $Z^o_{2,n}(y_1; \alpha, \xi_H)\mathbf{1}_{|y_1|<K_H} + Z^o_{2,n}(y_1; \alpha, \xi_L)\mathbf{1}_{|y_1|>K_L}$, where K_H represents the positive root of the equation $Z^o_{2,n}(y_1; \alpha, \xi_H) = 0$. Since $K_H < K_L$, this strategy has two no-trade zones, $[-K_L, -K_H]$ and $[K_H, K_L]$, and three trading zones, $[-K_H, K_H]$, $[-\infty, -K_L]$, and $[K_L, \infty]$. The solution is shown by the red solid line in the right panel of Figure 11. It is too complicated to be practically relevant.

Note that K_L and $K(\xi_L)$ are very close but not identical. K_L is the positive point where $Z_{2,n}(y_1; \alpha, K(\xi_L))$ crosses the horizontal y_1-axis, whereas $K(\xi_L)$ is the horizontal intercept of the asymptote of $Z_{2,n}(y_1; \alpha, K(\xi_L))$. The difference between K_H and $K(\xi_H)$ is similar.

A.6 Proof of Theorem 2

First, we verify that, under (C1), (C2), and (C3), each arbitrageur will not deviate from the strategy (27) at $t = 2$. By symmetry, it suffices to consider the positive domain.

For any realized order flow $y_1 \in [0, K(\xi_w)]$, the n-th arbitrageur will not deviate to buy any share of this asset, because choosing $Z'_{2,n} > 0$ may lose money under the lowest prior ξ_L:

$$\begin{aligned} \mathrm{E}^{\mathcal{A}}[\Delta\tilde{\pi}_{z,n}|y_1, \tilde{\xi} = \xi_L] &= \mathrm{E}^{\mathcal{A}}\left[\left(\tilde{v} - \lambda_1 y_1 - \lambda_2[\beta_2(\tilde{v} - \lambda_1 y_1) + Z'_{2,n} + \tilde{u}_2]\right) Z'_{2,n}|y_1, \tilde{\xi} = \xi_L\right] - 0 \\ &= (1 - \beta_2\lambda_2)\widehat{\theta}(y_1; \alpha, \xi_L) Z'_{2,n} - \lambda_2 Z'^2_{2,n} < 0. \end{aligned} \tag{A36}$$

The inequality is due to $(1-\beta_2\lambda_2)\widehat{\theta}(y_1; \alpha, \xi_L)=\lambda_2(N+1)Z_{2,n}(y_1; \alpha, \xi_L) < 0$ for $y_1 \in [0, K(\xi_w)]$, which is implied by (C1) $K(\xi_w) < K_L$, (C3) $\xi_L < \xi_c$, and Corollary 4.1(d). Similarly, each arbitrageur would not deviate by choosing any $Z'_{2,n} < 0$ as it may lose money under ξ_H:

$$\mathrm{E}^{\mathcal{A}}[\Delta\tilde{\pi}_{z,n}|y_1, \tilde{\xi} = \xi_H] = (1 - \beta_2\lambda_2)\widehat{\theta}(y_1; \alpha, \xi_H) Z'_{2,n} - \lambda_2 Z'^2_{2,n} < 0. \tag{A37}$$

The inequality is due to $(1 - \beta_2\lambda_2)\widehat{\theta}(y_1; \alpha, \xi_H)=\lambda_2(N + 1)Z_{2,n}(y_1; \alpha, \xi_H) > 0$ since $\xi_H > \xi_c$. By the max-min criterion, arbitrageurs will not deviate from no trading for $y_1 \in [0, K(\xi_w)]$.

For $y_1 \in (K(\xi_w), \infty)$, each arbitrageur will not trade less than the amount of $Z^\infty(y_1; \alpha, \xi_w)$, because doing this would violate either (C1) or (C2) or both. Each arbitrageur will not trade more that $Z^\infty(y_1; \alpha, \xi_w)$ either, because choosing any $Z'_{2,n}(y_1) > Z^\infty(y_1; \alpha, \xi_w)$ may lose more or earn less than the buying decision made along $Z^\infty(y_1; \alpha, \xi_w)$. Define $Z_\Delta \coloneqq Z'_{2,n}(y_1) - Z^\infty(y_1; \alpha, \xi_w)$. The difference of payoffs from this unilateral deviation under ξ_L is

$$\begin{aligned} \mathrm{E}^{\mathcal{A}}[\Delta\tilde{\pi}_{z,n}|y_1, \xi_L] &= \mathrm{E}^{\mathcal{A}}\left[\left(\tilde{v} - \lambda_1 y_1 - \lambda_2[\beta_2(\tilde{v} - \lambda_1 y_1) + Z'_{2,n} + Z_{2,-n} + \tilde{u}_2]\right) Z'_{2,n}|y_1, \xi_L\right] \\ &\quad -\mathrm{E}^{\mathcal{A}}\left[(\tilde{v} - \lambda_1 y_1 - \lambda_2[\beta_2(\tilde{v} - \lambda_1 y_1) + Z^\infty + Z_{2,-n} + \tilde{u}_2]) Z^\infty|y_1, \xi_L\right] \\ &= \mathrm{E}^{\mathcal{A}}\left[Z_\Delta[(1 - \beta_2\lambda_2)\tilde{\theta} - \lambda_2(Z^\infty + Z_{2,-n})]\big|y_1, \xi_L\right] - \lambda_2 Z'_{2,n} Z_\Delta. \end{aligned} \tag{A38}$$

Let $Z_L \coloneqq (1 - \beta_2\lambda_2)\frac{\widehat{\theta}(y_1;\alpha,\xi_L)}{\lambda_2(N+1)}$. Since everyone else follows Z^∞ and $Z_L < Z^\infty < Z'_{2,n}$, we have

$$\begin{aligned} \mathrm{E}^{\mathcal{A}}[\Delta\tilde{\pi}_{z,n}|y_1, \xi_L] &= \lambda_2 Z_\Delta[(N + 1)Z_L - Z^\infty - (N - 1)Z^\infty] - \lambda_2 Z'_{2,n} Z_\Delta \\ &= \lambda_2 Z_\Delta[(N + 1)Z_L - NZ^\infty - Z'_{2,n}] < 0. \end{aligned} \tag{A39}$$

Thus, each arbitrageur will only trade $Z^\infty(y_1; \alpha, \xi_w)$ when $y_1 > K(\xi_w)$. By symmetry, the full equilibrium strategy is $Z^\infty(y_1; \alpha, \xi_w)\mathbf{1}_{|y_1|>K(\xi_w)}$. No one will deviate from it at $t = 2$.

The equilibrium condition (C1) is central to Theorem 2. This is endogenously implied by the existence of a debiased equilibrium. If (C1) is absent, (C3) alone is sufficient to direct the economy to the equilibrium in Corollary 4.2 and the presence of (C2) makes no difference. When (C1) holds, (C2) plays an important role in regularizing the robust optimization problem. Without (C2), the other conditions (C1) and (C3) cannot support the equilibrium in Theorem 2. This is because if candidate strategies were allowed to be non-convex in the positive domain of y_1 or non-concave in the negative domain, then for any strategy that satisfies (C1), arbitrageurs would always find some deviations that trade more conservatively than that strategy. Such deviations are permitted by the gap between $Z^{\infty}(y_1;\alpha,\xi_w)$ and $Z^{\infty}(y_1;\alpha,\xi_L)$, which is an implication of (C1).

Under (C1) and (C2), all admissible strategies must be within the shaded area in Figure 2, enclosed by $Z^o_{2,n}(y_1;\alpha,\xi\to 0)$, $Z^o_{2,n}(y_1;\alpha,\xi\to\infty)$, and $Z^{\infty}(y_1;\alpha,\xi_w)$. This can be understood by just looking at the positive domain of y_1. It is apparently irrational for any strategy to go above $Z^o_{2,n}(y_1;\alpha,\xi\to\infty)$ or below $Z^o_{2,n}(y_1;\alpha,\xi\to 0)$. Moreover, (C2) means that the first derivative of any admissible strategy is non-decreasing in the positive domain of y_1. Such a strategy cannot cross the asymptote $Z^{\infty}(y_1;\alpha,\xi_w)$ without violating (C1).

It remains to verify that no arbitrageur would find it utility-improving to trade at $t=1$, given that other arbitrageurs only trade at $t=2$ using the strategy (27). The proof here is similar to A.2 for Theorem 1. Intuitively, each arbitrageur would not trade at $t=1$ since it would risk trading on the opposite side of the true fat-tail signal. But it might be profitable if other arbitrageurs were overly misled by the secret "disruptive" trading. Suppose the n-th arbitrageur is an instigator who considers to trade $Z'_{1,n}=z_1\neq 0$, in an effort to confuse other arbitrageurs (now momentum traders). To save notations, we will use K to represent $K(\xi_w)$. Given her market power and unilateral deviation, the instigator understands the composition of order flows: $\tilde{y}'_1=\beta_1(\tilde{v}-p_0)+z_1+\tilde{u}_1$ and $\tilde{y}'_2=\beta_2(\tilde{v}-\lambda_1\tilde{y}'_1)+Z'_{2,n}(\tilde{y}_1,z_1)+Z_{2,-n}(\tilde{y}'_1;\alpha,K)+\tilde{u}_2$. Here, $Z_{2,-n}\coloneqq\sum_{n'\neq n}Z_{2,n'}(y'_1;\alpha,K)$ is the total order flow placed by other arbitrageurs who will estimate $\tilde{\theta}=\tilde{v}-\lambda_1 y'_1$ based on the observed y'_1 without knowing that y'_1 contains an uninformed order flow z_1 from the instigator. Of course, the instigator's estimate of $\tilde{\theta}$ is correctly based on $y_1=\beta_1(\tilde{v}-p_0)+u_1$ instead of y'_1, because she knows the order z_1 placed by herself. Being averse to the model risk, the instigator has the objective function:

$$\max_{z_2}\min_{\xi}\mathrm{E}^{\mathcal{A}}[(\tilde{v}-\lambda_1\tilde{y}'_1-\lambda_2\tilde{y}'_2)z_2|y_1,z_1,\alpha,\xi],\tag{A40}$$

The instigator can exploit the possibility that other arbitrageurs are (unknowingly) biased by her uninformed trade z_1 since they follow the strategy $Z_{2,n'}(y'_1=y_1+z_1;\alpha,K)$ for each $n'\neq n$. For a given prior $\xi\in[\xi_L,\xi_H]$, the instigator's optimal strategy at $t=2$ is

$$Z^o_{2,n}(y_1,z_1;\alpha,\xi)=(1-\beta_2\lambda_2)\frac{\mathrm{E}^{\mathcal{A}}[\tilde{v}|y_1,\alpha,\xi]-\lambda_1(y_1+z_1)}{2\lambda_2}-\frac{1}{2}\sum_{n'\neq n}Z_{2,n'}(y_1+z_1;\alpha,K).\tag{A41}$$

Since the last term is independent of her prior ξ, it will remain in her max-min strategy:

$$Z'_{2,n}(y_1, z_1) = \arg\max_{z_2} \min_{\xi} \mathrm{E}^{\mathcal{A}}\left[(\tilde{v} - \lambda_1 y'_1 - \lambda_2(y'_1 + z_2))z_2 | y_1, z_1, \alpha, \xi\right] - \frac{1}{2}\sum_{n' \neq n} Z_{2,n'}(y'_1; \alpha, K).$$

The first term is actually the problem we solved earlier as if there was only one arbitrageur and the first-period order flow was y'_1 instead of y_1. Nonetheless, the instigator's estimate of $\tilde{v}$ is correctly based on y_1 (not y'_1). By deducting her own price impact $\lambda_1 z_1$ from this problem, she can find that the max-min solution to the remaining problem is simply

$$Z_{2,n}(y_1; \alpha, K)\big|_{N=1} = \frac{\alpha(1-\beta_1\lambda_1)(1-\beta_2\lambda_2)}{2\beta_1\lambda_2}\left[y_1 - \mathrm{sign}(y_1)K\right]\mathbf{1}_{|y_1|>K} = \frac{N+1}{2} Z_{2,n}(y_1; \alpha, K).$$

Taking into account all those results, she will find her max-min strategy given by

$$\begin{aligned} Z'_{2,n}(y_1, z_1) &= \frac{(N+1)}{2} Z_{2,n}(y_1; \alpha, K) - (1-\beta_2\lambda_2)\frac{\lambda_1 z_1}{2\lambda_2} - \frac{N-1}{2} Z_{2,n}(y'_1; \alpha, K) \\ &= Z_{2,n}(y_1; \alpha, K) - (1-\beta_2\lambda_2)\frac{z_1}{2\delta} - \frac{\alpha(1-\beta_1\lambda_1)(1-\beta_2\lambda_2)(N-1)}{2\beta_1\lambda_2(N+1)} D(y_1, z_1), \end{aligned} \tag{A42}$$

where we define the difference between two soft-thresholding functions:

$$D(y_1, z_1) \coloneqq [y_1 + z_1 - \mathrm{sign}(y_1 + z_1)K]\mathbf{1}_{|y_1+z_1|>K} - [y_1 - \mathrm{sign}(y_1)K]\mathbf{1}_{|y_1|>K}. \tag{A43}$$

Ex ante, the expected trading profit of this instigator (the n-th arbitrageur) is

$$\Pi^d_{z,n}(z_1) = \mathrm{E}^{\mathcal{A}}[(\tilde{v} - \lambda_1\tilde{y}'_1)z_1 + (\tilde{v} - \lambda_1\tilde{y}'_1 - \lambda_2\tilde{y}'_2) \cdot Z'_{2,n}(\tilde{y}_1, z_1)], \tag{A44}$$

and hence the extra profit attributable to her unilateral deviation $\langle Z'_{1,n}, Z'_{2,n}\rangle$ is

$$\Delta\Pi^d_{z,n}(z_1) = \Pi^d_{z,n}(z_1) - \mathrm{E}^{\mathcal{A}}[(\tilde{v} - \lambda_1\tilde{y}_1 - \lambda_2\tilde{y}_2) \cdot Z_{2,n}(\tilde{y}_1; \alpha, K)], \tag{A45}$$

where $\tilde{y}_1 = \beta_1(\tilde{v} - p_0) + \tilde{u}_1$ and $\tilde{y}_2 = \beta_2(\tilde{v} - \lambda_1\tilde{y}_1) + \sum_{n=1}^N Z_{2,n}(\tilde{y}_1; \alpha, K) + \tilde{u}_2$. We can derive

$$\begin{aligned} \Delta\Pi^d_{z,n} &= -\lambda_1 z_1^2 + \lambda_2\mathrm{E}^{\mathcal{A}}[(Z'_{2,n}(\tilde{y}_1, z_1))^2] - \lambda_2\mathrm{E}^{\mathcal{A}}[(Z_{2,n}(\tilde{y}_1; \alpha, K))^2] \\ &= -\lambda_1 z_1^2 + \lambda_2\mathrm{E}^{\mathcal{A}}\Bigg[\left((1-\beta_2\lambda_2)\frac{z_1}{2\delta} + \frac{\alpha(1-\beta_1\lambda_1)(1-\beta_2\lambda_2)(N-1)}{2\beta_1\lambda_2(N+1)} D\right)^2 \\ &\quad - \frac{\alpha(1-\beta_1\lambda_1)(1-\beta_2\lambda_2)(N-1)}{\beta_1\lambda_2(N+1)} Z_{2,n} D\Bigg]. \end{aligned} \tag{A46}$$

By symmetry, $\Delta\Pi^d_{z,n}(-z_1) = \Delta\Pi^d_{z,n}(z_1)$. We can simply focus on the case of $z_1 > 0$. It takes some straightforward but tedious calculation to arrive at the following result (when $z_1 > 0$):

$$Z_{2,n}(y_1;\alpha,K)\cdot(D-z_1) = \begin{cases} \max\{A(y_1,z_1),0\} & \text{if } z_1 \leq 2K \\ \max\{A(y_1,z_1)\mathbf{1}_{y_1<K-z_1} + B(y_1,z_1)\mathbf{1}_{y_1\geq K-z_1},0\} & \text{if } z_1 > 2K \end{cases}$$

where $A(y_1,z_1) \coloneqq -\frac{\alpha(1-\beta_1\lambda_1)(1-\beta_2\lambda_2)}{\beta_1\lambda_2(N+1)}(y_1+K)(y_1+z_1+K)$ is a quadratic function of y_1 for any given z_1 and $B(y_1,z_1) \coloneqq A(K-z_1,z_1)\frac{y_1+K}{2K-z_1}$ is a linear function of y_1 for any given z_1. The above result proves that $Z_{2,n}(\tilde{y}_1;\alpha,K)D(\tilde{y}_1,z_1) \geq Z_{2,n}(\tilde{y}_1;\alpha,K)z_1$, which further implies

$$\begin{aligned} \Delta\Pi^d_{z,n} &\leq -\lambda_1 z_1^2 + \lambda_2 \mathrm{E}^{\mathcal{A}}\left[\left((1-\beta_2\lambda_2)\frac{z_1}{2\delta} + \frac{\alpha(1-\beta_1\lambda_1)(1-\beta_2\lambda_2)(N-1)}{2\beta_1\lambda_2(N+1)}D\right)^2\right. \\ &\leq -\lambda_1 z_1^2 + \lambda_2\left[\frac{1-\beta_2\lambda_2}{2\delta} + \frac{\alpha(1-\beta_1\lambda_1)(1-\beta_2\lambda_2)(N-1)}{2\beta_1\lambda_2(N+1)}\right]^2 z_1^2. \end{aligned} \tag{A47}$$

The second step is by the property that $0 \leq D \leq z_1$ if $z_1 \geq 0$ and $z_1 \leq D \leq 0$ if $z_1 \leq 0$ and $\mathrm{E}^{\mathcal{A}}[Z_{2,n}]z_1 = 0$. The equality in (A47) holds when $\xi_w \to \infty$ so that $K(\xi_w) \to 0$ and $D \to z_1$.

It is not a profitable deviation (i.e., $\Delta\Pi^d_{z,n} < 0$) if the coefficient of z_1^2 in (A47) is negative. This coefficient condition leads to the same equilibrium condition (22) in Proposition 1:

$$1 + \frac{\alpha(1-\beta_1\lambda_1)}{\beta_1\lambda_1}\cdot\frac{N-1}{N+1} < \frac{2\sqrt{\lambda_2/\lambda_1}}{1-\beta_2\lambda_2}, \tag{A48}$$

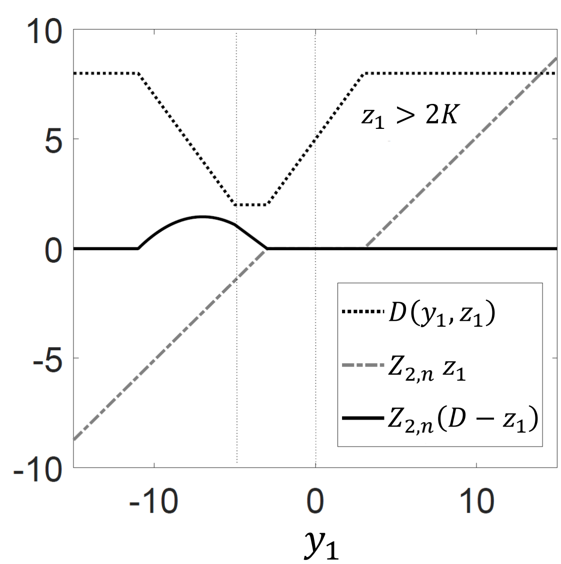

Figure 12. $D(y_1,z_1)$, $Z_{2,n}(y_1;\alpha,K)z_1$, and $Z_{2,n}(y_1;\alpha,K)\cdot(D-z_1)$ when $z_1 > 2K$.

A.7 Proof of Theorem 3

By definition, if a strategy can be written as a soft-thresholding function, then we may map it into a LASSO strategy which takes some LASSO estimates of economic variables. For $\alpha > 0$, the robust strategy (27) in Theorem 2 is a soft-thresholding function of y_1:

$$Z_{2,n}(y_1; \alpha, K(\xi_w)) = Z^{\infty}(y_1; \alpha, \xi_w)\mathbf{1}_{|y_1|>K(\xi_w)} = \frac{\alpha(1-\beta_1\lambda_1)(1-\beta_2\lambda_2)}{\beta_1\lambda_2(N+1)}\mathcal{S}(y_1; K(\xi_w)). \tag{A49}$$

Similar to the optimal strategy (18), the robust strategy (A49) can be written as

$$Z_{2,n}(y_1; \alpha, K(\xi_w)) = \frac{\alpha(1-\beta_2\lambda_2)}{\lambda_2(N+1)} \cdot \widehat{\theta}^{\text{imp}}(y_1; \xi_w), \tag{A50}$$

where $\widehat{\theta}^{\text{imp}} = \frac{1-\beta_1\lambda_1}{\beta_1}\mathcal{S}(y_1; K(\xi_w))$ is the strategy-implied estimator. It remains to show that $\widehat{\theta}^{\text{imp}}$ corresponds to some LASSO estimate of $\tilde{\theta}$. If we define

$$\widehat{\theta}^{\text{lasso}}(y_1; \xi_w) \coloneqq \arg\min_{\theta} \left\{ \frac{1}{2}\left| y_1 - \frac{\beta_1\theta}{1-\beta_1\lambda_1} \right|^2 + \frac{\sigma_u^2|\theta|}{(1-\beta_1\lambda_1)^2\xi_w} \right\}, \tag{A51}$$

then, using equations (2), (3), and (5), we can show that

$$\widehat{\theta}^{\text{lasso}} = \frac{1-\beta_1\lambda_1}{\beta_1}[y_1 - \text{sign}(y_1)K(\xi_w)]\mathbf{1}_{|y_1|>K(\xi_w)} = \frac{1-\beta_1\lambda_1}{\beta_1}\mathcal{S}(y_1; K(\xi_w)) = \widehat{\theta}^{\text{imp}}. \tag{A52}$$

This proves that the robust strategy (27) is a LASSO strategy since we have

$$Z_{2,n}(y_1; \alpha, K(\xi_w)) = \frac{\alpha(1-\beta_2\lambda_2)}{\lambda_2(N+1)} \cdot \widehat{\theta}^{\text{lasso}}(y_1; \xi_w). \tag{A53}$$

We can also define the LASSO estimate of $\tilde{v}$ as

$$\widehat{v}^{\text{lasso}}(y_1; \xi_w) \coloneqq \arg\min_{v} \left\{ \frac{1}{2}|y_1 - \beta_1 v|^2 + \frac{\sigma_u^2}{\xi_w}|v| \right\} = \frac{1}{\beta_1}\mathcal{S}(y_1; \kappa\sigma_u). \tag{A54}$$

Since $\kappa(\xi_w)\sigma_u < K(\xi_w) = \frac{\kappa(\xi_w)\sigma_u}{1-\beta_1\lambda_1}$, it follows that $\widehat{\theta}^{\text{lasso}} = (\widehat{v}^{\text{lasso}} - \lambda_1 y_1)\mathbf{1}_{|y_1|>K(\xi_w)}$ and thus

$$Z_{2,n}(y_1; \alpha, K(\xi_w)) = \frac{\alpha(1-\beta_2\lambda_2)(\widehat{v}^{\text{lasso}} - \lambda_1 y_1)\mathbf{1}_{|y_1|>K(\xi_w)}}{\lambda_2(N+1)}. \tag{A55}$$

Eq. (A53) and Eq. (A55) together prove Eq. (28) in Theorem 3.

A.8 Proof of Proposition 3

When $\alpha = 1$, a fixed Laplacian-Gaussian mixture prior $\mathcal{LG}(\alpha, \xi)$ reduces to a pure Laplace prior $\mathcal{L}(0, \xi)$. Conditional on this prior and the order flow $y_1 = \beta_1(\tilde{v} - p_0) + \tilde{u}_1$ with $p_0 = 0$, the posterior distribution $f(v|y_1)$ is equal to $f(y_1|v)f(v)/f(y_1)$ by Bayes' rule and hence the *maximum a posteriori* (MAP) estimate of $\tilde{v}$ is given by

$$\widehat{v}^{\text{map}} = \arg\max_{v} \ \exp\left[-\frac{(y_1 - \beta_1 v)^2}{2\sigma_u^2} - \frac{|v|}{\xi}\right] = \arg\min_{v} \left\{\frac{|y_1 - \beta_1 v|^2}{2} + \frac{\sigma_u^2}{\xi}|v|\right\}. \tag{A56}$$

When $\xi = \xi_w$, Eq. (A56) becomes the same LASSO objective function (29) that defines $\widehat{v}^{\text{lasso}}$. The first order condition of (A56) is $y_1(v) = \beta_1 v + \text{sign}(v)\kappa(\xi)\sigma_u$, where $\kappa(\xi) := \sigma_u/(\beta_1\xi)$. Inverting this function $y_1(v)$ yields the MAP estimator which has a learning threshold $\kappa\sigma_u$. When $\xi = \xi_w$, it coincides with the soft-thresholding expression of $\widehat{v}^{\text{lasso}}$ in Eq. (29):

$$\widehat{v}^{\text{map}}(y_1; \alpha = 1, \xi_w) = \frac{1}{\beta_1}\left[y_1 - \text{sign}(y_1)\kappa\sigma_u\right]\mathbf{1}_{|y_1| > \kappa\sigma_u} = \frac{1}{\beta_1}\mathcal{S}(y_1; \kappa\sigma_u) = \widehat{v}^{\text{lasso}}(y_1; \xi_w). \tag{A57}$$

(A56) and (A57) demonstrate the statistical interpretation of LASSO by Tibshirani (1996). Figure 13 plots the posterior distribution $f(v|y_1; \alpha = 1, \xi)$ for three values of y_1 with $\sigma_u = 1$. It illustrates the effect of a sharply peaked Laplace prior on suppressing nonzero estimates. When y_1 is at or below the learning threshold $\kappa\sigma_u$, the posterior remains sharply peaked at the origin such that $\widehat{v}^{\text{map}} = 0$. When y_1 exceeds the threshold $\kappa\sigma_u$, the posterior mode shifts to the right such that $\widehat{v}^{\text{map}} > 0$. The l_1 penalty term in LASSO has the same effect.

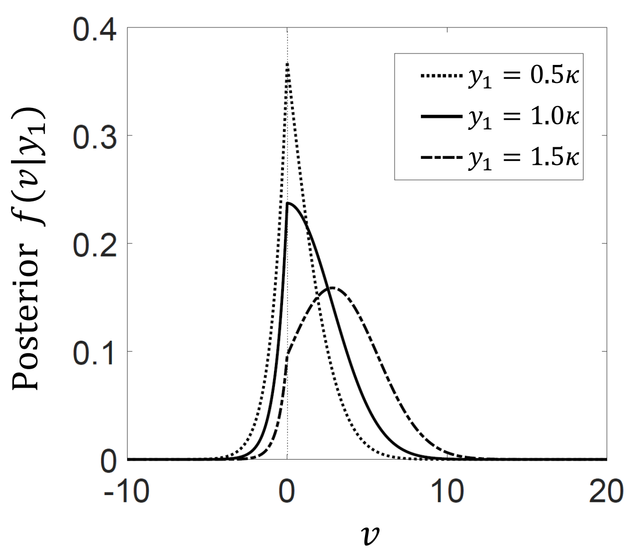

Figure 13. The posterior distributions $f(v|y_1; \alpha = 1, \xi)$ for $y_1 = 0.5\kappa$, $y_1 = \kappa$, and $y_1 = 1.5\kappa$.

Now consider an otherwise identical economy where arbitrageurs naively use the MAP rule to estimate $\tilde{v}$ based on a pure Laplace prior $\tilde{v} \sim \mathcal{L}(0, \xi_w)$. We can easily show that their heuristic feedback trading strategy coincides with the robust LASSO strategy (27):

$$
\begin{aligned}
& Z_{2,n}^{\text{map}}(y_1; \alpha = 1, \xi_w) = \frac{(1-\beta_2\lambda_2)}{\lambda_2(N+1)}[\widehat{v}^{\text{map}}(y_1; \alpha = 1, \xi_w) - \lambda_1 y_1]\mathbf{1}_{|y_1|>K(\xi_w)} \\
= \;& Z_{2,n}(y_1; \alpha = 1, K(\xi_w)) = \frac{(1-\beta_1\lambda_1)(1-\beta_2\lambda_2)}{\beta_1\lambda_2(N+1)}\left[y_1 - \text{sign}(y_1)K(\xi_w)\right]\mathbf{1}_{|y_1|>K(\xi_w)}, \quad \text{(A58)}
\end{aligned}
$$

where the term $\mathbf{1}_{|y_1|>K(\xi_w)}$ is imposed to guarantee the practice of positive feedback trading. Eq. (A58) demonstrates the observational equivalence between the heuristic MAP strategy $Z_{2,n}^{\text{map}}(y_1; \alpha = 1, \xi_w)$ and the robust LASSO strategy $Z_{2,n}(y_1; \alpha = 1, \xi_w)$.

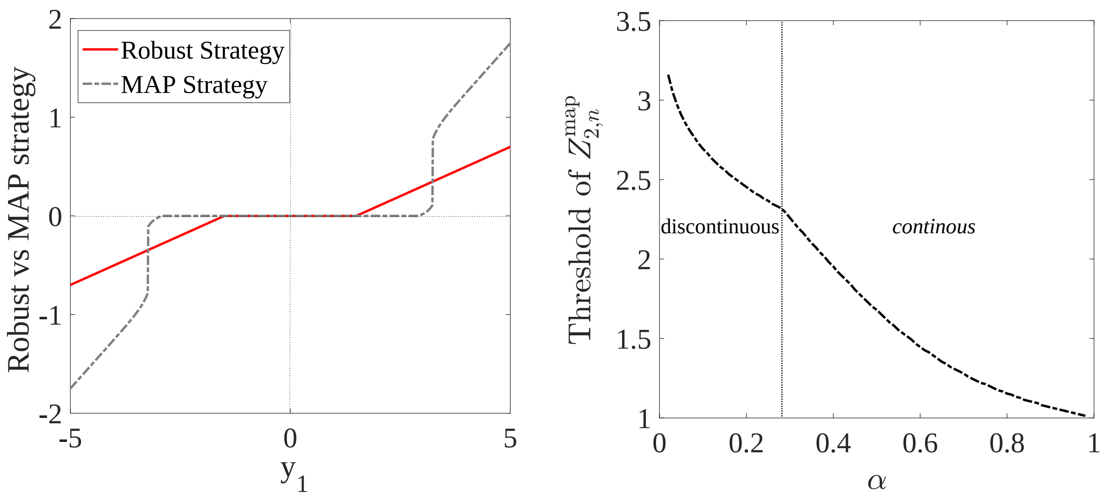

Figure 14. Left: the LASSO strategy $Z_{2,n}(y_1)$ and the MAP strategy $Z_{2,n}^{\text{map}}(y_1)$ when $\alpha = 0.5$. Right: the trading threshold of MAP strategy versus α.

For any $\alpha \in (0, 1)$, this observational equivalence fails because the MAP estimate becomes

$$
\widehat{v}^{\text{map}} := \arg\max_{v} \left[\frac{\alpha}{2\xi}\exp\left(-\frac{(y_1-\beta_1 v)^2}{2\sigma_u^2} - \frac{|v|}{\xi}\right) + \frac{1-\alpha}{\sqrt{2\pi\sigma_v^2}}\exp\left(-\frac{(y_1-\beta_1 v)^2}{2\sigma_u^2} - \frac{v^2}{2\sigma_v^2}\right)\right].
$$

which is always different from the LASSO estimate $\widehat{v}^{\text{lasso}}$. There is no analytical solution to this problem, but we can numerically determine the MAP estimator and the corresponding feedback trading strategy. A numerical example is shown in Figure 14 (left). When $\alpha \in (0, 1)$, the heuristic MAP strategy has a trading threshold always larger than that of the robust LASSO strategy; see the right panel of Figure 14. We find that when α is sufficiently small, the MAP strategy $Z_{2,n}^{\text{map}}(y_1; \alpha, \xi_w)$ becomes discontinuous at its trading thresholds. This can also be seen from Figure 3 in the main text. Unlike the robust LASSO strategy which scales linearly with α, the MAP strategy has the same asymptotes as they are independent of α. All the above results show the significant discrepancy between these two strategies.

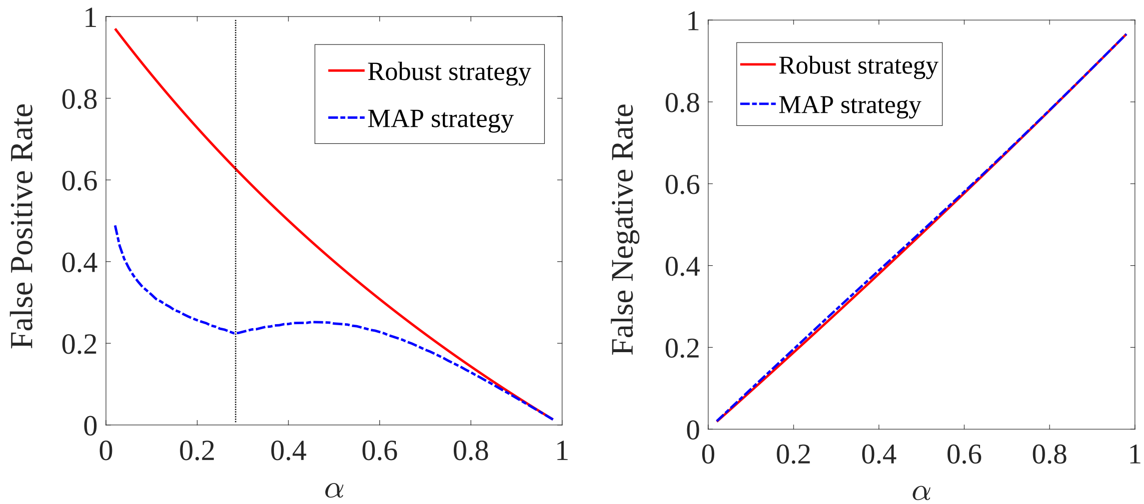

Figure 15. Left: false positive rates for the robust LASSO strategy $Z_{2,n}$ and the heuristic MAP strategy $Z_{2,n}^{\text{map}}$. Right: false positive rates for both types of strategies.

Next we compare the performances of the two strategies over the entire range of $\alpha \in (0, 1)$. The MAP strategy is not Bayesian rational as it tends to produce the *all-or-none* responses: whenever the order flow y_1 exceeds its trading thresholds, the MAP strategy treats y_1 as if it contains a Laplacian signal for sure and thus commits to the maximal trading intensity. This feature of binary classification stems from the heuristic MAP rule as it simply picks up the posterior mode and brutally ignores all the remaining posterior information. When the frequency of fat-tail shocks declines (i.e., as α decreases), the MAP strategy uses a larger threshold to achieve its binary classification. In contrast, the robust LASSO strategy does not alter the threshold since $K(\xi_w)$ is independent of α. Of course, a wider inaction zone implies a lower rate of Type I errors (false positives), as shown in Figure 15. In terms of Type II errors (false negatives), the two strategies are almost identical in performance.

The false positive rate can only reflect the accuracy in the direction of responses, not in the magnitude of responses. The MAP strategy may overly trade given its all-or-none feature. Figure 4 shows the expected total trading profits at $\alpha = 0.5$ under three types of strategies: the robust LASSO strategy $Z_{2,n}(y_1; \alpha, K(\xi_w))$, the optimal strategy $Z_{2,n}^o(y_1; \alpha, \xi_w)$, and the heuristic MAP strategy $Z_{2,n}^{\text{map}}(y_1; \alpha, \xi_w)$. One can see that for a wide range of ξ_w (relative to the true prior ξ_v), the LASSO strategy is much more profitable than the MAP strategy. The MAP strategy can easily lose money when it has estimate bias or faces intense competition.

A.9 Proof of Theorem 4

In the symmetric equilibrium of Theorem 1 where arbitrageurs follow the same optimal strategy $Z^o_{2,n}(\tilde{y}_1;\alpha,\xi_w)$, their expected total profits under the physical measure is given by

$$
\begin{aligned}
& \mathrm{E}\left[\sum_{n=1}^{N}(\tilde{v}-\tilde{p}_2)Z^o_{2,n}(\tilde{y}_1;\alpha,\xi_w)\right] \\
= \; & \mathrm{E}\left[N\left(\tilde{v}-\lambda_1\tilde{y}_1-\lambda_2\beta_2(\tilde{v}-\lambda_1\tilde{y}_1)-\sum_{n=1}^{N}\lambda_2 Z^o_{2,n}\right)Z^o_{2,n}\right] \\
= \; & (1-\beta_2\lambda_2)^2\mathrm{E}\left[\frac{N(N+1)\tilde{\theta}-N^2\widehat{\theta}(\tilde{y}_1;\alpha,\xi_w)}{(N+1)}\cdot\frac{\widehat{\theta}(\tilde{y}_1;\alpha,\xi_w)}{(N+1)\lambda_2}\right] \\
= \; & (1-\beta_2\lambda_2)^2\frac{N\mathrm{E}\left[\widehat{\theta}(\tilde{y}_1;\alpha,\xi_v)\widehat{\theta}(\tilde{y}_1;\alpha,\xi_w)\right]}{(N+1)\lambda_2}-(1-\beta_2\lambda_2)^2\frac{N^2\mathrm{E}\left[\widehat{\theta}(\tilde{y}_1;\alpha,\xi_w)^2\right]}{(N+1)^2\lambda_2}, \quad \text{(A59)}
\end{aligned}
$$

where we have used Eq. (18) as well as the following property

$$
\mathrm{E}\left[\tilde{\theta}\cdot\widehat{\theta}(\tilde{y}_1;\alpha,\xi_w)\right]=\mathrm{E}\left[\mathrm{E}\left[\tilde{\theta}\middle|\tilde{y}_1;\alpha,\xi_v\right]\widehat{\theta}(\tilde{y}_1;\alpha,\xi_w)\right]=\mathrm{E}\left[\widehat{\theta}(\tilde{y}_1;\alpha,\xi_v)\cdot\widehat{\theta}(\tilde{y}_1;\alpha,\xi_w)\right]. \quad \text{(A60)}
$$

When $\xi_w=\xi_v$ (unbiased case), it is easy to verify that

$$
\lim_{N\to\infty}\mathrm{E}\left[\sum_{n=1}^{N}(\tilde{v}-\tilde{p}_2)Z^o_{2,n}|\xi_w=\xi_v\right]=\lim_{N\to\infty}\frac{N(1-\beta_2\lambda_2)^2}{\lambda_2(N+1)^2}\mathrm{E}[\widehat{\theta}(\tilde{y}_1;\alpha,\xi_v)^2]=0. \quad \text{(A61)}
$$

Arbitrageurs eventually compete away their aggregate trading profit if $\xi_w=\xi_v$ and $N\to\infty$.

In the symmetric equilibrium of Theorem 2 where arbitrageurs use the same robust LASSO strategy $Z_{2,n}(\tilde{y}_1;\alpha,K(\xi_w))$, the expectation of their total profit is derived to be

$$
\begin{aligned}
& \mathrm{E}\left[\sum_{n=1}^{N}(\tilde{v}-\tilde{p}_2)Z_{2,n}(\tilde{y}_1;\alpha,K(\xi_w))\right] \\
= \; & \alpha(1-\beta_2\lambda_2)^2\mathrm{E}\left[\frac{N(N+1)(\tilde{v}-\lambda_1\tilde{y}_1)-\alpha N^2\widehat{\theta}^{\text{lasso}}(\tilde{y}_1;\xi_w)}{(N+1)}\cdot\frac{\widehat{\theta}^{\text{lasso}}(\tilde{y}_1;\xi_w)}{(N+1)\lambda_2}\right] \\
= \; & \alpha(1-\beta_2\lambda_2)^2\frac{\mathrm{E}\left[N(N+1)(\tilde{\theta}-\alpha\widehat{\theta}^{\text{lasso}}+\alpha\widehat{\theta}^{\text{lasso}})\widehat{\theta}^{\text{lasso}}-\alpha N^2(\widehat{\theta}^{\text{lasso}})^2\right]}{(N+1)^2\lambda_2} \\
= \; & \alpha(1-\beta_2\lambda_2)^2\frac{N\mathrm{E}\left[(\widehat{\theta}(\tilde{y}_1;\alpha,\xi_v)-\alpha\widehat{\theta}^{\text{lasso}})\widehat{\theta}^{\text{lasso}}\right]}{(N+1)\lambda_2}+\alpha^2(1-\beta_2\lambda_2)^2\frac{N\mathrm{E}\left[(\widehat{\theta}^{\text{lasso}})^2\right]}{(N+1)^2\lambda_2}, \quad \text{(A62)}
\end{aligned}
$$

where we have used Eq. (28) and the following result similar to (A60),

$$
\mathrm{E}\left[\tilde{\theta}\cdot\widehat{\theta}^{\text{lasso}}(\tilde{y}_1;\alpha,\xi_w)\right]=\mathrm{E}\left[\mathrm{E}\left[\tilde{\theta}\middle|\tilde{y}_1;\alpha,\xi_v\right]\widehat{\theta}^{\text{lasso}}\right]=\mathrm{E}\left[\widehat{\theta}(\tilde{y}_1;\alpha,\xi_v)\cdot\widehat{\theta}^{\text{lasso}}(\tilde{y}_1;\alpha,\xi_w)\right]. \quad \text{(A63)}
$$

Given the expressions of (18), (19), (29), and (30), one can further verify that

$$\mathrm{E}\left[[\widehat{\theta}(\tilde{y}_1;\alpha,\xi_v)-\alpha\widehat{\theta}^{\text{lasso}}]\widehat{\theta}^{\text{lasso}}\right]=\alpha\mathrm{E}\left[[\widehat{v}(\tilde{y}_1;\alpha=1,\xi_v)-\widehat{v}^{\text{lasso}}]\widehat{\theta}^{\text{lasso}}\right], \tag{A64}$$

where we have use the scaling property $\widehat{\theta}(\tilde{y}_1;\alpha,\xi_v)=\alpha(\widehat{v}[\tilde{y}_1;\alpha=1,\xi_v)-\lambda_1\tilde{y}_1]$ and the result that $\widehat{\theta}^{\text{lasso}}=(\widehat{v}^{\text{lasso}}-\lambda_1\tilde{y}_1)\mathbf{1}_{|\tilde{y}_1|>K}$ must satisfy $(\widehat{v}^{\text{lasso}}-\lambda_1\tilde{y}_1)\cdot\widehat{\theta}^{\text{lasso}}=(\widehat{\theta}^{\text{lasso}})^2$. In the limit $N\to\infty$, we find that the expected total profit Eq. (A62) is strictly positive:

$$\lim_{N\to\infty}\mathrm{E}\left[\sum_{n=1}^{N}(\tilde{v}-\tilde{p}_2)Z_{2,n}\right]=\frac{\alpha^2(1-\beta_2\lambda_2)^2}{\lambda_2}\mathrm{E}\left[[\widehat{v}(\tilde{y}_1;\xi_v)-\widehat{v}^{\text{lasso}}(\tilde{y}_1;\xi_w)]\cdot\widehat{\theta}^{\text{lasso}}(\tilde{y}_1;\xi_w)\right]>0. \tag{A65}$$

Here, we have used the notation $\widehat{v}(\tilde{y}_1;\xi_v)$ to stand for $\widehat{v}(\tilde{y}_1;\alpha=1,\xi_v)$. When $0<\xi_w\le\xi_v$ and $\alpha\in(0,1]$, one can verify that $\widehat{v}(\tilde{y}_1;\xi_v)>\widehat{v}^{\text{lasso}}(\tilde{y}_1;\xi_w)$ for $y_1>0$ and by symmetry, we also have $\widehat{v}(\tilde{y}_1;\xi_v)<\widehat{v}^{\text{lasso}}(\tilde{y}_1;\xi_w)$ for $y_1<0$. This means that $\widehat{v}(\tilde{y}_1;\xi_v)-\widehat{v}^{\text{lasso}}(\tilde{y}_1;\xi_w)$ and $\widehat{\theta}^{\text{lasso}}(\tilde{y}_1;\xi_w)$ have the same sign for $|\tilde{y}_1|>K(\xi_w)$, thus proving the positive sign of Eq. (A65) when $0<\xi_w\le\xi_v$. Because competition drives down profits, it also implies that Eq. (A62) is strictly positive for any $N\ge 1$ when $0<\xi_w\le\xi_v$.

We have used Monte-Carlo simulations to verify the analytical results in Theorem 4. It is also of interest to compare these two strategies in other performance measures (Table 1).

Table 1. Performance comparison based on Monte-Carlo simulations with $\alpha=1$.

$\xi_w=\xi_v, N=10$	Optimal $Z^o_{2,n}(y_1;\xi_w)$	Robust $Z_{2,n}(y_1;K(\xi_w))$
Gain-Loss Ratio	0.9545	1.3403
Profit Std Dev	0.1882	0.1107
Profit Skewness	3.9425	8.3697
Win-Loss Ratio	1.0011	1.5778
Ave Profit Per Trade	0.0014	0.1101
$\xi_w=0.9\xi_v, N=10$	Optimal $Z^o_{2,n}(y_1;\xi_w)$	Robust $Z_{2,n}(y_1;K(\xi_w))$
Gain-Loss Ratio	0.9959	1.1266
Profit Std Dev	0.1936	0.1503
Profit Skewness	5.0032	23.724
Win-Loss Ratio	1.0343	1.1593
Ave Profit Per Trade	-0.0024	0.2694

A.10 Proof of Proposition 4

When $N \to \infty$, if there is ever a finite mass, denoted $\phi \in (0,1]$, of arbitrageurs constrained by model risk as in Theorem 2, then the asset price at $t=2$ will be

$$
\begin{aligned}
\tilde{p}_2 &= \lambda_1\tilde{y}_1 + \lambda_2\Bigg[\beta_2(\tilde{v}-\lambda_1\tilde{y}_1) + \sum_{n=1}^{(1-\phi)N} Z^o_{2,n}(\tilde{y}_1;\alpha,\xi_w) + \sum_{n=1}^{\phi N} Z_{2,n}(\tilde{y}_1;\alpha,K(\xi_w)) + \tilde{u}_2\Bigg] \\
&= \lambda_1\tilde{y}_1 + \lambda_2\Bigg[\beta_2(\tilde{v}-\lambda_1\tilde{y}_1) + \sum_{n=1}^{N} Z^o_{2,n}(\tilde{y}_1;\alpha,\xi_w) + \sum_{n=1}^{\phi N}(Z_{2,n}-Z^o_{2,n}) + \tilde{u}_2\Bigg] \\
&\to \beta_2\lambda_2\tilde{v} + (1-\beta_2\lambda_2)\left(\widehat{v}(\tilde{y}_1;\alpha,\xi_w) + \phi\left[\alpha\widehat{\theta}^{\text{lasso}}(\tilde{y}_1;\xi_w) - \widehat{\theta}(\tilde{y}_1;\alpha,\xi_w)\right]\right) + \lambda_2\tilde{u}_2.
\end{aligned}
$$

We will keep using $\widehat{v}(\tilde{y}_1;\xi)$ to stand for $\widehat{v}(\tilde{y}_1;\alpha=1,\xi)$ and use $\widehat{\theta}(\tilde{y}_1;\xi)$ for $\widehat{\theta}(\tilde{y}_1;\alpha=1,\xi)$. The expectation of $\tilde{p}_2-\tilde{v}$ conditional on the price history $\{\tilde{p}_1=\lambda_1\tilde{y}_1,p_0\}$ is

$$
\begin{aligned}
&\lim_{N\to\infty}\mathrm{E}[\tilde{p}_2-\tilde{v}|\tilde{p}_1,p_0] = \lim_{N\to\infty}\mathrm{E}[\lambda_1\tilde{y}_1+\lambda_2\tilde{y}_2-\tilde{v}|\tilde{y}_1] \\
=\;& (1-\beta_2\lambda_2)\left(\widehat{v}(\tilde{y}_1;\alpha,\xi_w) - \mathrm{E}[\tilde{v}|\tilde{y}_1;\alpha,\xi_v]\right) + \phi(1-\beta_2\lambda_2)\left[\alpha\widehat{\theta}^{\text{lasso}}(\tilde{y}_1;\xi_w) - \widehat{\theta}(\tilde{y}_1;\alpha,\xi_w)\right] \\
=\;& (1-\beta_2\lambda_2)\left(\widehat{v}(\tilde{y}_1;\alpha,\xi_w) - \widehat{v}(\tilde{y}_1;\alpha,\xi_v)\right) + \phi(1-\beta_2\lambda_2)\left[\alpha\widehat{\theta}^{\text{lasso}}(\tilde{y}_1;\xi_w) - \widehat{\theta}(\tilde{y}_1;\alpha,\xi_w)\right] \\
=\;& (1-\beta_2\lambda_2)\left(\left[\widehat{v}(\tilde{y}_1;\alpha,\xi_w) - \widehat{v}(\tilde{y}_1;\alpha,\xi_v)\right] + \alpha\phi\left[\widehat{\theta}^{\text{lasso}}(\tilde{y}_1;\xi_w) - \widehat{\theta}(\tilde{y}_1;\xi_w)\right]\right) \\
=\;& \alpha(1-\beta_2\lambda_2)\left(\left[\widehat{v}(\tilde{y}_1;\xi_w) - \widehat{v}(\tilde{y}_1;\xi_v)\right] + \phi\left[\widehat{\theta}^{\text{lasso}}(\tilde{y}_1;\xi_w) - \widehat{\theta}(\tilde{y}_1;\xi_w)\right]\right). \qquad \text{(A66)}
\end{aligned}
$$

In deriving Eq. (A66), we have used the scaling property that $\widehat{\theta}(\tilde{y}_1;\alpha,\xi_w) = \alpha\widehat{\theta}(\tilde{y}_1;\alpha=1,\xi_w)$ and the observation that $\widehat{v}(\tilde{y}_1;\alpha,\xi_w) - \widehat{v}(\tilde{y}_1;\alpha,\xi_v) = \alpha\left[\widehat{v}(\tilde{y}_1;\alpha=1,\xi_w) - \widehat{v}(\tilde{y}_1;\alpha=1,\xi_v)\right]$, which follows from the expression of (19). Eq. (A66) is not equal to zero almost everywhere (for almost any realizations of $\tilde{y}_1$), unless both $\xi_w=\xi_v$ and $\phi=0$ hold simultaneously.

Note that the economy in either Theorem 1 or Theorem 2 can hold an infinite number of arbitrageurs when the following inequality holds: $\alpha+(1-\alpha)\beta_1\lambda_1 < \frac{2\beta_1\sqrt{\lambda_1\lambda_2}}{1-\beta_2\lambda_2}$. This is from the equilibrium condition (22) in Proposition 1 by taking the limit $N\to\infty$.

Figure 16 plots the aggregate trading profiles of arbitrageurs when they follow the robust LASSO strategy. The red solid line is for the homogeneous case when they use identical trading thresholds $K(\xi_v)$. The black dash-dot line is for the heterogeneous case when they use different trading thresholds $K(\xi_{w,n})$ for $n=1,...,N$. It has a narrower no-trade region determined by the most optimistic trader whose effective prior is $\xi^*_w := \max\{\xi_{w,1},...,\xi_{w,N}\}$. We impose $\frac{1}{N}\sum_{n=1}^N \xi^{-1}_{w,n} = \xi^{-1}_v$ so that the two aggregate trading strategies converge.

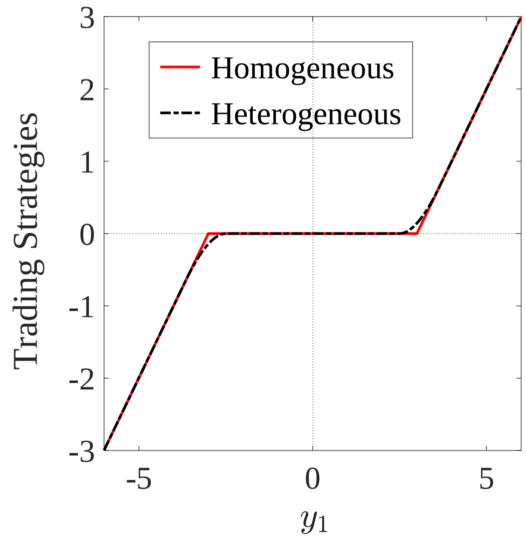

Figure 16. Comparison of the robust strategy $Z_{2,n}(y_1; \alpha, K(\xi_v))$ homogeneous for each n and the population average of the robust strategy over traders with heterogeneous thresholds, $\frac{1}{N}\sum_{n=1}^{N} Z_{2,n}(y_1; \alpha, K(\xi_{w,n}))$, where $\frac{1}{N}\sum_{n=1}^{N} \xi_{w,n}^{-1} = \xi_v^{-1}$ is imposed for a fair comparison.

A.11 Proof Proposition 6

Consider $J \geq 1$ stocks whose liquidation values have independent Laplacian-Gaussian mixture distributions, $\tilde{v}_j \sim \mathcal{LG}(\alpha_j, \xi_j)$, for $j = 1, ..., J$. Suppose arbitrageurs solve the same problem as in Theorem 2 for each asset based on their uncertain fat-tail prior $\tilde{v}_j \sim \mathcal{LG}(\alpha_j, \tilde{\xi}_j)$. When (C1)-(C3) hold for each asset, their robust strategy at $t = 2$ is similar to Eq. (27):

$$Z_{2,n}(y_{1,j}; \alpha_j, \xi_{w,j}) = \frac{\alpha_j(1-\beta_{1,j}\lambda_{1,j})(1-\beta_{2,j}\lambda_{2,j})}{\beta_{1,j}\lambda_{2,j}(N+1)} \left[y_{1,j} - \text{sign}(y_{1,j})K_j(\xi_{w,j})\right] \mathbf{1}_{|y_{1,j}|>K_j(\xi_{w,j})} \quad \text{(A67)}$$

which is a soft-thresholding function of $y_{1,j}$. With $p_{1,j} = 0$ and $p_{1,j} = \lambda_{1,j} y_{1,j}$, we can define the LASSO estimates $\widehat{\theta}_j^{\text{lasso}}(p_{1,j}; \xi_{w,j})$ of the pricing error for each asset $j \in \{1, ..., J\}$ as

$$\widehat{\theta}_j^{\text{lasso}} := \arg\min_{\theta_j} \; \frac{1}{2}\left|p_{1,j} - \frac{\beta_{1,j}\lambda_{1,j}\theta_j}{1-\beta_{1,j}\lambda_{1,j}}\right|^2 + \left(\frac{\lambda_{1,j}\sigma_{u,j}}{1-\beta_{1,j}\lambda_{1,j}}\right)^2 \frac{|\theta_j|}{\xi_{w,j}} = \frac{1-\beta_{1,j}\lambda_{1,j}}{\beta_{1,j}\lambda_{1,j}} \mathcal{S}(p_{1,j}; \lambda_{1,j}K_j). \quad \text{(A68)}$$

By Theorem 3, the robust strategy for each asset $j \in \{1, ..., J\}$ is a LASSO strategy:

$$Z_{2,n}(p_{1,j}; \alpha_j, \xi_{w,j}) = \frac{\alpha_j(1-\beta_{2,j}\lambda_{2,j})}{\lambda_{2,j}(N+1)} \widehat{\theta}_j^{\text{lasso}}(p_{1,j}; \xi_{w,j}). \quad \text{(A69)}$$

By proposition 5, if we consider the uncertainty of $\tilde{\alpha}_j$, we can simply replace α_j by its prior mean $\overline{\alpha}_j$ in the above equation and the new LASSO strategy will be $Z_{2,n}(p_{1,j}; \overline{\alpha}_j, \xi_{w,j})$.

A.12 Proof of Proposition 7

All other things being equal, we replace Eq. (9) by the mixture Gaussian prior below:

$$f(v;\alpha,\zeta)=\frac{\alpha}{\sqrt{2\pi\zeta^2}}\exp\left(-\frac{v^2}{2\zeta^2}\right)+\frac{1-\alpha}{\sqrt{2\pi\sigma_v^2}}\exp\left(-\frac{v^2}{2\sigma_v^2}\right). \tag{A70}$$

Arbitrageurs know α but are uncertain about the volatility ζ of the Gaussian shocks. Since Bayesian learning will maintain the simple scaling property about α, it suffices to consider the case that $\alpha=1$. Each arbitrageur's prior in this case becomes $\tilde{v}\sim\mathcal{N}(0,\tilde{\zeta}^2)$, where $\tilde{\zeta}\in[\zeta_L,\zeta_H]$. For any fixed Gaussian prior $\tilde{\zeta}=\zeta$, arbitrageurs believe that $\tilde{y}_1=\beta_1\tilde{v}+\tilde{u}_1\sim\mathcal{N}\left(0,(\beta_1\zeta)^2+\sigma_u^2\right)$. Their posterior belief about $\tilde{v}$ conditional on $\tilde{y}_1$ is

$$f(v|y_1)=\frac{f(y_1|v)f(v;\zeta)}{f(y_1)}=\frac{1}{2\pi\zeta\sigma_u f(y_1)}\exp\left[-\frac{(y_1-\beta_1 v)^2}{2\sigma_u^2}-\frac{v^2}{2\zeta^2}\right]. \tag{A71}$$

By projection theorem, they will use the linear estimator under the Gaussian prior,

$$\widehat{v}^{\text{ridge}}(y_1;\zeta)=\mathrm{E}[\tilde{v}|y_1,\zeta]=\frac{\beta_1\zeta^2y_1}{(\beta_1\zeta)^2+\sigma_u^2}=\frac{y_1}{\beta_1+\sigma_u^2/(\beta_1\zeta^2)}, \tag{A72}$$

which is the simplest version of *ridge regression* (Hastie et al. (2009)) with an l_2 norm penalty. For any given ζ, the optimal strategy of arbitrageurs is always a linear function of y_1:

$$Z_{2,n}^o(y_1;\zeta)=\frac{1-\beta_2\lambda_2}{\lambda_2(N+1)}(\widehat{v}^{\text{ridge}}-\lambda_1y_1)=\frac{\beta_1(1-\beta_1\lambda_1)\zeta^2-\lambda_1\sigma_u^2}{(\beta_1\zeta)^2+\sigma_u^2}\cdot\frac{1-\beta_2\lambda_2}{\lambda_2(N+1)}y_1. \tag{A73}$$

Any uncertainty about the prior ζ only changes the slope of this linear strategy. Therefore, the robust strategy should be linear under the max-min choice criteria. It is easy to verify that $Z_{2,n}^o(y_1;\zeta=\sigma_v)=0$ since the price is supposed to be efficient when $\zeta=\sigma_v$. Hence, if $\zeta_L\le\sigma_v\le\zeta_H$, then $Z_{2,n}^o(y_1;\zeta_H)$ is upward sloping and $Z_{2,n}^o(y_1;\zeta_L)$ is downward sloping such that the max-min strategy is no trade at all. If $\sigma_v<\zeta_L\le\zeta_H$, then the max-min strategy is the upward sloping linear strategy $Z_{2,n}^o(y_1;\zeta_L)$. If $\zeta_L\le\zeta_H<\sigma_v$, then the max-min strategy is the downward-sloping linear strategy $Z_{2,n}^o(y_1;\zeta_H)$. In sum, the max-min robust strategy is

$$Z_{2,n}(y_1;\zeta)=\begin{cases}Z_{2,n}^o(y_1;\zeta_L), & \text{if } \sigma_v<\zeta_L\le\zeta_H,\\ 0, & \text{if } \zeta_L\le\sigma_v\le\zeta_H,\\ Z_{2,n}^o(y_1;\zeta_H), & \text{if } \zeta_L\le\zeta_H<\sigma_v.\end{cases} \tag{A74}$$

The above strategy is always a linear function of y_1, without any trading threshold.

References

Abreu, D., and M. K. Brunnermeier. 2002. Synchronization risk and delayed arbitrage. *Journal of Financial Economics* 66:341–360.

Ahn, D., S. Choi, D. Gale, and S. Kariv. 2014. Estimating ambiguity aversion in a portfolio choice experiment. *Quantitative Economics* 5:195–223.

Athey, S. 2018. The impact of machine learning on economics. In *The economics of artificial intelligence: An agenda*, pp. 507–547. University of Chicago Press.

Back, K. 1992. Insider trading in continuous time. *Review of Financial Studies* 5:387–409.

Back, K., C. H. Cao, and G. A. Willard. 2000. Imperfect competition among informed traders. *Journal of Finance* 55:2117–2155.

Banerjee, S., J. Davis, and N. Gondhi. 2019. Choosing to disagree in financial markets. *Available at SSRN* 3335257.

Barberis, N., R. Greenwood, L. Jin, and A. Shleifer. 2015. X-CAPM: An extrapolative capital asset pricing model. *Journal of Financial Economics* 115:1–24.

Barberis, N., R. Greenwood, L. Jin, and A. Shleifer. 2018. Extrapolation and bubbles. *Journal of Financial Economics* .

Barberis, N., and R. Thaler. 2003. A survey of behavioral finance. *Handbook of the Economics of Finance* 1:1053–1128.

Behr, A., and U. Pötter. 2009. Alternatives to the normal model of stock returns: Gaussian mixture, generalised logF and generalised hyperbolic models. *Annals of Finance* 5:49–68.

Biais, B., T. Foucault, and S. Moinas. 2015. Equilibrium fast trading. *Journal of Financial economics* 116:292–313.

Bollerslev, T. 1986. Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics* 31:307–327.

Bollerslev, T., and V. Todorov. 2011a. Estimation of jump tails. *Econometrica* 79:1727–1783.

Bollerslev, T., and V. Todorov. 2011b. Tails, fears, and risk premia. *Journal of Finance* 66:2165–2211.

Bossaerts, P., P. Ghirardato, S. Guarnaschelli, and W. R. Zame. 2010. Ambiguity in asset markets: Theory and experiment. *Review of Financial Studies* 23:1325–1359.

Brogaard, J., T. Hendershott, and R. Riordan. 2014. High-frequency trading and price discovery. *Review of Financial Studies* 27:2267–2306.

Budish, E., P. Cramton, and J. Shim. 2015. The high-frequency trading arms race: Frequent batch auctions as a market design response. *The Quarterly Journal of Economics* 130:1547–1621.

Camerer, C. F. 2018. Artificial intelligence and behavioral economics. In *The economics of artificial intelligence: An agenda*, pp. 587–608. University of Chicago Press.

Chinco, A., A. D. Clark-Joseph, and M. Ye. 2019. Sparse signals in the cross-section of returns. *Journal of Finance* 74:449–492.

Colliard, J.-E., T. Foucault, and S. Lovo. 2022. Algorithmic Pricing and Liquidity in Securities Markets. *HEC Paris Research Paper* .

Collin-Dufresne, P., and V. Fos. 2016. Insider trading, stochastic liquidity, and equilibrium prices. *Econometrica* 84:1441–1475.

Da, R., S. Nagel, and D. Xiu. 2022. The statistical limit of arbitrage. Tech. rep., Technical Report, Chicago Booth.

DeLong, J. B., A. Shleifer, L. H. Summers, and R. J. Waldmann. 1990. Positive feedback investment strategies and destabilizing rational speculation. *Journal of Finance* 45:379–395.

Deng, W., L. Gao, B. Hu, and G. Zhou. 2020. Seeing is Believing: Annual Report "Graphicity" and Stock Returns Predictability. *Available at SSRN 3723126* .

Didisheim, A., S. Ke, B. T. Kelly, and S. Malamud. 2023. Complexity in Factor Pricing Models. *Swiss Finance Institute Research Paper* .

Domingos, P., and M. Pazzani. 1997. On the optimality of the simple Bayesian classifier under zero-one loss. *Machine Learning* 29:103–130.

Dong, X., Y. Li, D. E. Rapach, and G. Zhou. 2022. Anomalies and the expected market return. *Journal of Finance* 77:639–681.

Dou, W. W., I. Goldstein, and Y. Ji. 2023. AI-Powered Trading, Algorithmic Collusion, and Price Efficiency. *Available at SSRN 4452704* .

Dow, J., and S. R. d. C. Werlang. 1992. Uncertainty aversion, risk aversion, and the optimal choice of portfolio. *Econometrica* pp. 197–204.

Easley, D., and M. O'Hara. 2009. Ambiguity and nonparticipation: The role of regulation. *Review of Financial Studies* 22:1817–1843.

Easley, D., and M. O'Hara. 2010. Microstructure and ambiguity. *Journal of Finance* 65:1817–1846.

Embrechts, P., C. Klüppelberg, and T. Mikosch. 2013. *Modelling extremal events: for insurance and finance*, vol. 33. Springer Science & Business Media.

Engle, R. F. 1982. Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica* pp. 987–1007.

Epstein, L. G., and M. Schneider. 2008. Ambiguity, information quality, and asset pricing. *Journal of Finance* 63:197–228.

Epstein, L. G., and M. Schneider. 2010. Ambiguity and Asset Markets. *Annual Review of Financial Economics* 2:10–1.

Fama, E. F. 1963. Mandelbrot and the stable Paretian hypothesis. *Journal of Business* 36:420–429.

Fama, E. F. 1965. The behavior of stock-market prices. *Journal of Business* 38:34–105.

Foster, F. D., and S. Viswanathan. 1994. Strategic trading with asymmetrically informed traders and long-lived information. *Journal of Financial and Quantitative Analysis* 29:499–518.

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

Freyberger, J., A. Neuhierl, and M. Weber. 2020. Dissecting characteristics nonparametrically. *Review of Financial Studies* 33:2326–2377.

Gabaix, X. 2014. A sparsity-based model of bounded rationality. *Quarterly Journal of Economics* 129:1661–1710.

Gabaix, X., P. Gopikrishnan, V. Plerou, and H. E. Stanley. 2006. Institutional investors and stock market volatility. *Quarterly Journal of Economics* 121:461–504.

Gabaix, X., A. Krishnamurthy, and O. Vigneron. 2007. Limits of arbitrage: theory and evidence from the mortgage-backed securities market. *Journal of Finance* 62:557–595.

Garlappi, L., R. Uppal, and T. Wang. 2006. Portfolio selection with parameter and model uncertainty: A multi-prior approach. *Review of Financial Studies* 20:41–81.

Giglio, S., B. Kelly, and D. Xiu. 2022. Factor models, machine learning, and asset pricing. *Annual Review of Financial Economics* 14:337–368.

Gilboa, I., and D. Schmeidler. 1989. Maxmin expected utility with a unique set of priors. *Journal of Mathematical Economics* 18:385–439.

Goto, S., and Y. Xu. 2015. Improving mean variance optimization through sparse hedging restrictions. *Journal of Financial and Quantitative Analysis* 50:1415–1441.

Granger, C. W. J., and Z. Ding. 1995. Some properties of absolute return: An alternative measure of risk. *Annales d'Economie et de Statistique* pp. 67–91.

Greig, D. M., B. T. Porteous, and A. H. Seheult. 1989. Exact maximum a posteriori estimation for binary images. *Journal of the Royal Statistical Society: Series B (Methodological)* 51:271–279.

Grinold, R. C., and R. N. Kahn. 2000. Active Portfolio Management .

Gromb, D., and D. Vayanos. 2010. Limits of arbitrage. *Annual Review of Financial Economics* 2:251–275.

Gu, S., B. Kelly, and D. Xiu. 2020. Empirical asset pricing via machine learning. *Review of Financial Studies* 33:2223–2273.

Haas, M., S. Mittnik, and M. S. Paolella. 2006. Modelling and predicting market risk with Laplace–Gaussian mixture distributions. *Applied Financial Economics* 16:1145–1162.

Han, Y., Y. Liu, G. Zhou, and Y. Zhu. 2021. Technical Analysis in the Stock Market: A Review. *Available at SSRN 3850494* .

Hansen, L., and T. J. Sargent. 2001. Robust control and model uncertainty. *American Economic Review* 91:60–66.

Hansen, L. P., and T. J. Sargent. 2008. *Robustness.* Princeton University Press.

Hasbrouck, J., and D. J. Seppi. 2001. Common factors in prices, order flows, and liquidity. *Journal of Financial Economics* 59:383–411.

Hastie, T., R. Tibshirani, J. H. Friedman, and J. H. Friedman. 2009. *The elements of statistical learning: data mining, inference, and prediction*, vol. 2. Springer.

Hastie, T., R. Tibshirani, and M. Wainwright. 2015. Statistical learning with sparsity. *Monographs on statistics and applied probability* 143:143.

Hendershott, T., C. M. Jones, and A. J. Menkveld. 2011. Does algorithmic trading improve liquidity? *Journal of Finance* 66:1–33.

Holden, C. W., and A. Subrahmanyam. 1992. Long-lived private information and imperfect competition. *Journal of Finance* 47:247–270.

Huang, J.-Z., and Z. Shi. 2022. Machine-Learning-Based Return Predictors and the Spanning Controversy in Macro-Finance. *Management Science* .

Huberman, G., and W. Stanzl. 2004. Price manipulation and quasi-arbitrage. *Econometrica* 72:1247–1275.

Huddart, S., J. S. Hughes, and C. B. Levine. 2001. Public disclosure and dissimulation of insider trades. *Econometrica* 69:665–681.

Illeditsch, P. K. 2011. Ambiguous information, portfolio inertia, and excess volatility. *Journal of Finance* 66:2213–2247.

Jordan, M. I., and T. M. Mitchell. 2015. Machine learning: Trends, perspectives, and prospects. *Science* 349:255–260.

Kelly, B., S. Malamud, and L. H. Pedersen. 2023. Principal portfolios. *Journal of Finance* 78:347–387.

Kelly, B. T., S. Malamud, and K. Zhou. 2022. The virtue of complexity in return prediction. Tech. rep., National Bureau of Economic Research.

Kelly, B. T., and D. Xiu. 2023. Financial machine learning. Tech. rep., National Bureau of Economic Research.

Khandani, A. E., and A. W. Lo. 2011. What happened to the quants in August 2007? Evidence from factors and transactions data. *Journal of Financial Markets* 14:1–46.

Klibanoff, P., M. Marinacci, and S. Mukerji. 2005. A smooth model of decision making under ambiguity. *Econometrica* 73:1849–1892.

Kondor, P. 2009. Risk in dynamic arbitrage: the price effects of convergence trading. *Journal of Finance* 64:631–655.

Kotz, S., T. Kozubowski, and K. Podgórski. 2001. *The Laplace distribution and generalizations: a revisit with applications to communications, economics, engineering, and finance.* 183. Springer Science & Business Media.

Kou, S. G. 2002. A jump-diffusion model for option pricing. *Management Science* 48:1086–1101.

Kozak, S., S. Nagel, and S. Santosh. 2020. Shrinking the cross-section. *Journal of Financial Economics* 135:271–292.

Kyle, A. S. 1985. Continuous auctions and insider trading. *Econometrica* pp. 1315–1335.

Kyle, A. S., and A. A. Obizhaeva. 2016. Market microstructure invariance: Empirical hypotheses. *Econometrica* 84:1345–1404.

Lewis, M. 2014. *Flash Boys: a Wall Street Revolt.* WW Norton & Company.

Lillo, F., and R. N. Mantegna. 2000. Variety and volatility in financial markets. *Physical Review E* 62:6126.

Lim, J. S., and A. V. Oppenheim. 1979. Enhancement and bandwidth compression of noisy speech. *Proceedings of the IEEE* 67:1586–1604.

Lo, A. W., H. Mamaysky, and J. Wang. 2000. Foundations of technical analysis: Computational algorithms, statistical inference, and empirical implementation. *Journal of Finance* 55:1705–1765.

Loeb, T. F. 1983. Trading cost: the critical link between investment information and results. *Financial Analysts Journal* pp. 39–44.

Malikov, G., and P. Pasquariello. 2021. Quants, Strategic Speculation, and Financial Market Quality. *Available at SSRN 3890275* .

Mantegna, R. N., and H. E. Stanley. 1999. *Introduction to econophysics: Correlations and complexity in finance.* Cambridge University Press.

Martin, I. W., and S. Nagel. 2022. Market efficiency in the age of big data. *Journal of Financial Economics* 145:154–177.

McQueen, G., and K. Vorkink. 2004. Whence GARCH? A preference-based explanation for conditional volatility. *Review of Financial Studies* 17:915–949.

Mullainathan, S., and J. Spiess. 2017. Machine learning: an applied econometric approach. *Journal of Economic Perspectives* 31:87–106.

Mussalli, G. 2018. Quant Meltdown: 10 Years Later. *PanAgora Asset Management.* .

Nagel, S. 2021. *Machine Learning in Asset Pricing*, vol. 8. Princeton University Press.

Plerou, V., P. Gopikrishnan, X. Gabaix, and H. E. Stanley. 2002. Quantifying stock-price response to demand fluctuations. *Physical Review E* 66:027104.

Rachev, S., and A. SenGupta. 1993. Laplace-Weibull mixtures for modeling price changes. *Management Science* 39:1029–1038.

Rapach, D. E., J. K. Strauss, and G. Zhou. 2013. International stock return predictability: what is the role of the United States? *Journal of Finance* 68:1633–1662.

Robert, C. P., et al. 2007. *The Bayesian choice: from decision-theoretic foundations to computational implementation*, vol. 2. Springer.

Schmeidler, D. 1989. Subjective probability and expected utility without additivity. *Econometrica* pp. 571–587.

Shleifer, A., and R. W. Vishny. 1997. The limits of arbitrage. *Journal of Finance* 52:35–55.

Silva, A. C., R. E. Prange, and V. M. Yakovenko. 2004. Exponential distribution of financial returns at mesoscopic time lags: a new stylized fact. *Physica A* 344:227–235.

Sims, C. A. 2003. Implications of rational inattention. *Journal of Monetary Economics* 50:665–690.

Tibshirani, R. 1996. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)* pp. 267–288.

Tversky, A., and D. Kahneman. 1974. Judgment under Uncertainty: Heuristics and Biases: Biases in judgments reveal some heuristics of thinking under uncertainty. *Science* 185:1124–1131.

Van Kervel, V., and A. J. Menkveld. 2019. High-frequency trading around large institutional orders. *Journal of Finance* 74:1091–1137.

Vayanos, D. 1999. Strategic trading and welfare in a dynamic market. *Review of Economic Studies* 66:219–254.

Vayanos, D. 2001. Strategic trading in a dynamic noisy market. *Journal of Finance* 56:131–171.

Xiong, W. 2001. Convergence trading with wealth effects: an amplification mechanism in financial markets. *Journal of Financial Economics* 62:247–292.

Xu, H., C. Caramanis, and S. Mannor. 2008. Robust regression and lasso. *Advances in Neural Information Processing Systems* 21.

Yang, L., and H. Zhu. 2020. Back-running: Seeking and hiding fundamental information in order flows. *Review of Financial Studies* 33:1484–1533.

Zhu, Y., and G. Zhou. 2009. Technical analysis: An asset allocation perspective on the use of moving averages. *Journal of Financial Economics* 92:519–544.