Law of Small Numbers in Financial Markets: Theory and Evidence*

Lawrence J. Jin[a] and Cameron Peng[b]

[a] *Cornell SC Johnson College of Business and NBER*

[b] *Department of Finance, the London School of Economics and Political Science*

ABSTRACT

We build a model of the law of small numbers (LSN)—the incorrect belief that even small samples represent the properties of the underlying population—to study its implications for trading behavior and asset prices. In the model, a belief in the LSN induces investors to expect short-term price trends to revert and long-term price trends to continue. As a result, asset prices exhibit excess volatility, short-term momentum, and long-term reversals. The model makes additional predictions about investor behavior, including the coexistence of the disposition effect and return extrapolation, a weakened disposition effect for long-term holdings, "doubling down" in buying, consistency between doubling down and the disposition effect, and heterogeneous trading propensities to past returns. By testing these predictions using account-level transaction data, we show that the LSN provides a parsimonious way for understanding a variety of puzzles about investor behavior and asset prices.

JEL classification: G02, G11, G12

Keywords: Law of Small Numbers, Trading Behavior, Disposition Effect, Return Extrapolation

*We thank Nicholas Barberis, James Choi, Cary Frydman, Marco Sammon, Dimitri Vayanos, and seminar participants at the City University of Hong Kong, Cornell University, the London School of Economics, the University of Florida, and the University of Notre Dame for helpful discussions and comments. We are grateful to Terrance Odean for providing the data. This paper supersedes an earlier working paper "Investor Behavior Under the Law of Small Numbers." Peng acknowledges financial support received during Ph.D. from a Whitebox Advisors research grant. Andrew Wang provided excellent research assistance.

E-mail addresses: lawrence.jin@cornell.edu (L. Jin), c.peng9@lse.ac.uk (C. Peng).

1. Introduction

When making forecasts about a random outcome, a common mistake people make is the "gambler's fallacy." For example, when a fair coin is tossed multiple times, a streak of heads makes it more likely for people to expect the next toss to be a tail, even though the objective probability remains constant at 50% (Rapoport and Budescu, 1992, 1997). The gambler's fallacy is often seen as indicative of the "law of small numbers (LSN)"—the incorrect belief that even a small, local sample represents the characteristics of the underlying population (Tversky and Kahneman, 1971).[1] More generally, along with other heuristics such as overreaction and base-rate neglect, the LSN falls under the broad notion that people often jump to conclusions too quickly by relying on too little data.

An immediate consequence of the LSN is that people behave as contrarians and, when predicting outcomes of a random sequence, they tend to expect immediate reversals in trends. However, it has also been suggested that the LSN can simultaneously lead to "hot hands" whereby people expect a streak of similar outcomes to continue (Rabin, 2002; Rabin and Vayanos, 2010). For example, a basketball player on a hot streak is often believed to be more likely to make the next shot, although the actual outcome appears uncorrelated with the previous streak (Gilovich, Vallone, and Tversky, 1985; Camerer, 1989; Tversky and Gilovich, 1989a,b). The two seemingly inconsistent phenomena can be reconciled based on people's prior knowledge about the data-generating process: when people know about the data-generating process in advance, the LSN results in the gambler's fallacy; but when they do not, they tend to rely too much on the few data points they have to make inferences, leading to "hot hands" instead.

In this paper, we develop a model of the LSN to study its implications for trading behavior and asset prices. We view the setting of trading in financial markets as one in which the LSN can be playing an important role, because investors constantly observe past trends in prices and fundamentals, and they need to make forecasts about future prices and fundamentals—a problem that resembles predicting outcomes of a random sequence. In these decisions, investors' beliefs about serial correlation, potentially fallacious and characterized by the LSN, can have a significant impact on their trading behavior and on asset prices. While existing papers have modeled the

[1] The same idea has also been labelled "local representativeness" (Bar-Hillel and Wagenaar, 1991).

LSN in general economic settings (e.g., Rabin, 2002; Rabin and Vayanos, 2010), our paper applies this belief structure in a financial setting with equilibrium asset prices. We derive new testable predictions about trading behavior and asset prices, and importantly, we also empirically test these predictions in the data.

We start with a tractable, continuous-time model of portfolio choice and asset prices. The model features two types of investors, rational arbitrageurs and LSN investors. Both have mean-variance preferences, and they allocate wealth between a risk-free asset and a risky asset. The risky asset has an exogenous dividend process, and its price process is determined endogenously in equilibrium. Rational arbitrageurs correctly understand the dividend process and the price process. However, LSN investors do not directly observe the true price process and need to make forecasts about future price changes based on their information set. In particular, we assume that they use an incorrect yet intuitive mental model to make inferences: they believe that the risky asset's price change is determined by a quality term—which is time-varying and unobservable—and a noise term, and they make inferences about the asset's quality by observing its past prices. Under this basic setup, good past returns indicate high asset quality; as such, LSN investors behave as return extrapolators.

We then introduce the LSN into investor beliefs. Specifically, following Rabin (2002) and Rabin and Vayanos (2010), we assume that, when making inferences about the underlying price process, LSN investors erroneously believe that the noise term is *negatively* auto-correlated. Intuitively, this assumption captures the gambler's fallacy, in that LSN investors expect short-term deviations from the mean to quickly revert in the near future. Compared to the earlier case without the LSN assumption, LSN investors' belief structure changes in two significant ways. First, different from simple return extrapolation in which beliefs about future price changes depend *positively* on all past price changes, LSN investors' beliefs depend *negatively* on recent price changes—they expect strong and immediate reversals for short-term price trends. This result follows directly from the assumption that LSN investors believe the noise term to be negatively auto-correlated. Second, consistent with return extrapolation, LSN investors' beliefs depend *positively* on price changes from the distant past. Strikingly, the degree of return extrapolation is stronger than in the case without the LSN assumption. Therefore, the same force that generates short-term contrarian beliefs also leads to stronger tendencies of return extrapolation based on long-term price trends.

Given mean-variance preferences, the above belief structure directly translates into investors' trading behaviors. On the one hand, LSN investors exhibit the disposition effect, selling when asset prices have recently gone up. On the other hand, they are also return extrapolators, buying when asset prices have gone up consistently over a long period of time. In this way, the model leads to the coexistence of the disposition effect and return extrapolation. These trading responses further feed back into asset price dynamics: in the short run, the disposition effect induces short-term momentum; in the longer run, return extrapolation results in long-term reversals. Overall, asset prices exhibit excess volatility, in that prices move more than in a benchmark model without LSN investors.

In the above model, LSN investors form incorrect beliefs about future price changes by looking at past price changes. They then use these beliefs about price changes to form their share demand of the risky asset. The direct mapping between past price changes and expectations of future price changes make this this thought process psychologically simple and realistic. We further consider an alternative specification in which LSN investors form incorrect beliefs about future *dividend* changes by looking at past dividend changes. In this scenario, before making investment decisions, LSN investors take an extra step to derive beliefs about prices from their beliefs about dividends. Because of this extra step investors need to take, we view this thought process as less realistic. Nonetheless, under this alternative specification, we again observe a similar dichotomy in belief formation: LSN investors' beliefs about future price changes depend negatively on recent price changes but positively on price changes from the distant past.

After analyzing the model's implications for investor beliefs, we examine and test the model's additional predictions using the Odean data (Odean, 1998; Barber and Odean, 2000). First, the model makes predictions about the degree of the disposition effect as a function of one's holding period. Specifically, the contrarian leg of investors' belief structure primarily concerns recent periods, so the model predicts a stronger disposition effect for positions with a short holding period. For positions with a longer holding period, return extrapolation starts to kick in, working in the opposite direction of the disposition effect. This prediction is supported by the Odean data. Among stocks bought within the last month, the probability of selling a winner is almost twice as high as the probability of selling a loser. In contrast, for positions held for more than a year, the propensities of selling winners and losers are virtually the same.

Second, the model suggests that investors not only display a disposition effect in selling, but also tend to "double down" in buying. That is, when they increase holdings of an existing position, they tend to buy shares that have gone down in value and are less likely to buy shares that have gone up in value. We confirm this prediction in the Odean data: on average, investors are 50% more likely to buy loser stocks than winner stocks, a result that is consistent with Odean (1998).

Third, the model not only predicts the coexistence of the disposition effect and doubling down at the *aggregate* level, but also proposes a strong association between these two phenomena at the *individual* level. Specifically, those who are more likely to double down in buying are also expected to exhibit a stronger disposition effect in selling, as the LSN belief structure underlies both behaviors. To test this prediction, we categorize investors into five groups based on their tendencies to double down in buying, and then compare the degrees of the disposition effect observed in selling across the five groups. Consistent with our hypothesis, the degree of the disposition effect increases monotonically with the tendency to double down, lending support to the idea that the LSN drives both buying and selling decisions.

Fourth, the model predicts that an individual's trading propensity, based on past returns, depends on their LSN beliefs. In an extension of the model, we consider LSN investors who believe noise is negatively auto-correlated and pure extrapolators who believe noise is i.i.d. LSN investors' selling propensity increases with recent returns, while their buying propensity decreases. The trading propensities of pure extrapolators display the opposite patterns. Our findings support these predictions and demonstrate the importance of investor heterogeneity in studying trading behavior. Ben-David and Hirshleifer (2012)'s "V-shaped" pattern in investors' buying and selling propensities is not observed in the trading of LSN investors, and a more careful consideration of investor heterogeneity is needed to account for this phenomenon.

Lastly, we also examine the model's prediction on asset prices regarding the sources of momentum and long-term reversals. The model suggests that individual stocks associated with stronger LSN beliefs should show stronger short-term momentum and long-term reversals. To test this, we analyze mutual fund holdings data, assigning each fund-quarter observation based on past stock returns. We consider funds holding stocks with good long-term performance but poor recent performance to have stronger LSN beliefs. Aggregating these measures at the stock level, we find supportive evidence that stocks with underlying funds prone to LSN beliefs exhibit stronger short-term momentum and long-term reversals.

We present a model of the LSN to account for trading behavior and asset prices. Previous work has built models to study belief formation under the LSN in partial-equilibrium settings (Rabin, 2002; Rabin and Vayanos, 2010). We introduce this belief structure in an equilibrium asset pricing model to study its implications for trading behavior and asset prices. This requires studying a more complex economic environment by specifying investor preferences and portfolio problems, introducing other market participants such as rational arbitrageurs, and analyzing the determination of asset prices in equilibrium. In this regard, the closest to our model is Teguia (2017), who also develops an equilibrium model that features LSN investors and rational traders. Our paper and Teguia (2017) differ in two important aspects. First, our paper explores novel predictions that are not considered by Teguia (2017): the degree of the disposition effect as a function of one's holding period, "doubling down" in buying, consistency between doubling down and the disposition effect, and heterogeneous trading propensities to past returns. Second and more importantly, we provide empirical tests of our model's predictions using account-level transaction data, and we find strong consistency between the data and the model's predictions.

Our model can simultaneously generate return extrapolation and the disposition effect. Importantly, the production of both phenomena relies on the single ingredient that random errors are erroneously perceived to be negatively auto-correlated. Other papers, such as Barberis, Shleifer, and Vishny (1998) and Liao, Peng, and Zhu (2022), can also generate both phenomena, but they require multiple psychological ingredients.[2] Therefore, we show theoretically that the LSN can be a unifying psychological root for both return extrapolation and the disposition effect.

Our empirical analysis has important implications for the study of investor behavior. First, we propose the LSN as a belief-based explanation for the disposition effect. Existing papers have proposed explanations based on non-traditional preferences such as prospect theory and realization utility (Barberis and Xiong, 2009, 2012; Ingersoll and Jin, 2013), or other psychological phenomena such as cognitive dissonance (Chang, Solomon, and Westerfield, 2016) and mental accounting (Frydman, Hartzmark, and Solomon, 2018). We show and confirm that the disposition effect can also arise from contrarian beliefs over short-term trends, which in turn can be derived from the

[2] More precisely, Barberis et al. (1998) require both conservatism and representativeness while Liao et al. (2022) require both extrapolation and realization utility.

LSN. Consistent with our model, we find that investors are particularly likely to sell assets whose price has only *recently* gone up—a phenomenon that most existing explanations of the disposition effect (e.g., non-traditional preferences where utility is a function of holding-period returns) do not speak to.

Second, we show that, on the flip side of the disposition effect, there also exits doubling down in buying behavior. More importantly, these two phenomena are tightly linked to each other: those who are more likely to double down are also those who display a stronger disposition effect in selling. This link enriches our understanding of investor trading by considering the buying side together with the selling side. Moreover, it raises the bar for explanations of the disposition effect: given the tight link between the buying behavior and the selling behavior, a unifying explanation should be able to account for both.

Finally, our results have implications for the well-documented "V-shape" trading propensities (Ben-David and Hirshleifer, 2012). Previously, the V-shape is often considered an aggregate phenomenon that applies to the average investor in the population. We uncover additional heterogeneity on the strength of the V-shape in the cross-section of investors: it is close to nonexistent among LSN investors, but is much stronger among extrapolators. Therefore, our results call for bringing in heterogeneity of investor beliefs to further understand the V-shape.

The rest of the paper proceeds as follows. Section 2 presents motivating evidence for the LSN from experimental and field settings. Section 3 presents the model and discusses its predictions. Section 4 tests the model's predictions on trading behavior using the Odean data. Section 5 tests the model's predictions on asset prices using the mutual fund holdings data. Section 6 concludes. Additional details are in the Appendix.

2. Motivating evidence

The law of small numbers (LSN) refers to the incorrect belief that even small samples represent the characteristics of the underlying population (Tversky and Kahneman, 1971; Rabin, 2002). According to the LSN, people expect good and bad outcomes to balance out over a short streak, so that the empirical distribution revealed by the short streak mimics the theoretical distribution of the population. For example, when a fair coin is tossed, after seeing several heads in a row, people tend to overestimate the probability of seeing a tail in the next toss, even though the objective probability remains constant at 50% (Rapoport and Budescu, 1992, 1997). This phenomenon, termed the "gambler's fallacy," has been robustly documented in many experimental settings and is commonly viewed as direct evidence of the law of small numbers. For example, additional evidence on the gambler's fallacy has been obtained in other experiments, such as those based on production tasks and recognition tasks, as reviewed by Bar-Hillel and Wagenaar (1991).

In parallel with the gambler's fallacy, researchers have also documented a different phenomenon called "the hot-hand fallacy:" in some settings, after seeing a streak of similar outcomes, people expect trend to continue rather than to reverse (Gilovich et al., 1985; Camerer, 1989; Tversky and Gilovich, 1989a,b). For example, a basketball player on a hot streak is often believed to be more likely to make the next shot, although the actual outcome appears uncorrelated with the previous streak. The two fallacies may initially appear to contradict to each other, but it has become clear that they can, in fact, be generated by the *same* psychology underpinning of the LSN. Indeed, as argued by Camerer (1989) and Rabin (2002), for outcomes of a random sequence, people prone to the gambler's fallacy expect more alternations than they actually occur. Consequently, when they *do* observe a long streak of positive outcomes, they overly attribute it to a positive mean rather than pure randomness, and this mistaken belief of a positive mean subsequently leads to "hot hands." Rabin and Vayanos (2010) show formally that the hot-hand fallacy can be derived from a model of the gambler's fallacy. In their model, a key conditional variable for belief formation is the length of the streak: with short streaks, people expect mean reversion, consistent with the gambler's fallacy; with longer streaks, people expect trend continuation, consistent with the hot-hand fallacy.

In addition to experimental evidence, field studies provide further support for the gambler's fallacy. For example, Chen, Moskowitz, and Shue (2016) find evidence of the gambler's fallacy in three separate high-stake settings: refugee asylum court decisions, loan application reviews, and Major League Baseball umpire pitch calls. More recently, Weber, Laudenbach, Wohlfart, and Weber (2023) survey retail investors at an online bank in Germany and find that the majority of them believe in a negative autocorrelation in stock returns.

3. The model

In this section, we develop an equilibrium model to study the trading and asset pricing implications of the LSN. We first describe the model's setup, then provide the model's solution, and finally discuss the model's implications.

3.1. Model setup

Asset space. We consider an infinite-horizon continuous-time model with two assets: a riskless asset with a constant interest rate r, and a risky asset. The risky asset has a fixed per-capita supply of Q, and its dividend payment evolves according to

$$dD_t = g_D dt + \sigma_D d\omega_t^D, \tag{1}$$

where ω_t^D is a standard Brownian motion. The price of the risky asset, denoted by P_t, is endogenously determined in equilibrium. In comparison, the riskless asset is in perfectly elastic supply.

Investor beliefs. We consider two types of investors: LSN investors and rational arbitrageurs. Rational arbitrageurs make up a fraction μ of the total population; LSN investors make up the remaining fraction of $1-\mu$.

To model beliefs under the LSN, we start by assuming that LSN investors do not directly observe the true price process and hence need to adopt a mental model and make inferences about future price changes. Specifically, we assume that LSN investors follow the belief structure proposed in Rabin and Vayanos (2010) to form a mental model about the risky asset's price. They perceive the price process as

$$dP_t = \theta_t dt + \sigma_P d\tilde{\omega}_t^P, \tag{2}$$

where θ_t represents the perceived quality of the asset and evolves according to

$$d\theta_t = \kappa(\overline{\theta} - \theta_t)dt + \sigma_\theta d\tilde{\omega}_t^\theta, \tag{3}$$

and $d\tilde{\omega}_t^P$ represents an innovation component. In equation (3), $\kappa > 0$ is a persistence parameter, $\overline{\theta}$ is the long-run mean of asset quality, and $d\tilde{\omega}_t^\theta$ represents a shock that is perceived by LSN investors to be independent of $d\tilde{\omega}_t^P$. Intuitively, parameter κ measures how quickly the asset's perceived quality θ_t changes over time: when κ increases, asset quality is expected to revert back to its long-run mean more quickly. Parameter σ_θ captures the size of perceived shocks to asset quality: when σ_θ increases, asset quality is more subject to random shocks and hence exhibits higher variability.

It is worth noting that equations (2) and (3) represent an incorrect mental model on the part of LSN investors; later in Section 3.3, we analyze investor beliefs and show that LSN investors and rational arbitrageurs hold distinct beliefs about future price changes. At the same time, such a mental model is also intuitive: when investors do not directly observe the true price process, they might naturally think of future price changes as coming from a persistent yet time-varying quality component and a transitory noise component. Moreover, this mental model serves as a foundation for LSN beliefs; if investors were able to directly observe the true price process, then no room is left for them to form incorrect beliefs.

We now introduce the LSN into investor beliefs. We follow Rabin (2002) and Rabin and Vayanos (2010) to assume that, in the perceived price process (2), the innovation term $d\tilde{\omega}_t^P$ is specified by

$$d\tilde{\omega}_t^P = d\tilde{\omega}_t - \alpha \left(\int_{-\infty}^{t} \delta e^{-\delta(t-s)} d\tilde{\omega}_s^P \right) dt. \tag{4}$$

That is, $d\tilde{\omega}_t^P$ contains two components: the first component, $d\tilde{\omega}_t$, is perceived by LSN investors to be a standard i.i.d. shock; the second component, $\int_{-\infty}^{t} \delta e^{-\delta(t-s)} d\tilde{\omega}_s^P$, is a weighted average of perceived price innovations from the past. Note that when $\alpha > 0$, $d\tilde{\omega}_t^P$ depends negatively on perceived price innovations from the past, capturing the gambler's fallacy in that any trends in the realization of past innovations are expected to revert in the near future. Further note that parameters α and δ measure two different aspects of the LSN. Parameter α measures the *strength* of the gambler's fallacy: a larger α means a stronger belief in trend reversion. Parameter δ measures the relative *weight* put on recent versus distant past realizations of $d\tilde{\omega}_s^P$: a larger δ implies higher relative weight placed on recent realizations, in which case perceived trend reversion applies primarily to recent trends as opposed to longer-term trends.

Equations (2) to (4) fully specify the beliefs of LSN investors. Below in Section 3.6, we consider a variant of the above belief system in which LSN investors form incorrect beliefs about future dividend changes.[3] We show that, under this alternative specification, the model's implications for investor beliefs remain similar.

Next, we turn to the rational arbitrageurs, who hold fully rational beliefs: they understand the dividend process in equation (1); they observe parameter μ and hence know the population fraction of LSN investors; and they are fully aware of the way in which LSN investors form beliefs about the risky asset price, as described by equations (2) to (4). Given their information set, rational arbitrageurs form correct beliefs about the evolution of the risky asset price. Given that P_t is endogenously determined in equilibrium, rational arbitrageurs' beliefs are also endogenously determined, in that they respond to the beliefs of LSN investors.

Investor preferences. Given that our focus is on investor beliefs rather than preferences, we adopt a parsimonious formalization of investor preferences: both LSN investors and rational arbitrageurs maximize instantaneous mean-variance preferences as in Greenwood and Vayanos (2014), specified by

$$\max_{N_t^i} \left(\mathbb{E}_t^i[dW_t^i] - \frac{\gamma}{2} \mathbb{V}\mathrm{ar}_t^i[dW_t^i] \right), \tag{5}$$

subject to the budget constraint on their wealth W_t^i

$$dW_t^i = rW_t^i dt - rN_t^i P_t dt + N_t^i dP_t + N_t^i D_t dt, \tag{6}$$

where N_t^i represents the per-capita share demand on the risky asset from investor i. Here, $i \in \{l, r\}$, where superscripts "l" and "r" represent LSN investors and rational arbitrageurs, respectively. Parameter γ represents risk aversion. For simplicity, γ is assumed to be the same for the two types of investors.

A common assumption made in the literature, one that is compatible with instantaneous mean-variance preferences, is that investors form overlapping generations (e.g., He and Krishnamurthy, 2013 and Greenwood and Vayanos, 2014). Specifically, for each generation of investor type i, it is endowed with Q shares of the risky asset and $W_t^i - Q \times P_t$ dollars of the riskless asset, lasts for dt period, and its wealth is then transferred to the next generation of the same investor type at the

[3] A large literature in behavioral finance directly specifies investors' incorrect beliefs about asset fundamentals; see, for example, Barberis et al. (1998), Scheinkman and Xiong (2003), Basak (2005), Hirshleifer, Li, and Yu (2015), and Nagel and Xu (2022).

end of the period.[4]

Market clearing. The share demands from LSN investors and rational arbitrageurs satisfy the following market clearing condition

$$\mu N_t^r + (1-\mu) N_t^l = Q \tag{7}$$

at each point in time t.

3.2. Model solution

We first note that LSN investors' beliefs, specified by equations (2) to (4), can be equivalently written as

$$dP_t = (\theta_t - \sigma_P \alpha \overline{\omega}_t)dt + \sigma_P d\tilde{\omega}_t \tag{8}$$

and

$$d\theta_t = \kappa(\overline{\theta} - \theta_t)dt + \sigma_\theta d\tilde{\omega}_t^\theta, \tag{9}$$

$$d\overline{\omega}_t = -(\alpha\delta + \delta)\overline{\omega}_t dt + \delta d\tilde{\omega}_t, \tag{10}$$

where $\overline{\omega}_t \equiv \int_{-\infty}^{t} \delta e^{-\delta(t-s)} d\tilde{\omega}_s^P$ and $\mathbb{E}_t^l[d\tilde{\omega}_t \cdot d\tilde{\omega}_t^\theta] = 0$. This alternative expression shows that the LSN enters the belief-formation process in two ways. First, in equation (8), LSN investors' perceived expected price change includes not only the perceived quality of the risky asset, θ_t, but also a contrarian component $-\sigma_P \alpha \overline{\omega}_t$. This contrarian term is directly derived from the assumption we have made in equation (4) about the gambler's fallacy. Second, in equation (10), $\overline{\omega}_t$ decays at the rate of $\alpha\delta + \delta$ rather than δ: $\overline{\omega}_t$ is constructed as a weighted average of past $d\tilde{\omega}_s^P$, where the declining weight leads to a baseline decay rate of δ in $\overline{\omega}_t$; moreover, the gambler's fallacy implies that LSN investors expect a negative serial autocorrelation in $d\tilde{\omega}_t^P$, causing an additional decay rate of $\alpha\delta$.

Note that, in the above belief-formation process, LSN investors do not observe θ_t and $\overline{\omega}_t$; as

[4] Alternatively, investors can be thought of as being infinitely-lived; but they reset their demand to Q shares every dt period.

in Rabin and Vayanos (2010), they use Bayesian inference to estimate both quantities.[5] Specifically, the information set at time t, $\mathcal{F}_t^P$, is defined using past risky asset prices $\{P_s, s \leq t\}$—that is, LSN investors update their beliefs about θ_t and $\overline{\omega}_t$ using past prices as informative signals. The conditional means and variances of $\boldsymbol{\theta}_t \equiv (\theta_t, \overline{\omega}_t)$ are denoted as

$$\boldsymbol{m}_t = (m_{t,1}, m_{t,2}) \equiv \mathbb{E}^l[(\theta_t, \overline{\omega}_t)|\mathcal{F}_t^P],$$

$$\boldsymbol{\gamma}_t = \begin{pmatrix} \gamma_{t,11} & \gamma_{t,12} \\ \gamma_{t,21} & \gamma_{t,22} \end{pmatrix} \equiv \mathbb{E}^l[(\boldsymbol{\theta}_t - \boldsymbol{m}_t)^T(\boldsymbol{\theta}_t - \boldsymbol{m}_t)|\mathcal{F}_t^P]. \tag{11}$$

We then apply Theorem 12.7 from Lipster and Shiryaev (2001) to the belief system of equations (8) to (10) and obtain

$$dP_t = (m_{t,1} - \sigma_P \alpha m_{t,2})dt + \sigma_P d\tilde{\omega}_t^l \tag{12}$$

and

$$dm_{t,1} = \kappa(\overline{\theta} - m_{t,1})dt + (\gamma_{11}\sigma_P^{-1} - \gamma_{12}\alpha)d\tilde{\omega}_t^l, \tag{13}$$

$$dm_{t,2} = -(\alpha\delta + \delta)m_{t,2}dt + (\delta + \gamma_{12}\sigma_P^{-1} - \gamma_{22}\alpha)d\tilde{\omega}_t^l, \tag{14}$$

where $d\tilde{\omega}_t^l$ is a Brownian shock perceived by LSN investors, and γ_{11}, γ_{12}, and γ_{22} are the stationary solutions for $\gamma_{t,11}$, $\gamma_{t,12}$, and $\gamma_{t,22}$, respectively. Note from equation (11) that $m_{t,1}$ and $m_{t,2}$ represent the inferred quantities of θ_t and $\overline{\omega}_t$.

Equations (12) to (14) allow us to directly link the evolution of past prices to LSN investors' inference process. Suppose that there is a large and positive price change. According to equation (12), LSN investors will attribute this positive price change to a positive perceived Brownian shock $d\tilde{\omega}_t^l$. Then, according to equation (13), this positive Brownian shock will lead LSN investors to infer a higher quality of the risky asset. At the same time, according to equation (14), the same shock will also lead LSN investors to infer stronger reversion in future price changes, since past prices have deviated substantially from the perceived trends. Therefore, in equation (12), the term $m_{t,1}$ represents an extrapolative component of LSN investors' beliefs as it depends positively on

[5] These estimated quantities in turn guide LSN investors' trading decisions.

price changes from the recent past, while the term $-\sigma_P \alpha m_{t,2}$ represents a contrarian component as it depends negatively on price changes from the recent past. Together, equations (12) to (14) fully characterize the inferences about the evolutions of P_t, $m_{t,1}$, and $m_{t,2}$ made by LSN investors; the derivation of these equations and the expressions of γ_{11}, γ_{12}, and γ_{22} are given in Appendix A.

Finally, we summarize the model's solution in the following proposition.

Proposition 1. (Model solution.) In the heterogeneous-agent model described above, the equilibrium price of the risky asset is

$$P_t = A + B \cdot m_{t,1} + C \cdot m_{t,2} + \frac{D_t}{r}. \tag{15}$$

The risky asset price P_t and the inferred means of the two state variables, $m_{t,1}$ and $m_{t,2}$, evolve according to

$$dP_t = \left[m_{t,1} - \sigma_P \alpha m_{t,2} + \sigma_P \cdot (l_0 + l_1 m_{t,1} + l_2 m_{t,2})\right] dt + \sigma_P d\omega_t^D, \tag{16}$$

$$dm_{t,1} = \left[\kappa(\overline{\theta} - m_{t,1}) + \sigma_{m1} \cdot (l_0 + l_1 m_{t,1} + l_2 m_{t,2})\right] dt + \sigma_{m1} d\omega_t^D, \tag{17}$$

and

$$dm_{t,2} = \left[-(\alpha\delta + \delta) m_{t,2} + \sigma_{m2} \cdot (l_0 + l_1 m_{t,1} + l_2 m_{t,2})\right] dt + \sigma_{m2} d\omega_t^D, \tag{18}$$

where ω_t^D is the standard Brownian motion from equation (1), $l_0 \equiv \sigma_D^{-1}(g_D + r\kappa B\overline{\theta})$, $l_1 \equiv -\sigma_D^{-1} r(1+\kappa B)$, $l_2 \equiv \sigma_D^{-1} r[\sigma_P \alpha - C(\alpha\delta + \delta)]$, $\sigma_{m1} \equiv \gamma_{11}\sigma_P^{-1} - \gamma_{12}\alpha$, $\sigma_{m2} \equiv \delta + \gamma_{12}\sigma_P^{-1} - \gamma_{22}\alpha$, and

$$\sigma_P = \frac{\sigma_D}{r} + \sigma_{m1} B + \sigma_{m2} C. \tag{19}$$

To solve for coefficients A, B, C and the price volatility σ_P, we first derive the optimal share demands for the risky asset from LSN investors and from the rational arbitrageurs

$$\begin{aligned} N_t^l &= \eta_0^l + \eta_1^l m_{t,1} + \eta_2^l m_{t,2}, \\ N_t^r &= \eta_0^r + \eta_1^r m_{t,1} + \eta_2^r m_{t,2}, \end{aligned} \tag{20}$$

where η_0^l, η_1^l, η_2^l, η_0^r, η_1^r, and η_2^r are expressed as functions of A, B, C, and σ_P. We then substitute equation (20) into the market clearing conditions in equation (7), which allows us to solve for A,

B, C, and σ_P through a system of simultaneous equations. ■

The proof of Proposition 1, the expressions of η_0^l, η_1^l, η_2^l, η_0^r, η_1^r, and η_2^r, and the numerical procedure that solves for A, B, C, and σ_P are given in Appendix B. In equation (15), A is a constant term, capturing investor risk aversion; B and C represent, respectively, the price impacts of the extrapolative and contrarian components of LSN investors' beliefs; and finally, $\frac{D_t}{r}$ represents a fundamental component of the risky asset price.

3.3. Model implications: investor beliefs

We start by examining the model's implications for investor beliefs. We first discuss parameter values. For asset parameters, we set: $g_D = 0.05$, $\sigma_D = 0.25$, $r = 0.025$, $Q = 1$. For risk preferences, we set $\gamma = 0.01$. Moreover, we set $\mu = 0.3$, so rational arbitrageurs make up 30% of the total population. We discuss our choice of belief parameters below.

No gambler's fallacy. We start with the benchmark case when there is no gambler's fallacy by setting $\alpha = 0$. In this case, equation (12) is reduced to $dP_t = m_{t,1}dt + \sigma_P d\tilde{\omega}_t^l$. Therefore, only the extrapolative component is at work. Furthermore, equation (13) is reduced to

$$dm_{t,1} = \kappa(\overline{\theta} - m_{t,1})dt + \gamma_{11}\sigma_P^{-1}d\tilde{\omega}_t^l, \tag{21}$$

where $\gamma_{11} = -\kappa\sigma_P^2 + \sqrt{(\kappa\sigma_P^2)^2 + \sigma_\theta^2\sigma_P^2}$ and is decreasing in κ. For belief parameters, we set $\overline{\theta} = g_D/r = 2$ and vary the values of κ and σ_θ for comparative statics. We first discretize the continuous-time model and simulate a time series of 10,000 years at the monthly frequency.[6] We then examine the properties of the model.

[Place Fig. 1 about here]

First, we analyze how, in the absence of the gambler's fallacy, investors beliefs about the future price change respond to past price changes in the model. Fig. 1 shows the sensitivity of beliefs to past price changes under different parameters of κ and σ_θ. Specifically, each line plots the coefficients from regressing LSN investors' beliefs about the future price change, $\mathbb{E}_t^l(dP_t)/dt = m_{t,1}$,

[6] In all simulation exercises, we use a value of 10 for the initial dividend level. Different initial dividend levels do not affect our model's implications.

on price changes over the past 60 months. In all these plots, beliefs load positively on past price changes, consistent with price change extrapolation. The intuition is straightforward: investors make inferences about the asset's quality by observing past price changes as informative signals.

In Panel A, we vary the value of κ between 0.01 and 1. In these plots, a smaller κ is associated with a higher degree of extrapolation. In other words, when investors perceive the asset's quality to be more persistent, they also extrapolate more from past price changes. The intuition can be seen from equations (9) and (21). With a small κ, the investors believe that the asset quality θ_t can persistently deviate from its long-term mean $\overline{\theta}$ and hence exhibit high variability. As such, when the investors observe a positive price change, they infer a large increase in $m_{t,1}$ and forecast a high price change moving forward. Conversely, with a large κ, the investors believe that θ_t tends to quickly mean-revert towards $\overline{\theta}$ and hence exhibits low variability. In this case, investors do not learn much about asset quality from price changes; when they observe a positive price change, they attribute most of it to a transitory shock—the term $\sigma_P d\tilde{\omega}_t^l$ in equation (12)—and only infer a small increase in $m_{t,1}$. As such, the investors do not significantly adjust their forecast of the future price change. In Panel B, we vary parameter σ_θ between 2.5 and 10. In these plots, a larger σ_θ is associated with a higher degree of extrapolation. When σ_θ is high, the investors perceive high variability of θ_t. Therefore, upon observing a positive price change, the investors infer a large increase in $m_{t,1}$, hence forecasting a high price change moving forward.

With gambler's fallacy. We now introduce the gambler's fallacy back to the model. Specifically, we set $\alpha = 0.2$, so that investors perceive random errors to be negatively autocorrelated. For the rest of the parameters, we set: $\kappa = 0.05$, $\overline{\theta} = g_D/r = 2$, $\sigma_\theta = 5$, and $\delta = 2.77$, where this value of δ indicates a look-back window of about six months; specifically, when forming beliefs about $\overline{\omega}_t$ in equation (10), LSN investors assign a 25% weight on an past innovation term from six months ago relative to the most recent past innovation. Given the above parameter values, we solve the model and obtain the following results. From Bayesian inference specified by equation (A.3) in Appendix A, we obtain $\gamma_{11} = 85.55$, $\gamma_{12} = -3.15$, and $\gamma_{22} = 0.12$. For the equilibrium price in equation (15), we obtain $A = -177.5$, $B = 4.01$, $C = -0.50$, and $\sigma_P = 24.97$. Finally, for the share demands described in equation (20), we obtain $\eta_0^l = 0.71$, $\eta_1^l = 0.14$, $\eta_2^l = -0.80$, $\eta_0^r = 1.67$, $\eta_1^r = -0.34$, and $\eta_2^r = 1.86$.

[Place Fig. 2 about here]

Fig. 2 shows the dependence of LSN beliefs on past price changes: the solid line plots the coefficients from regressing the LSN beliefs about the future price change on price changes over the past 60 months; here $\alpha = 0.2$. Consistent with the gambler's fallacy, LSN beliefs depend negatively on recent price changes, indicating that LSN investors expect recent trends to quickly reverse. At the same time, over longer horizons, the coefficients become positive, indicating extrapolative beliefs. To better understand the effect of the gambler's fallacy on investor beliefs, the dashed line plots the coefficients of the same regression for an investor with $\alpha = 0$. The comparison between the solid line and the dashed line shows that, over longer horizons, the coefficients under the $\alpha = 0.2$ case are more positive than those under the $\alpha = 0$ case. This suggests that, consistent with the result in Rabin and Vayanos (2010), the gambler's fallacy simultaneously generates contrarian beliefs over short-term trends and extrapolative beliefs over longer-term trends.

To further understand the extent to which these contrarian and extrapolative beliefs are biased, the dash-dot line plots the coefficients for regressing the *rational* beliefs about the future price in an economy where a fraction μ of investors are rational and the remaining fraction $1 - \mu$ have the LSN beliefs with $\alpha = 0.2$. The comparison between the solid line and the dash-dot line shows that, the contrarian beliefs lead LSN investors to underreact to short-term trends; moreover, the extrapolative beliefs lead LSN investors to overreact to longer-term trends.

[Place Fig. 3 about here]

Fig. 3 examines how the two belief parameters regulating the LSN, α and δ, affect the dependence of investor beliefs on past price changes. Panel A is concerned with α, which measures the overall strength of the gambler's fallacy. When α increases, not only does short-run mean-reversion increase in magnitude, longer-run extrapolation also increases. The simultaneous increases both short-term contrarian and long-term extrapolation behaviors confirms the LSN as the common driver of both phenomena. Panel B is concerned with δ, which measures the relative weight put on recent versus distant past innovation terms. When δ increases, investors believe that more recent trends tend to mean-revert more strongly. As such, after observing a long sequence of positive price changes, investors infer more strongly that the quality of the risky asset is high; in other words, they exhibit stronger extrapolative beliefs over long-term trends.

3.4. Model implications: trading behavior

We now turn to the model's implications for LSN investors' trading behavior. First, we examine how trading responds to past price changes. Next, we connect LSN investors' selling behavior to the disposition effect and describe a "doubling down" pattern in their buying behavior. Finally, we study the role of heterogeneous beliefs in driving different patterns of buying and selling behavior.

3.4.1. Trading responses to past price changes

To examine how trading responds to past price changes, we regress LSN investors' demand change, $N_t^l - Q$, on price changes over the past 60 months; Fig. 4 plots the regression coefficients. Given the assumption of mean-variance preferences, the sensitivity of trading to past price changes goes hand in hand with the sensitivity of beliefs to past price changes. In particular, Fig. 4 shows that LSN investors increase their holdings of the risky asset when the asset has recently gone down in value or when the asset has done well over a longer period of time.

[Place Fig. 4 about here]

Another way to establish the same intuition is by examining the price pattern before an investor buys or sells. In Fig. 5, Panel A plots the median price changes over the past 36 months prior to a buy; Panel B plots the median price changes over the past 36 months prior to a sell. Indeed, LSN investors tend to buy assets that have recently gone down in value but have done well over a longer period of time. Conversely, they sell assets that have recently gone up in value but have performed poorly over a longer period of time.

[Place Fig. 5 about here]

We summarize these findings in the following model prediction.

Prediction 1. (Trading response.) In the model described in Section 3.1, LSN investors, on average, buy assets with a negative short-term return and a positive long-term return, and sell assets with a positive short-term return and a negative long-term return.

3.4.2. The disposition effect

Given the contrarian beliefs over short-term trends, LSN can naturally generate the disposition effect, which says that investors tend to sell stocks trading at a gain and hold on to stocks trading at a loss. To examine the model's implications for trading behavior, we again discretize the model and simulate 10,000 years of monthly data. We adopt the baseline parameters specified in Section 3.3: $g_D = 0.05$, $\sigma_D = 0.25$, $r = 0.025$, $Q = 1$, $\gamma = 0.01$, $\mu = 0.3$, $\kappa = 0.05$, $\overline{\theta} = 2$, $\sigma_\theta = 5$, $\alpha = 0.2$, and $\delta = 2.77$. Then, at each point in time in this simulated time series, we check whether an LSN investor has a positive or negative demand change: a positive demand change counts as a "buy" and a negative one counts as a "sell."

In the prior literature studying the disposition effect, gain and loss are typically defined based on the purchase price or other plausible reference prices. In our model, however, investors continuously trade and almost never fully liquidate their positions in the risky asset. Given this, we look at the price change of the risky asset over four different horizons: the price change over the past month ("1M"), from one quarter ago to one month ago ("1M to 1Q"), from one year ago to one quarter ago ("1Q to 1Y"), and from five years ago to one year ago ("1Y to 5Y"). A positive price change counts as a "gain" and a negative one counts as a "loss." Combining the LSN investor's demand change with the price change of the risky asset, each point in time belongs to one of the four categories: "buy at gain," "sell at gain," "buy at loss," or "sell at loss." We can then compare the selling propensities between gains and losses to study the disposition effect in our model.

Table 1 shows that LSN investors display a disposition effect when gains and losses are defined based on price changes over the past month to the past quarter. This is because contrarian beliefs dominate investors' reactions to short-term trends. In comparison, investors display a reverse disposition effect when price changes are measured over a horizon that is longer than one year, because extrapolative beliefs dominate investors' reactions towards long-term trends. These findings lead to the following prediction about the disposition effect.

[Place Table 1 about here]

Prediction 2. (Disposition effect.) In the model described in Section 3.1, LSN investors display a disposition effect over short horizons: on average, they sell winners and hold on to losers, where winners and losers are defined by price changes over the last month to the last quarter.

Predictions 1 and 2 together suggest that a belief in the LSN can give rise to the coexistence of return extrapolation *and* the disposition effect. In particular, LSN investors hold extrapolative beliefs over long-term trends, causing them to have extrapolative demand. At the same time, they hold contrarian beliefs over short-term trends, causing them to display a disposition effect, in particular over short horizons. Taken together, these model implications mean that, return extrapolation and the disposition effect are not necessarily in conflict with each other. Instead, they are operating over different horizons and can be microfounded by the law of small numbers.

A related observation from Table 1 is that, over short horizons, LSN investors exhibit a "doubling down" pattern in buying: on average, their propensity to buy losers is significantly higher than their propensity to buy winners, where winners and losers are defined by price changes over the last month to the last quarter. This is an intuitive result—as discussed above, contrarian beliefs dominate investors' reactions to short-term trends—and we summarize it below.

Prediction 3. ("Doubling down" in buying behavior.) In the model described in Section 3.1, LSN investors exhibit a "doubling down" pattern in their buying behavior: on average, their propensity to buy losers is significantly higher than their propensity to buy winners, where winners and losers are defined by price changes over the last month to the last quarter.

Together, Predictions 2 and 3 establish that LSN investors are "contrarians" over short-term price trends. This trading pattern is supported by growing evidence from the field. For example, Kaniel, Saar, and Titman (2008) show that individuals tend to buy stocks following declines in the previous month and sell following price increases. More recently, Luo, Ravina, Sammon, and Viceira (2021) show that many retail investors trade as contrarians after large earnings surprises, especially for loser stocks, and that such contrarian trading contributes to post earnings announcement drift and price momentum; Kogan, Makarov, Niessner, and Schoar (2023) show that retail investors are contrarian when trading stocks but extrapolative when trading cryptos.

3.4.3. Heterogeneity

We now study the role of heterogeneous beliefs in driving different patterns of buying and selling behavior. We start by examining how the model-implied disposition effect varies as the two key belief parameters of LSN investors, α and δ, vary. Table 2 shows that a higher degree of the gambler's fallacy—measured by an increase in α—is associated with a stronger disposition effect when price changes are measured over the past month to the past quarter. In addition, when the look-back window is shorter—that is, when δ is higher—we also find a stronger disposition effect. These findings lead to the following prediction.

[Place Table 2 about here]

Prediction 4. (Disposition effect and the LSN.) In the model described in Section 3.1, investors with a stronger degree of the LSN beliefs, measured by either a higher α or a higher δ, display a stronger disposition effect.

We also note that, in our model, LSN beliefs are driving *both* buying and selling behavior. As such, there exists testable consistency between buying and selling behavior. On the one hand, Fig. 3 suggests that "doubling down" in buying behavior is more pronounced for investors with a stronger degree of the LSN beliefs, measured by either a higher α or a higher δ. On the other hand, Table 2 and Prediction 4 show that investors with a stronger degree of the LSN beliefs also display a stronger disposition effect. Taken together, our model makes the following prediction.

Prediction 5. (Consistency between buying and selling behavior.) In the model described in Section 3.1, investors who exhibit a stronger "doubling down" pattern in buying also exhibit a stronger disposition effect.

So far, we have looked at investors' buying and selling propensities separately for winning stocks and losing stocks. When computing these propensities—as presented in Tables 1 and 2—we have only checked whether a recent price change is positive or negative. We have not yet looked at how the *magnitude* of the recent price change affects investors' buying and selling propensities. We now examine the role of heterogeneous beliefs in driving the relationship between investors' buying or selling propensity and the magnitude of the recent price change. To do so, we analyze a more generalized model with three types of investors: LSN investors with $\alpha = 0.2$, LSN investors with $\alpha = 0$, and rational arbitrageurs.[7] We refer to LSN investors with $\alpha = 0$ as "extrapolators," because their beliefs about the future price change depend positively on past price changes. We

[7]The procedure that solves this more generalized model is given in Appendix C.

then refer to LSN investors with $\alpha = 0.2$ simply as "LSN investors."

[Place Fig. 6 about here]

Fig. 6 Panel A plots, separately for LSN investors and the extrapolators, the relationship between their buying propensity and the price change over the past one month. Fig. 6 Panel B plots, again for LSN investors and the extrapolators, the relationship between their selling propensity and the price change over the past one month. Fig. 6 shows that, in this more generalized model with three types of investors, LSN investors' buying propensity tends to depend negatively on recent price changes, while the extrapolators' buying propensity tends to depend positively on recent price changes. At the same time, LSN investors' selling propensity tends to depend positively on recent price changes, while the extrapolators' selling propensity tends to depend negatively on recent price changes. We summarize these results in the following model prediction.

Prediction 6. (Heterogeneous trading responses to past price changes.) In the more generalized model with three types of investors, LSN investors' buying propensity tends to depend negatively on recent price changes, while the extrapolators' buying propensity tends to depend positively on recent price changes. At the same time, LSN investors' selling propensity tends to depend positively on recent price changes, while the extrapolators' selling propensity tends to depend negatively on recent price changes.

3.5. Model implications: asset prices

In our model, asset prices are determined by the interaction between LSN investors and rational arbitrageurs. As discussed in Section 3.3, LSN investors hold contrarian beliefs over short-term trends and extrapolative beliefs over longer-term trends. As a result of market clearing, rational arbitrageurs must then hold the opposite beliefs—they have extrapolative beliefs over short-term trends and contrarian beliefs over longer-term trends, as we have observed from the dash-dot line in Fig. 2. These beliefs, by being fully rational, imply that asset prices exhibit short-term momentum and long-term reversals. Fig. 7 confirms this model implication. Specifically, at each point in time, we compute the price change over the next n months and the price change over the past n months; we then compute the time-series correlation between these two price changes. The figure plots the correlation as a function of n, where n goes from 1 to 60. For $n \leq 8$, the correlation is positive, indicating short-term momentum; for $9 < n < 60$, the correlation is negative, indicating long-term reversals.

[Place Figs. 7 and 8 about here]

Fig. 8 further examines how changes in the two belief parameters, α and δ, affect asset prices. It shows that an increase in α or δ gives rise to stronger patterns of short-term momentum and long-term reversals. With a higher α or a higher δ, the LSN beliefs become more contrarian over short-term trends and more extrapolative over longer-term trends; we have shown these results in Fig. 3. In response to these more pronounced LSN beliefs, the rational beliefs become more extrapolative over short-term trends and more contrarian over longer-term trends, implying stronger patterns of short-term momentum and long-term reversals.

[Place Fig. 9 about here]

LSN beliefs also lead to excess volatility: with the model parameters specified in Section 3.3, and in particular, with $\alpha = 0.2$, the implied volatility of price change, $\sigma_P = 24.97$, is significantly higher than the fundamental volatility of $\sigma_D/r = 10$. Fig. 9 further shows that, for a wide range of values of α and δ, our model generates excess volatility: σ_P remains significantly higher than σ_D/r.

3.6. Alternative specification of LSN beliefs

The baseline model described above applies the LSN to the price process. However, a large literature in behavioral finance directly specifies investors' incorrect beliefs about asset fundamentals (Barberis et al., 1998; Scheinkman and Xiong, 2003; Basak, 2005; Hirshleifer et al., 2015; Nagel and Xu, 2022). In this section, we follow this literature and apply the LSN to the dividend process.

Specifically, the true evolution of the risky asset's dividend payment is assumed to be

$$dD_t = g_D dt + \sigma_D d\omega_t^D. \tag{22}$$

However, LSN investors are now assumed to perceive the following dividend process

$$dD_t = \theta_t dt + \sigma_D d\tilde{\omega}_t^D, \qquad d\theta_t = \kappa(\bar{\theta} - \theta_t)dt + \sigma_\theta d\tilde{\omega}_t^\theta,$$
$$d\tilde{\omega}_t^D = d\tilde{\omega}_t - \alpha \left(\delta \int_{-\infty}^{t} e^{-\delta(t-s)} d\tilde{\omega}_s^D \right) dt. \tag{23}$$

In words, LSN investors perceive future dividend changes as coming from two components: a persistent yet time-varying component, and a transitory noise component that exhibits a negative serial autocorrelation. The rest of the model is identical to the baseline model; we leave the detailed description of the model to Appendix D.

[Place Fig. 10 about here]

Fig. 10 plots the dependence of the LSN and rational beliefs about the future price change, implied by this alternative model, on past price changes. The comparison between Fig. 10 and Fig. 2 shows that, similar to the baseline model, the alternative model again produces a dichotomy in belief formation: LSN investors's beliefs about future price changes depend negatively on recent price changes but positively on price changes from the distant past. Moreover, the model's implications for trading behavior and asset prices should also be similar to those from the baseline model; in both models, trading behavior and asset prices are completely driven by investor beliefs.

Despite the similarity between the two models, we view the baseline model as psychologically more realistic. In that model, the LSN is directly applied to the price process: LSN investors form incorrect beliefs about future price changes by looking at past price changes. The investors then use these beliefs about price changes to form their share demand of the risky asset. In other words, LSN investors apply a belief heuristic to directly guide their trading decisions. By contrast, under the alternative model, LSN investors need to *derive* beliefs about price changes from their beliefs about dividend changes as specified in equation (23); they then use these beliefs about price changes to make trading decisions. This extra step of deriving beliefs about prices from beliefs about dividends may not resemble the actual behavior of real-world investors.

4. Evidence from investor behavior

4.1. Data

The primary data set is from a large discount brokerage firm and contains individual-level transaction records from 1991 to 1996; more details of this data set can be found in Odean (1998) and Barber and Odean (2000). The data set specifies the date, price, transaction type (buy or sell), quantity, security type, security code, and commission paid for each trade investors have made during the sample period. Many other papers have used this data set to study investor behavior (e.g., Odean, 1998; Barber and Odean, 2000; Ben-David and Hirshleifer, 2012; Hartzmark, 2015). The data on stock prices and returns are from the Center for Research in Security Prices (CRSP).

We apply several filters to the original data set to construct the sample of transactions, which we later use to recover daily portfolio holdings. First, we follow Odean (1998) and drop observations that 1) are outside of the window between 1991 and 1996, 2) are not common-share transactions, and 3) have negative commissions. Second, similar to Hartzmark (2015), we drop an investor's entire transaction history of a stock if its position in the portfolio ever becomes negative, thereby allowing subsequent analysis to focus on long positions. This filter also excludes any trading history that starts with selling, making it plausible to calculate the purchase price for each position. In this filtered sample, the summary statistics of the transaction size, price per share, monthly turnover, commission, and spread resemble those reported by Barber and Odean (2000).[8]

4.2. Trading behavior: short-term contrarian and long-term extrapolation

4.2.1. Aggregate patterns

We start by examining the return patterns for stocks that investors tend to trade. As outlined in Section 3.4, Prediction 1 posits that investors exhibit a tendency to buy stocks that are short-term losers but long-term winners, and sell stocks that are short-term winners but long-term losers. To test this, Fig. 11 plots the aggregate return patterns leading up to a trade. Panel A specifically focuses on buying behavior, where each individual purchase is considered as a separate observation. We aggregate the lagged monthly market-adjusted return before the purchase takes place across all

[8] In addition to the "Odean data," we complement our analysis using data from a large brokerage firm in China.

purchases. To minimize the effects of outliers, we report the median buy rather than the average buy.

[Place Fig. 11 about here]

In Fig. 11, Panel A, a median buy is associated with the following return pattern: the stock tends to exhibit strong returns from approximately 36 months prior to the purchase up until around 5 months prior, but then experiences a decline in returns, including some periods of negative returns. This decrease in return is particularly evident for the most recent month, with a median lagged one-month return of approximately -1%. In Panel B, which concerns selling behavior, a median sell is associated with a rather different return pattern: the stock experiences consistently positive but moderate returns from 36 months ago up to around 2 months ago. However, there is a sudden and substantial increase in return for the most recent month, which suggests that investors tend to be more inclined to sell stocks that have recently experienced a rise in price.

Comparing between Figs. 5 and 11 suggests that the aggregate trading patterns observed in the Odean data are generally consistent with the model's predictions. However, there is one notable discrepancy: the model suggests that investors should sell long-term losers, whereas in the actual data, investors tend to sell both long-term winners and losers. This points to other mechanisms such as realization utility as potential drivers of trading behavior (Barberis and Xiong, 2009, 2012; Ingersoll and Jin, 2013).

4.2.2. Stock-level evidence

To provide further evidence in support of Prediction 1, we run stock-level regressions. Specifically, on each date, we aggregate all buys and sells for each stock as Buy and Sell. We then consider two measures of trading propensity: the first one is measured by (Buy − Sell)/(Buy + Sell) and the second one is simply Buy − Sell. We regress measures of trading propensity on past stock returns, controlling for date and stock fixed effects. Table 3 reports the results, with double-clustered standard errors being reported.

[Place Tables 3 and 4 about here]

Column (1) shows that there is heightened selling activity associated with stocks that have recently experienced price increases. In Column (2), we observe that the trading propensity shifts from selling to buying in response to more distant returns. This finding is consistent with Prediction 1, which suggests that investors, on average, tend to purchase stocks that are long-term winners but short-term losers. In Column (3), a different measure of trading activity shows consistent evidence that investors tend to buy short-term losers. In Column (4), this trading propensity decays over a longer horizon, but does not turn positive as in Column (2). We conduct similar analyses using a Chinese dataset and discover that the patterns are strikingly similar; the results are presented in Table 4. The primary difference is that Chinese retail investors exhibit excessive trading behavior, resulting in a much shorter look-back window compared to investors in the previous dataset. For instance, while the trading propensity in Table 3 flips signs for the lagged stock return from three quarters ago, in Table 4, the sign flips for the lagged return from about three weeks ago. This observation is in line with the literature on Chinese retail investors' trading behavior (Liu, Peng, Xiong, and Xiong, 2022).

4.3. The disposition effect

4.3.1. Aggregate evidence

Under our model of the LSN, investors expect short-term trends to reverse in the future. According to Prediction 2, this on average leads to the disposition effect: because investors expect current winners to underperform and current losers to outperform in the future, they tend to sell winners and hold on to losers. Furthermore, Prediction 2 suggests that the disposition effect is more prominent for short holding periods. As the holding period increases, investors' extrapolative beliefs begin to have a more significant impact on their reactions to long-term returns, thereby reducing the disposition effect.

Fig. 12 tests the prediction that the disposition effect should be more pronounced over short horizons. Panel A displays the propensities of selling winners and losers for daily portfolio holdings, confirming the existence of the disposition effect. On average, the probability of selling a winner stock is around 0.32%, while the probability of selling a loser stock is 0.23%. Panel B plots the probability of selling a winner stock and a loser stock for different holding periods. The holding period is measured as the time since the position was initially established. The results indicate that the disposition effect is much stronger for recently bought positions. For positions bought within the last month, the probability of selling a winner (1.2%) is almost twice as much as the probability of selling a loser (0.7%). However, these differences become smaller for positions held over a longer period. For positions held for more than a year, the propensities of selling winners and losers are virtually the same.

[Place Fig. 12 about here]

4.3.2. Additional buying behavior

Our model predicts that investors will exhibit a similar pattern when they buy additional shares of stocks they already own in their portfolio, in addition to the disposition effect. Prediction 3 states that LSN investors have a higher propensity to buy stocks that have recently decreased in value. This behavior of "doubling down" has been previously documented in Odean (1998), and is replicated in Fig. 13 Panel A. Overall, The probability of buying a winning stock already in the portfolio is less than 0.1%, while the probability of buying a losing stock already in the portfolio is almost 0.15%.

[Place Fig. 13 about here]

Panel B further breaks down the buying propensity based on the position's holding period. Overall, doubling down is present across all holding periods—but it is most pronounced for positions with a holding period between a month and a quarter. This goes in the opposite direction of our model predictions.

4.4. Investor heterogeneity

4.4.1. Disposition effect

Our model not only predicts the disposition effect and "doubling down" behavior, but also suggests a connection between these two phenomena at the investor level. According to Predictions 4 and 5, investors who hold stronger beliefs in the LSN are more likely to engage in both doubling down and the disposition effect. To test this prediction, we sort investors based on their degrees of doubling down. Specifically, for all investors who have made at least ten buys, we calculate the average one-month return across all buys for each investor and sort all investors into five groups. In Fig. 14 Panel A, we verify the validity of our sorting approach. As designed, the tendency of doubling down monotonically increases from Group 1 to Group 5. In the context of our model, one way of interpreting the five different groups is that Group 5 is the most prone to the LSN while Group 1 the least.

[Place Fig. 14 about here]

In Panel B, we then compare the selling propensities of gains and losses for the same five groups. Consistent with Predictions 4 and 5, the degree of the disposition effect also monotonically increases: conditional on a sale, the probability of selling a winner increases from 0.6 for Group 1 to 0.75. In fact, if we condition on positions with a shorter holding period, the increase in disposition effects is even sharper, as shown in Fig. 15.

[Place Fig. 15 about here]

The consistency between buying and selling behavior has two implications. First, it suggests that contrarian beliefs—the type we model through the LSN—are important for understanding the disposition effect. Previously, the disposition effect has often been attributed to explanations based on nonstandard preferences and other cognitive forces (Odean, 1998; Barberis and Xiong, 2009, 2012; Ingersoll and Jin, 2013; Chang et al., 2016; Frydman et al., 2018). Here, we show that a promising candidate explanation is contrarian beliefs induced by the LSN. Importantly, the LSN not only explains the disposition effect in selling, but also "doubleing down" in buying. Second and more broadly, it suggests that LSN beliefs, in general, are important for understanding both buying and selling behavior.

4.4.2. Trading activity as a function of past returns

Our model also makes predictions about how investors trade as a function of past returns. According to Prediction 6 from Section 3.4, LSN investors' buying propensity should depend negatively on recent returns, while their selling propensity should depend positively on recent returns. The opposite patterns should be observed for extrapolators. At first glance, this seems to go against the well-documented "V-shape" trading propensities in Ben-David and Hirshleifer (2012), which suggest that selling and buying propensities increase in the extremeness of returns.

We again sort investors into five groups based on their tendencies of "doubling down," as we have done in the previous section. For each group of investors, we then examine their buying and selling propensities as a function of the past month's returns; here returns are broadly classified into six subgroups, each with an 10% return interval. When calculating trading propensities, we examine daily portfolios as in Ben-David and Hirshleifer (2012) and calculate the probability of trading a particular position on a given day. We also take out rank effects as documented in Hartzmark (2015), which suggest that investors are more likely to trade positions that rank top or bottom in their portfolio. We do so by first estimating the magnitudes of the rank effects and then taking them out when calculating each subgroups' trading propensities. In general, the consideration of rank effects has little effects on our analysis.

[Place Figs. 16 and 17 about here]

Fig. 16 shows the results on selling propensity. Panel A first confirms the existence of the "V-shape" in selling. Panel B compares across the five groups sorted on investors' "doubling down" behaviors, where Group 1 is considered the most extrapolative and Group 5 most prone to the LSN. We find that the V-shape is much weaker among LSN investors: for positions with the past month's returns that are above –20%, the selling propensity monotonically increases in returns; this is consistent with Prediction 6. There is still a salience effect, in that LSN investors are more likely to sell extreme losers—those with the past month's returns below –20%—but the size of the V-shape is much smaller than in the aggregate sample. Fig. 17 shows the results on buying propensity. Again, we first document the existence of the V-shape in Panel A, and Panel B further compares across the five groups of investors. Consistent with Prediction 6, buying propensity monotonically decreases in returns for Group 5, one that is most prone to the LSN.

Taken together, these results not only provide further support to our LSN model, but also shed light on the nature of the V-shape trading propensities documented by Ben-David and Hirshleifer (2012). As we have shown above, this phenomenon is not present in *all* investors. Interestingly, it is among the most extrapolative investors that the V-shape is the most pronounced. Future work on understanding the V-shape should also be able to speak to the heterogeneous results we document here.

5. Evidence from asset prices

5.1. Data

In this section, we test the model's prediction about asset prices. In particular, the model predicts that, in the cross-section of individual stocks, those associated with more pronounced LSN beliefs should exhibit both stronger short-term momentum *and* stronger long-term reversal. Instead of using the Odean data, we test this prediction using quarterly holdings of mutual funds data, since the coverage is much more comprehensive and the price impacts of mutual funds are likely to be greater.

Our data cover all US equity mutual funds from 1980 to 2019. Quarterly fund holdings data are from the Thomson/Refinitiv Mutual Fund Holdings (S12) database. We follow the same procedure used in Peng and Wang (2023), which contains more details. In a nutshell, we 1) focus on funds that specialize in US equities, 2) require the reporting date and the filing date to be sufficiently close, 3) require the ratio of equity holdings to TNA to be close to one, 4) require a minimum fund size of $1 million, and 5) require that the TNAs reported in the Thomson Reuters database and in the CRSP database do not differ by more than a factor of two.

5.2. Results

5.2.1. Measuring the LSN

To measure a fund's degree of LSN, we first construct two measures based on mutual fund holdings. First, we measure fund j's holding-based demand for *long-term* returns in quarter q as

$$LongRet_{j,q}^{fund} = \frac{\sum_i Dollar_{i,q} \times LongRet_{i,q}}{\sum_i Dollar_{i,j,q}}, \tag{24}$$

where $Dollar_{i,j,q}$ is the dollar amount of stock i held by fund j at the end of quarter q, and $LongRet_{i,q}$ is stock i's past five-year return by the end of quarter q. Second, we measure fund j's holding-based demand for *short-term* returns in quarter q as

$$ShortRet^{fund}_{j,q} = \frac{\sum_i Dollar_{i,q} \times ShortRet_{i,q}}{\sum_i Dollar_{i,j,q}}, \tag{25}$$

where $Dollar_{i,j,q}$ is the dollar amount of stock i held by fund j at the end of quarter q, and $ShortRet_{i,q}$ is stock i's past quarterly return by the end of quarter q.

A fund's degree of LSN, denoted by $FundLSN$, is then constructed as

$$FundLSN_{j,q} = LongRet^{fund}_{j,q} - ShortRet^{fund}_{j,q}. \tag{26}$$

The idea is that funds more prone to the LSN are more likely to hold stocks with good returns over the long-run but poor returns in more recent periods.

Next, we aggregate fund-level factor demand to the stock-level in each quarter as

$$\overline{LSN}_{i,q} = \frac{\sum_i shares_{i,j,q} \times FundLSN_{j,q}}{\sum_i shares_{i,j,q}}, \tag{27}$$

where $\overline{LSN}_{i,q}$ measures the degree of LSN of the underlying invesors holding stock i in quarter q.

5.2.2. Cross-sectional return predictability

To test the model's predictions on cross-sectional return predictability, at the end of each quarter, all stocks are independently sorted into 25 portfolios based on their past one-year returns and $\overline{LSN}_{i,q}$, where $\overline{LSN}_{i,q}$ measures underlying funds' degree of LSN. To address potential microstructure issues and focus on mutual fund behavior, we exclude stocks with a price below five dollars, a total mutual fund ownership below 1%, or a market capitalization in the bottom decile.

To illustrate the impact of LSN, we take the difference between the LSN momentum return—that is, the return of the winner-minus-loser strategy conditional on stocks in the highest decile based on $\overline{LSN}_{i,q}$—and the unconditional momentum return. Figure 18 shows the results. Consistent with the model's prediction, we see that the LSN momentum return is stronger initially and then falls down in later quarters.

6. Conclusion

The law of small numbers, a prominent type of incorrect belief, has received wide support from experimental and field studies. In this paper, we incorporate it into a tractable equilibrium asset pricing model. We study the implications of the LSN for trading behavior and asset prices.

We show that the LSN beliefs helps explain the coexistence of the disposition effect and return extrapolation: investors sell assets whose prices have recently gone up, but they buy assets whose prices have gone up for multiple periods in a row. The LSN beliefs also give rise to excess volatility, short-term momentum, and long-term reversals. Moreover, the model makes additional predictions: the disposition effect is more pronounced over shorter holding periods; investors exhibit "doubling down" in their buying behavior; investors who exhibit a stronger "doubling down" pattern in buying behavior also exhibit a stronger disposition effect; and trading responses to past returns vary significantly with investors' degree of the LSN beliefs. We empirically test and confirm each of these predictions using account-level transaction data.

References

Bar-Hillel, M., Wagenaar, W. A., 1991. The perception of randomness. Advances in Applied Mathematics 12, 428–454.

Barber, B. M., Odean, T., 2000. Trading is Hazardous to Your Wealth: The Common Stock Investment Performance of Individual Investors. Journal of Finance 55, 773–806.

Barberis, N., Shleifer, A., Vishny, R., 1998. A Model of Investor Sentiment. Journal of Financial Economics 49, 307–343.

Barberis, N., Xiong, W., 2009. What Drives the Disposition Effect? An Analysis of a Long-Standing Preference-Based Explanation. Journal of Finance 64, 751–784.

Barberis, N., Xiong, W., 2012. Realization Utility. Journal of Financial Economics 104, 251–271.

Basak, S., 2005. Asset Pricing with Heterogeneous Beliefs. Journal of Banking and Finance 29, 2849–2881.

Ben-David, I., Hirshleifer, D., 2012. Are Investors Really Reluctant to Realize Their Losses? Trading Responses to Past Returns and the Disposition Effect. Review of Financial Studies 25.

Camerer, C. F., 1989. Does the Basketball Market Believe in the 'Hot Hand,'? American Economic Review 79, 1257–1261.

Chang, T. Y., Solomon, D. H., Westerfield, M. M., 2016. Looking for Someone to Blame: Delegation, Cognitive Dissonance and the Disposition Effect. Journal of Finance 71, 267–302.

Chen, D. L., Moskowitz, T. J., Shue, K., 2016. Decision making under the gambler's fallacy: Evidence from asylum judges, loan officers, and baseball umpires. The Quarterly Journal of Economics 131, 1181–1242.

Frydman, C., Hartzmark, S. M., Solomon, D. H., 2018. Rolling mental accounts. The Review of Financial Studies 31, 362–397.

Gilovich, T., Vallone, R., Tversky, A., 1985. The hot hand in basketball: On the misperception of random sequences. Cognitive psychology 17, 295–314.

Greenwood, R., Vayanos, D., 2014. Bond Supply and Excess Bond Returns. Review of Financial Studies 27, 663–713.

Hartzmark, S. M., 2015. The Worst, the Best, Ignoring All the Rest: The Rank Effect and Trading Behavior. Review of Financial Studies 28, 1024–1059.

He, Z., Krishnamurthy, A., 2013. Intermediary Asset Pricing. American Economic Review 103, 732–770.

Hirshleifer, D., Li, J., Yu, J., 2015. Asset Pricing in Production Economies with Extrapolative Expectations. Journal of Monetary Economics 76, 87–106.

Ingersoll, J. E., Jin, L. J., 2013. Realization Utility with Reference-Dependent Preferences. Review of Financial Studies 26, 723–767.

Kaniel, R., Saar, G., Titman, S., 2008. Individual Investor Trading and Stock Returns. Journal of Finance 63, 273–310.

Kogan, S., Makarov, I., Niessner, M., Schoar, A., 2023. Are cryptos different? evidence from retail trading. Working paper .

Liao, J., Peng, C., Zhu, N., 2022. Extrapolative bubbles and trading volume. The Review of Financial Studies 35, 1682–1722.

Lipster, R., Shiryaev, A., 2001. Statistics of Random Processes I, General Theory. Springer, New York.

Liu, H., Peng, C., Xiong, W. A., Xiong, W., 2022. Taming the bias zoo. Journal of Financial Economics 143, 716–741.

Luo, C., Ravina, E., Sammon, M., Viceira, L. M., 2021. Retail investors' contrarian behavior around news and the momentum effect. Available at SSRN 3544949 .

Nagel, S., Xu, Z., 2022. Asset Pricing with Fading Memory. Review of Financial Studies 35, 2190–2245.

Odean, T., 1998. Are Investors Reluctant to Realize Their Losses? Journal of Finance 53, 1775–1798.

Peng, C., Wang, C., 2023. Factor demand and factor returns. Working paper .

Rabin, M., 2002. Inference by Believers in the Law of Small Numbers. Quarterly Journal of Economics 117, 775–816.

Rabin, M., Vayanos, D., 2010. The Gambler's and Hot-hand Fallacies: Theory and Applications. Review of Economic Studies 77, 730–778.

Rapoport, A., Budescu, D. V., 1992. Generation of random series in two-person strictly competitive games. Journal of Experimental Psychology: General 121, 352.

Rapoport, A., Budescu, D. V., 1997. Randomization in individual choice behavior. Psychological Review 104, 603.

Scheinkman, J., Xiong, W., 2003. Overconfidence and Speculative Bubbles. Journal of Political Economy 111, 1183–1220.

Teguia, A., 2017. Law of small numbers and hysteresis in asset prices and portfolio choices. Tech. rep.

Tversky, A., Gilovich, T., 1989a. The cold facts about the "hot hand" in basketball. Chance 2, 16–21.

Tversky, A., Gilovich, T., 1989b. The "hot hand": Statistical reality or cognitive illusion? Chance 2, 31–34.

Tversky, A., Kahneman, D., 1971. Belief in the law of small numbers. Psychological bulletin 76, 105.

Weber, A., Laudenbach, C., Wohlfart, J., Weber, R., 2023. Beliefs about the stock market and investment choices: Evidence from a field experiment. Working paper .

Panel A: Effect of κ

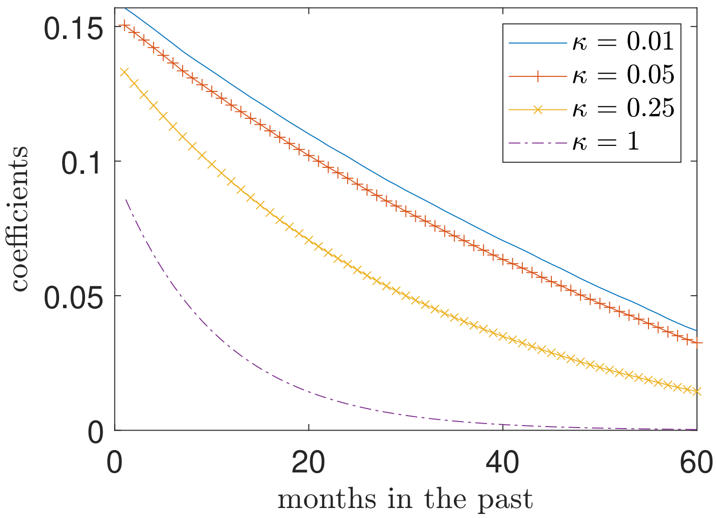

Panel B: Effect of σ_θ

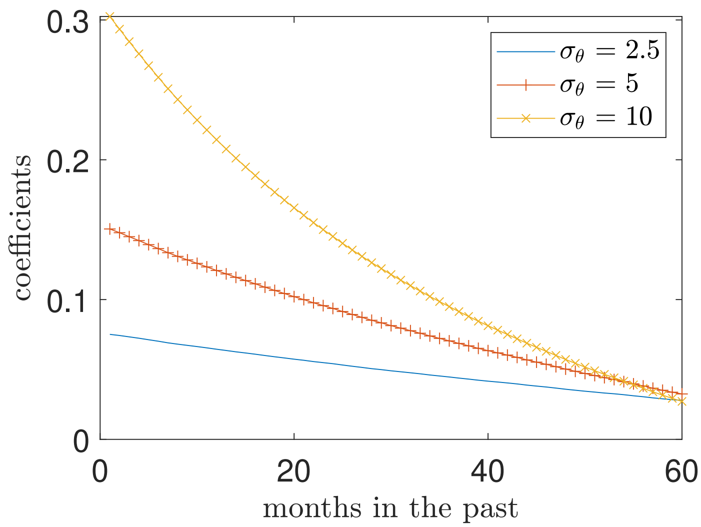

Fig. 1. Dependence of the LSN beliefs on past price changes: The $\alpha = 0$ case. The figure plots, for different values of κ and σ_θ, the coefficients from regressing LSN investors' beliefs about the future price change, $\mathbb{E}_t^l(dP_t)/dt = m_{t,1}$, on price changes over the past 60 months. The default values of κ and σ_θ are 0.05 and 5, respectively. The other parameters are: $g_D = 0.05$, $\sigma_D = 0.25$, $r = 0.025$, $Q = 1$, $\delta = 2.77$, $\alpha = 0$, $\overline{\theta} = 2$, $\gamma = 0.01$, and $\mu = 0.3$.

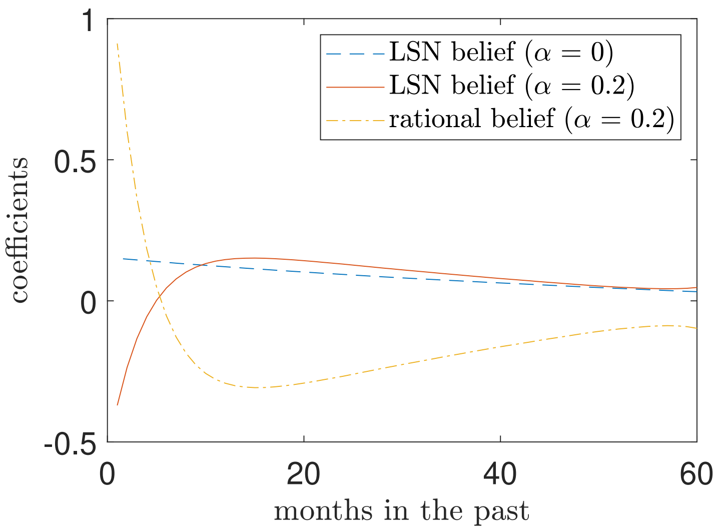

Fig. 2. Dependence of the LSN and rational beliefs on past price changes. The figure plots the coefficients from regressing either LSN investors' beliefs about the future price change—$\mathbb{E}_t^l(dP_t)/dt = m_{t,1} - \sigma_P \alpha m_{t,2}$—or the rational investors' beliefs about the future price change—$\mathbb{E}_t^r(dP_t)/dt = m_{t,1} - \sigma_P \alpha m_{t,2} + \sigma_P(l_0 + l_1 m_{t,1} + l_2 m_{t,2})$—on price changes over the past 60 months. We first consider an economy where a fraction μ of investors are rational and the remaining fraction $1 - \mu$ have the LSN belief with $\alpha = 0$; this is a benchmark case with no gambler's fallacy. For this case, the dashed line plots the coefficients for the LSN belief. We then consider an economy where a fraction μ of investors are rational and the remaining fraction $1 - \mu$ have the LSN belief with $\alpha = 0.2$. For this case, the solid line plots the coefficients for the LSN belief; as a comparison, the dash-dot line plots the coefficients for the rational belief. The other parameter values are: $g_D = 0.05$, $\sigma_D = 0.25$, $r = 0.025$, $Q = 1$, $\kappa = 0.05$, $\sigma_\theta = 5$, $\delta = 2.77$, $\overline{\theta} = 2$, $\gamma = 0.01$, and $\mu = 0.3$.

Panel A: Effect of α

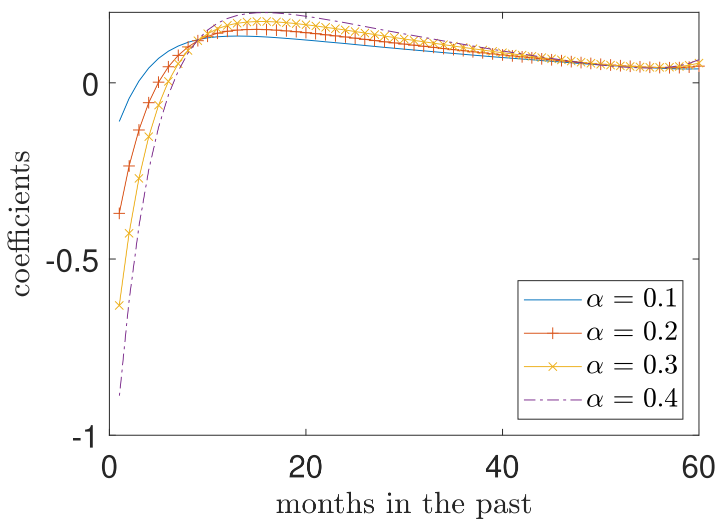

Panel B: Effect of δ

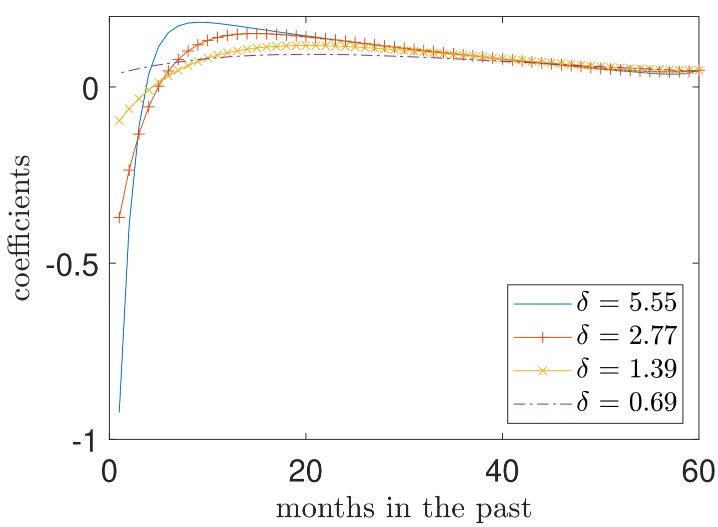

Fig. 3. Dependence of the LSN beliefs on past price changes: Comparative statics. The figure plots, for different values of α and δ, the coefficients from regressing LSN investors' beliefs about the future price change, $\mathbb{E}_t^l(dP_t)/dt = m_{t,1} - \sigma_P \alpha m_{t,2}$, on price changes over the past 60 months. The default values of α and δ are 0.2 and 2.77, respectively. The other parameters are: $g_D = 0.05$, $\sigma_D = 0.25$, $r = 0.025$, $Q = 1$, $\kappa = 0.05$, $\sigma_\theta = 5$, $\overline{\theta} = 2$, $\gamma = 0.01$, and $\mu = 0.3$.

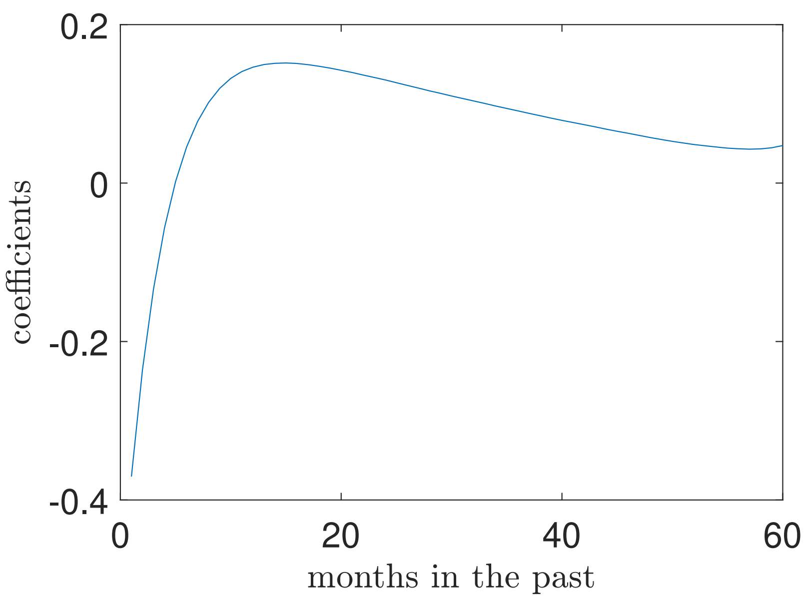

Fig. 4. Dependence of the change in LSN investors' demand on past price changes. The figure plots the coefficients from regressing the change in LSN investors' demand on the risky asset, $N_t^l - Q$, on price changes over the past 60 months. The parameter values are: $g_D = 0.05$, $\sigma_D = 0.25$, $r = 0.025$, $Q = 1$, $\kappa = 0.05$, $\sigma_\theta = 5$, $\alpha = 0.2$, $\delta = 2.77$, $\overline{\theta} = 2$, $\gamma = 0.01$, and $\mu = 0.3$.

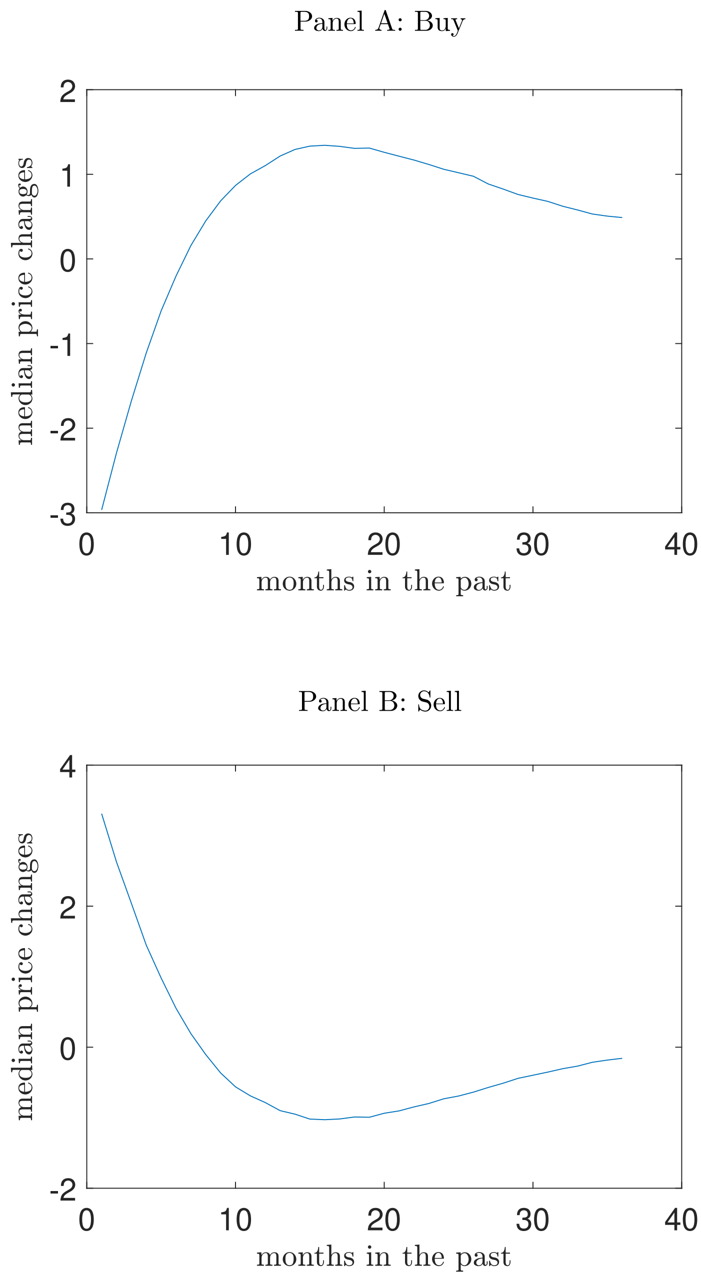

Fig. 5. Pattern of price changes before trading. Panel A plots the median price changes over the past 36 months prior to a buying decision. Panel B plots the median price changes over the past 36 months prior to a selling decision. The parameter values are: $g_D = 0.05$, $\sigma_D = 0.25$, $r = 0.025$, $Q = 1$, $\kappa = 0.05$, $\sigma_\theta = 5$, $\alpha = 0.2$, $\delta = 2.77$, $\overline{\theta} = 2$, $\gamma = 0.01$, and $\mu = 0.3$.

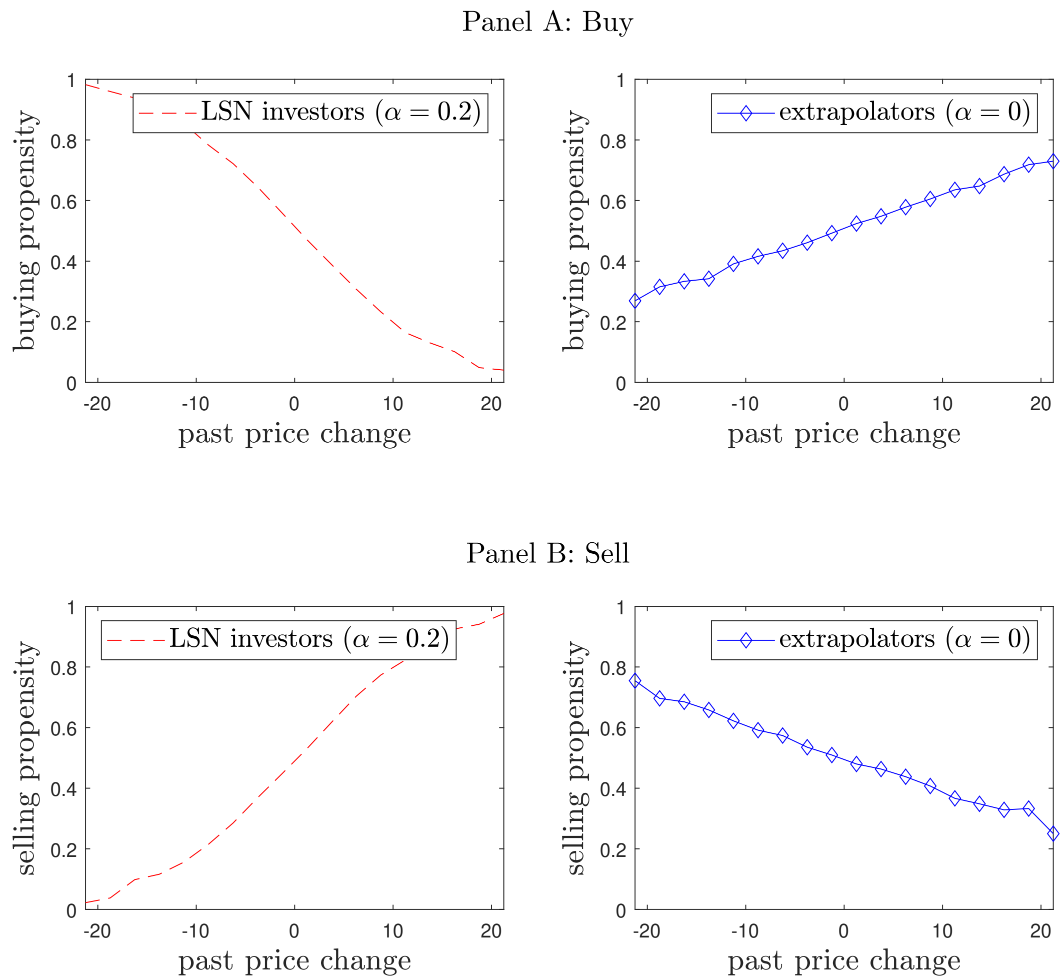

Fig. 6. Heterogeneous trading responses to past price changes. We analyze a model with three types of investors: LSN investors with $\alpha = 0.2$, LSN investors with $\alpha = 0$ (referred to as "extrapolators"), and rational arbitrageurs. Panel A plots, separately for LSN investors and extrapolators, the relationship between their buying propensity and the price change from the past one month. Panel B plots, again for LSN investors and extrapolators, the relationship between their selling propensity and the price change from the past one month. LSN investors make up 35% of the total population; the extrapolators make up 35%; and rational arbitrageurs make up the remaining 30%. The other parameter values are: $g_D = 0.05$, $\sigma_D = 0.25$, $r = 0.025$, $Q = 1$, $\kappa = 0.05$, $\sigma_\theta = 5$, $\delta = 2.77$, $\overline{\theta} = 2$, and $\gamma = 0.01$.

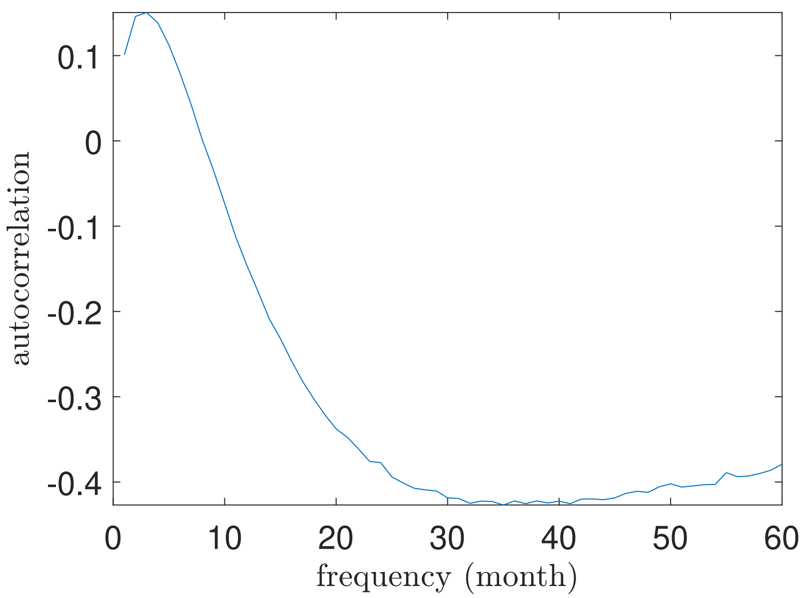

Fig. 7. Autocorrelation of price changes. At each point in time, we compute the price change over the next n months and the price change over the past n months; we then compute the time-series correlation between these two price changes. The figure plots the correlation as a function of n, where n goes from 1 to 60. The parameter values are: $g_D = 0.05$, $\sigma_D = 0.25$, $r = 0.025$, $Q = 1$, $\kappa = 0.05$, $\sigma_\theta = 5$, $\alpha = 0.2$, $\delta = 2.77$, $\overline{\theta} = 2$, $\gamma = 0.01$, and $\mu = 0.3$.

Panel A: Effect of α

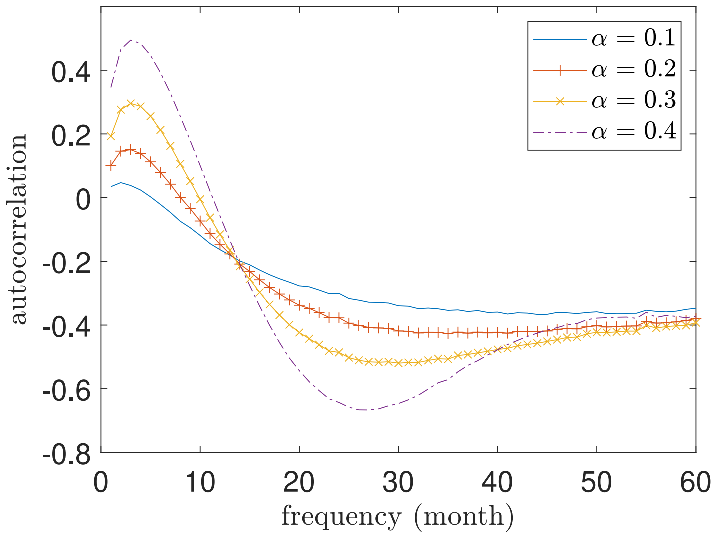

Panel B: Effect of δ

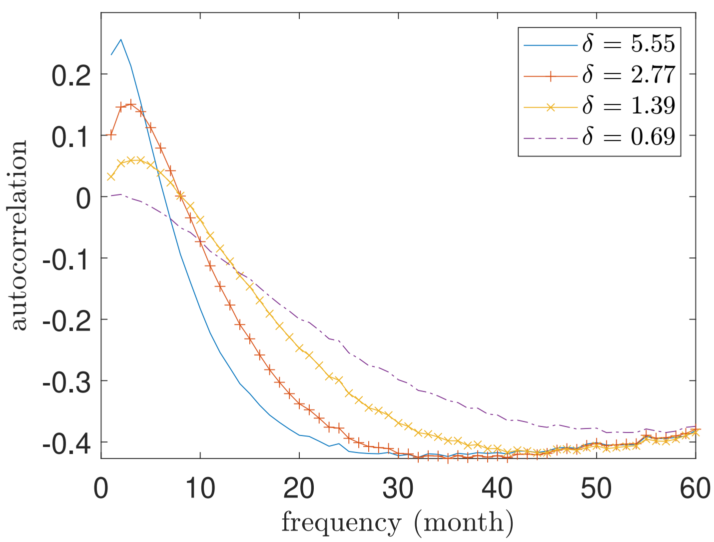

Fig. 8. Autocorrelation of price changes: Comparative statics. The figures plot, for different values of α and δ, the time-series correlation between the price change over the next n months and the price change over the past n months, where n goes from 1 to 60. The default values of α and δ are 0.2 and 2.77, respectively. The other parameter values are: $g_D = 0.05$, $\sigma_D = 0.25$, $r = 0.025$, $Q = 1$, $\kappa = 0.05$, $\sigma_\theta = 5$, $\overline{\theta} = 2$, $\gamma = 0.01$, and $\mu = 0.3$.

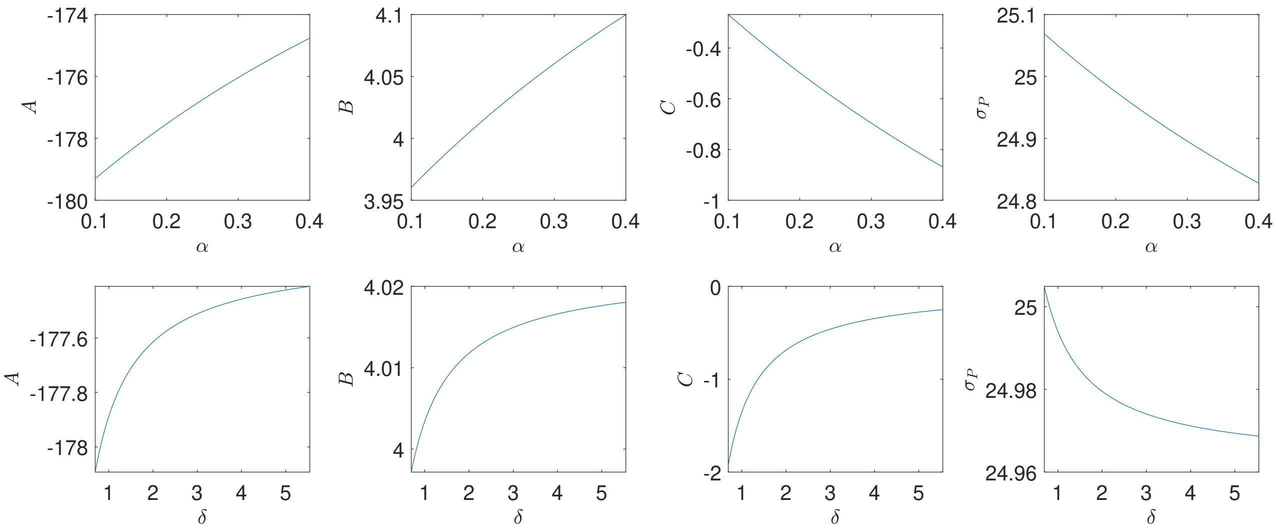

Fig. 9. Model solution as function of α and δ. The upper panel plots the model solution—the coefficients A, B, and C, and the price volatility σ_P—as function of α. The lower panel plots the same quantities as function of δ. The default values of α and δ are 0.2 and 2.77, respectively. The other parameters are: $g_D = 0.05$, $\sigma_D = 0.25$, $r = 0.025$, $Q = 1$, $\kappa = 0.05$, $\sigma_\theta = 5$, $\overline{\theta} = 2$, $\gamma = 0.01$, and $\mu = 0.3$.

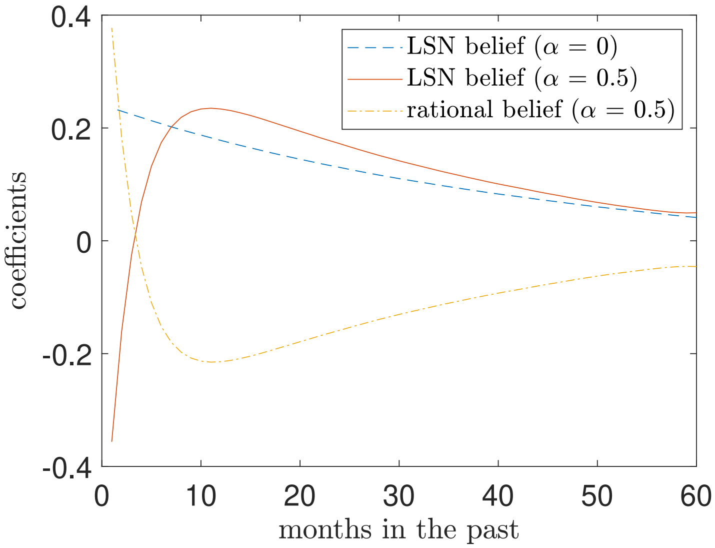

Fig. 10. Dependence of the LSN and rational beliefs about the future price change, implied by the alternative model specified in Section 3.6 and Appendix D, on past price changes. The figure plots the coefficients from regressing either LSN investors' beliefs about the future price change, $\mathbb{E}_t^l(dP_t)/dt$, or the rational investors' beliefs about the future price change, $\mathbb{E}_t^r(dP_t)/dt$, on price changes over the past 60 months. We first consider an economy where a fraction μ of investors are rational and the remaining fraction $1-\mu$ have the LSN belief with $\alpha = 0$; this is a benchmark case with no gambler's fallacy. For this case, the dashed line plots the coefficients for the LSN belief. We then consider an economy where a fraction μ of investors are rational and the remaining fraction $1-\mu$ have the LSN belief with $\alpha = 0.5$. For this case, the solid line plots the coefficients for the LSN belief; as a comparison, the dash-dot line plots the coefficients for the rational belief. The other parameter values are: $g_D = 0.05$, $\sigma_D = 0.25$, $r = 0.025$, $Q = 1$, $\kappa = 0.05$, $\sigma_\theta = 0.125$, $\delta = 2.77$, $\overline{\theta} = 0.05$, $\gamma = 0.01$, and $\mu = 0.5$.

Panel A: Buy

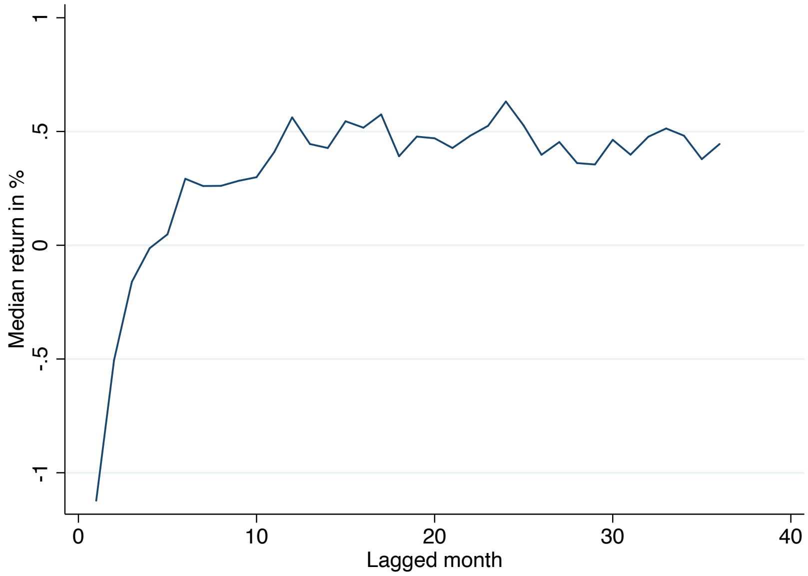

Panel B: Sell

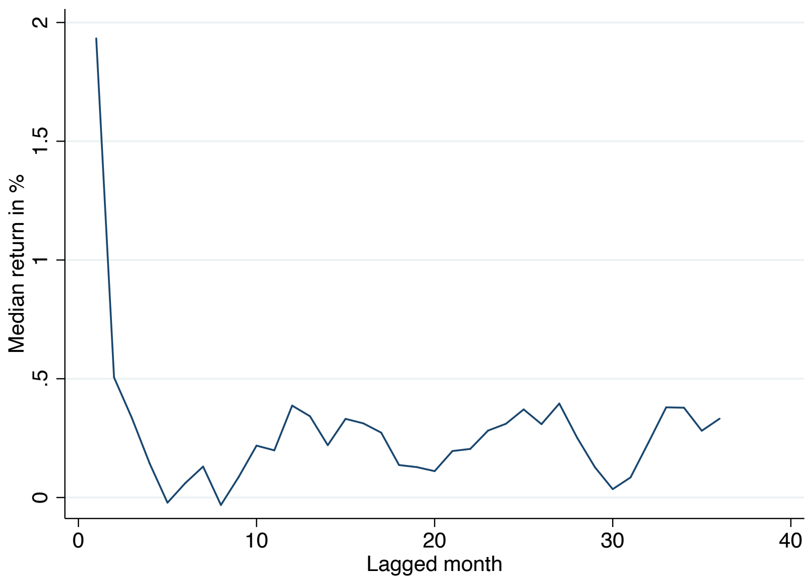

Fig. 11. Return patterns before trading. This figure plots the return patterns before buys and sells, using transactions observed in the Odean data. In Panel A, each buy is considered as a separate observation, and we aggregate across all buys the lagged monthly market-adjusted return before the buy takes place. The line plots the median monthly return across all observations. In Panel B, each sell is considered as a separate observation, and the line plots the median monthly market-adjusted return across all observations

Panel A: Disposition effect

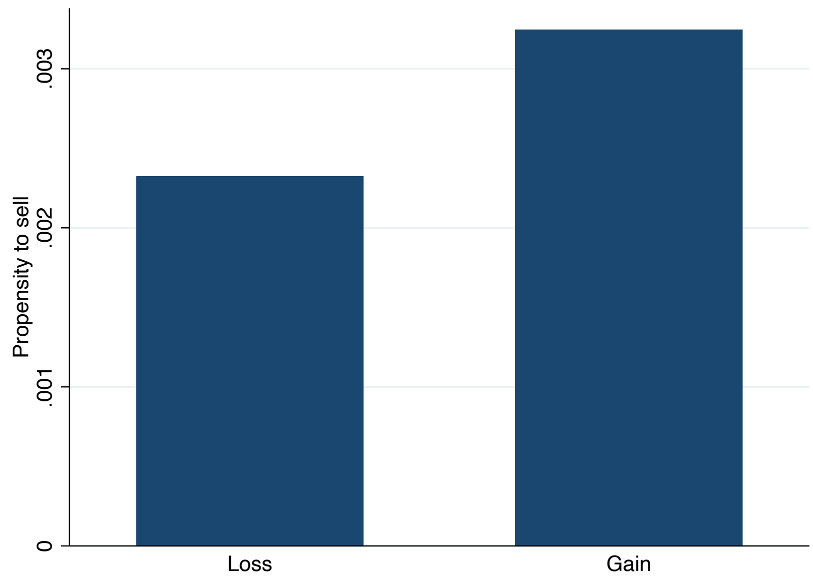

Panel B: Disposition effect at different holding periods

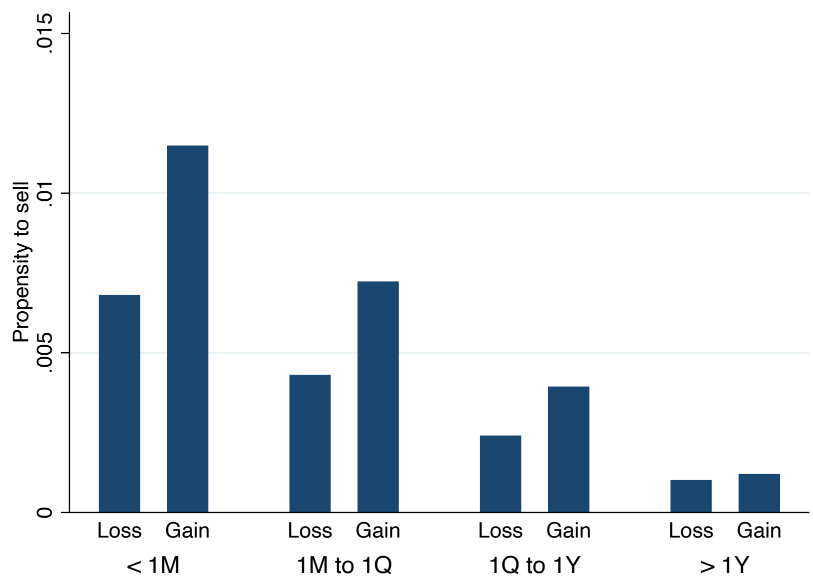

Fig. 12. Disposition effect.

Panel A: Additional buying

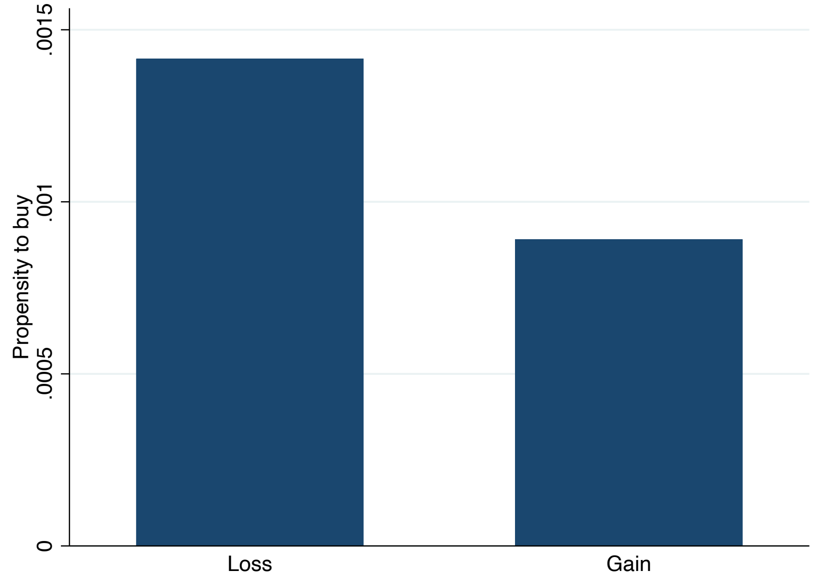

Panel B: Additional buying at different holding periods

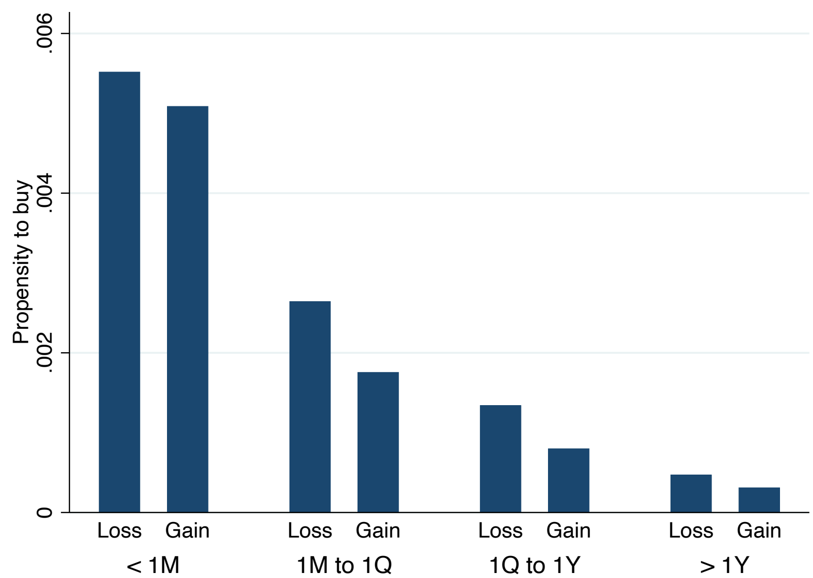

Fig. 13. Additional buying.

Panel A: Buying

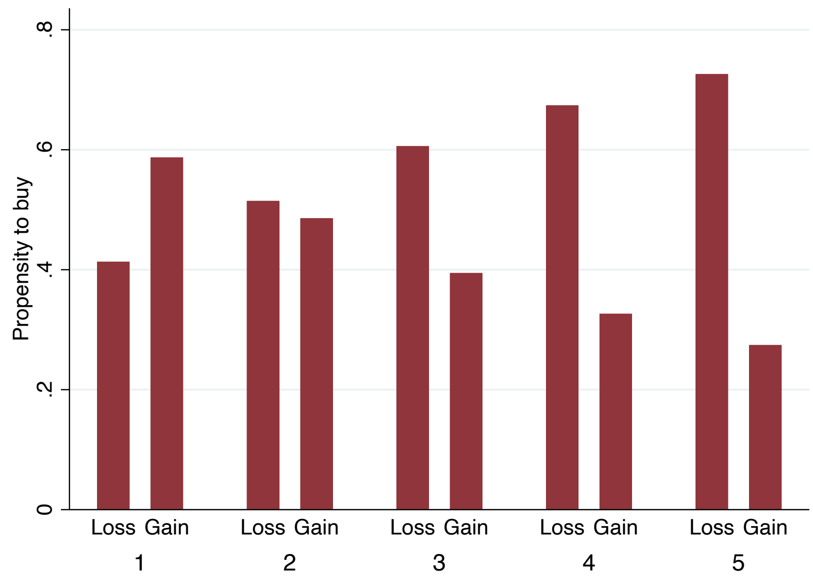

Panel B: Selling

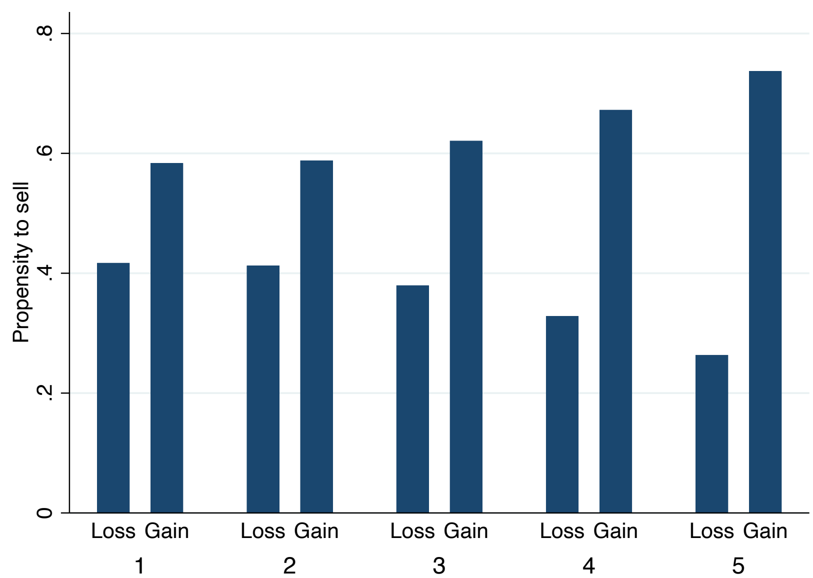

Fig. 14. Consistency between buying and selling behavior.

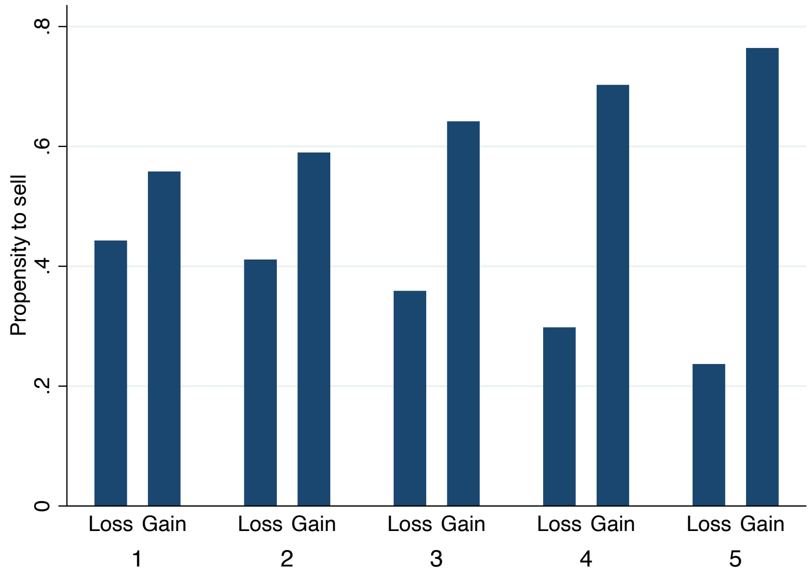

Fig. 15. Consistency between buying and selling behavior (one-month holding period).

Panel A: Aggregate

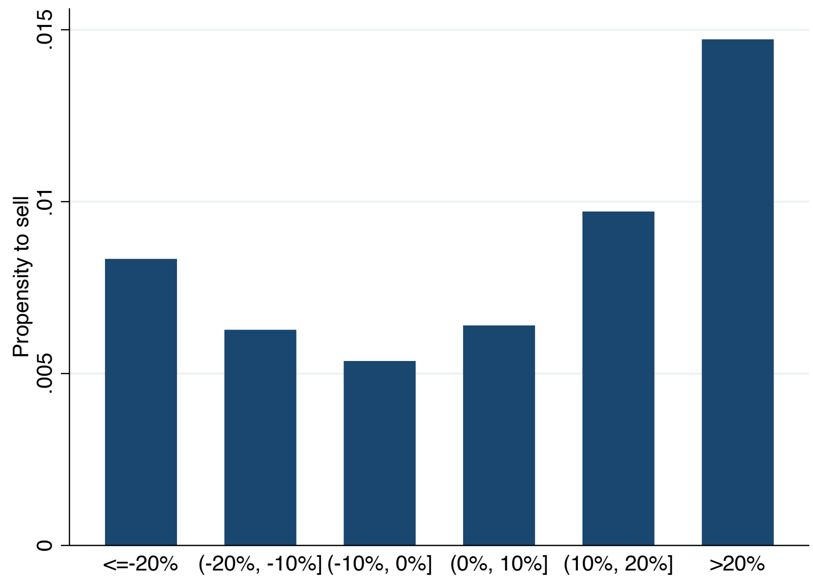

Panel B: By group

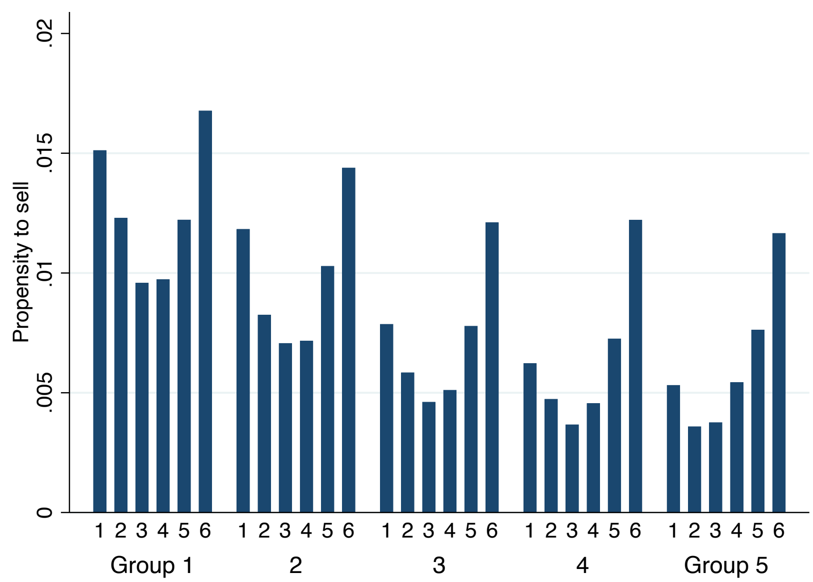

Fig. 16. V-shape selling behavior.

Panel A: Aggregate

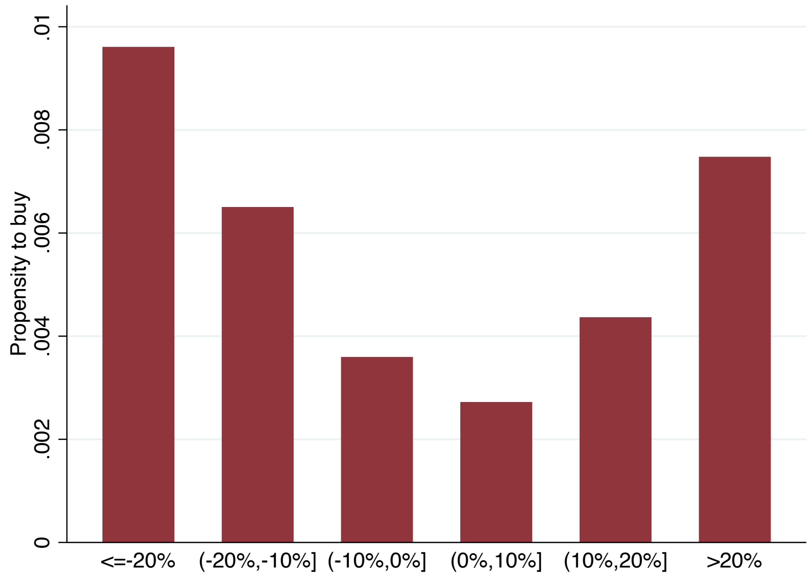

Panel B: By group

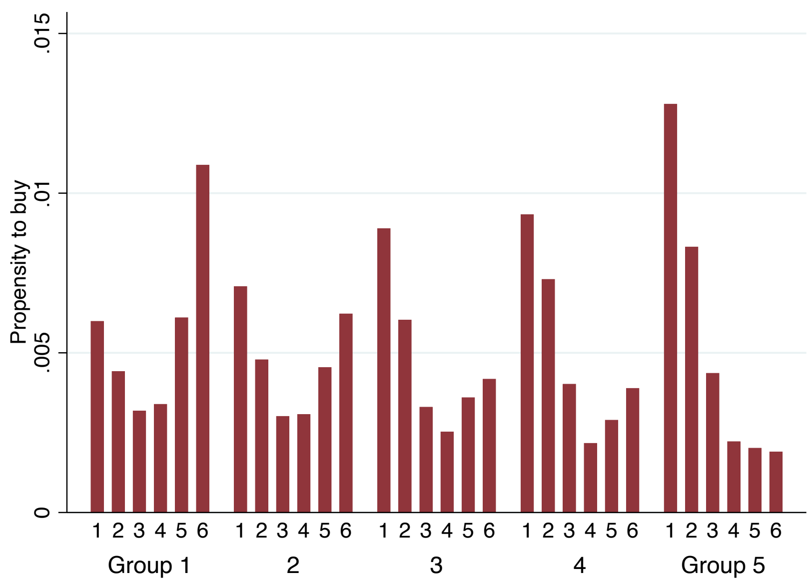

Fig. 17. V-shape buying behavior.

Panel A: Equal-weighted

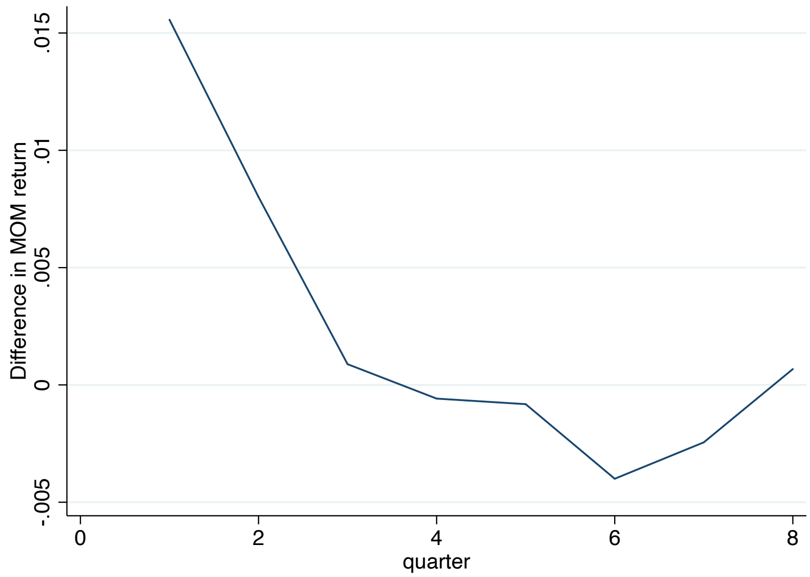

Panel B: Value-weighted

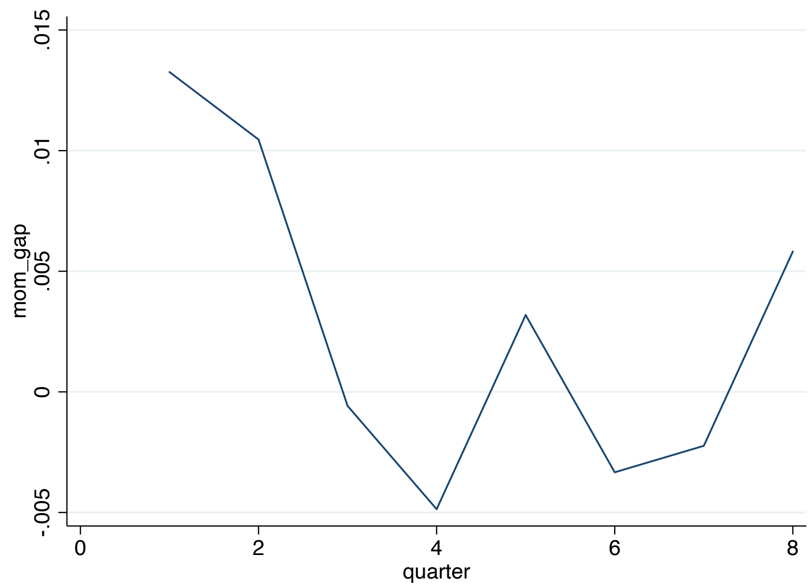

Fig. 18. Cross-sectional return predictability.

	Past horizon			
	1M	1M to 1Q	1Q to 1Y	1Y to 5Y
Buy at gain	19,566	20,844	32,329	46,469
Sell at gain	41,565	40,689	30,422	22,038
Propensity of selling at gain	*68.0%*	*66.1%*	*48.4%*	*32.2%*
Buy at loss	40,561	39,283	27,798	13,658
Sell at loss	18,248	19,124	29,391	37,775
Propensity of selling at loss	*31.0%*	*32.7%*	*51.4%*	*73.4%*
Disposition effect	*2.19*	*2.02*	*0.94*	*0.44*

Table 1. Measures of the disposition effect over different horizons.

We look at 10,000 years of monthly data simulated from the model. We adopt the baseline parameters that are specified in Section 3.3: $g_D = 0.05$, $\sigma_D = 0.25$, $r = 0.025$, $Q = 1$, $\gamma = 0.01$, $\mu = 0.3$, $\kappa = 0.05$, $\overline{\theta} = 2$, $\sigma_\theta = 5$, $\alpha = 0.2$, and $\delta = 2.77$. At each point in time in this simulated time series, we check whether the LSN investor has a positive or negative demand change. If she has a positive demand change, we count it as a "buy;" and if she has a negative demand change, we count it as a "sell." We then look at the price change of the risky asset over four different horizons: the price change over the past month ("1M"), the price change from one quarter ago to one month ago ("1M to 1Q"), the price change from one year ago to one quarter ago ("1Q to 1Y"), and the price change from five years ago to one year ago ("1Y to 5Y"). If the price change is positive, we count it as a "gain;" and if it is negative, we count it as a "loss." "*Propensity of selling at gain*" is calculated by dividing "Sell at gain" by the sum of "Sell at gain" and "Buy at gain." "*Propensity of selling at loss*" is calculated by dividing "Sell at loss" by the sum of "Sell at loss" and "Buy at loss." "*Disposition effect*" is then measured by the ratio of "*Propensity of selling at gain*" and "*Propensity of selling at loss.*"

	Past horizon			
	1M	1M to 1Q	1Q to 1Y	1Y to 5Y
Baseline: $\alpha = 0.2$, $\delta = 2.77$	2.19	2.02	0.94	0.44
Low α: $\alpha = 0.1$, $\delta = 2.77$	1.52	1.33	0.66	0.42
High α: $\alpha = 0.3$, $\delta = 2.77$	2.64	2.61	1.21	0.46
Low δ: $\alpha = 0.2$, $\delta = 1.39$	1.46	1.47	0.99	0.31
High δ: $\alpha = 0.2$, $\delta = 5.55$	3.83	2.09	0.73	0.63

Table 2. Measures of the disposition effect under different parametrizations of the LSN.

We look at 10,000 years of monthly data simulated from the model. The baseline parameters are: $g_D = 0.05$, $\sigma_D = 0.25$, $r = 0.025$, $Q = 1$, $\gamma = 0.01$, $\mu = 0.3$, $\kappa = 0.05$, $\overline{\theta} = 2$, $\sigma_\theta = 5$, $\alpha = 0.2$, and $\delta = 2.77$. Measures of the disposition effect are defined in Table 1. In each row, we vary one parameter from the baseline value and redo the entire simulation exercise to calculate the new measure of the disposition effect.

	(1)	(2)	(3)	(4)
	(Buy−Sell)/(Buy+Sell)		Buy−Sell	
Lagged stock return, 1M	–0.303***		–0.353*	
	(0.0198)		(0.191)	
Lagged stock return, 2M	–0.234***		–0.485***	
	(0.0135)		(0.0452)	
Lagged stock return, 3M	–0.129***		–0.255***	
	(0.0115)		(0.0502)	
Lagged stock return, 1Q		–0.172***		–0.317***
		(0.0118)		(0.0602)
Lagged stock return, 2Q		–0.0339***		–0.178***
		(0.00744)		(0.0271)
Lagged stock return, 3Q		0.0167**		–0.0928***
		(0.00765)		(0.0299)
Lagged stock return, 4Q		0.0161**		–0.0522
		(0.00737)		(0.0329)
Lagged stock return, 5Q		0.0120		–0.0582*
		(0.00755)		(0.0319)
Lagged stock return, 6Q		0.0231***		–0.0863***
		(0.00832)		(0.0307)
Lagged stock return, 7Q		0.0371***		–0.0285
		(0.00802)		(0.0338)
Lagged stock return, 8Q		0.0182**		–0.0115
		(0.00846)		(0.0289)
Lagged stock return, 9Q		0.0179**		–0.0487**
		(0.00818)		(0.0247)
Lagged stock return, 10Q		0.0281***		–0.00567
		(0.00875)		(0.0277)
Lagged stock return, 11Q		0.0226***		–0.00937
		(0.00824)		(0.0254)
Lagged stock return, 12Q		0.0186**		–0.0266
		(0.00797)		(0.0252)
Observations	577,488	577,488	577,488	577,488
R-squared	0.070	0.068	0.053	0.053

Robust standard errors in parentheses; *** p<0.01, ** p<0.05, * p<0.1

Table 3. Stock-level regressions results, the Odean data.

On each date, we aggregate all the transactions for each stock to get the total volume of buy and sell, denoted by Buy and Sell. Stock and date fixed effects are included. Standard errors are double-clustered by stock and date.

	(1)	(2)	(3)	(4)
	(Buy−Sell)/(Buy+Sell)		Buy−Sell	
Lagged stock return, 1W	–0.329***	–0.326***	–0.0104***	–0.0104***
	(0.0146)	(0.0146)	(0.00118)	(0.00119)
Lagged stock return, 2W	–0.0634***	–0.0600***	–0.00134**	–0.00134**
	(0.00969)	(0.00981)	(0.000620)	(0.000613)
Lagged stock return, 3W	0.00182	0.00516	–0.00152*	–0.00152*
	(0.00921)	(0.00928)	(0.000790)	(0.000800)
Lagged stock return, 4W	0.0333***	0.0381***	0.000957*	0.000987*
	(0.00885)	(0.00898)	(0.000531)	(0.000526)
Lagged stock return, 5W		0.0365***		0.000522
		(0.00857)		(0.000617)
Lagged stock return, 6W		0.0215**		0.000329
		(0.00844)		(0.000598)
Lagged stock return, 7W		0.0147*		–0.000465
		(0.00836)		(0.000548)
Lagged stock return, 8W		0.0288***		0.000175
		(0.00821)		(0.000484)
Lagged stock return, 9W		0.0282***		–0.000109
		(0.00777)		(0.000498)
Lagged stock return, 10W		0.0152**		–0.000144
		(0.00770)		(0.000405)
Lagged stock return, 11W		0.0277***		–0.000680
		(0.00793)		(0.000634)
Lagged stock return, 12W		0.0155**		–0.000912*
		(0.00718)		(0.000530)
Observations	2,754,207	2,754,207	2,754,207	2,754,207
R-squared	0.014	0.014	0.001	0.001

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Table 4. Stock-level regressions results, Chinese brokerage data.

Appendix A. Bayesian Inference

In this section, we analyze how LSN investors form beliefs about θ_t and $\overline{\omega}_t$ using Bayesian inference. By equations (8) to (10) of the main text and Theorem 12.7 of Lipster and Shiryaev (2001), we obtain

$$\begin{pmatrix} dm_{t,1} \\ dm_{t,2} \end{pmatrix} = \begin{pmatrix} \kappa\bar{\theta} - \kappa m_{t,1} \\ -(\alpha\delta + \delta)m_{t,2} \end{pmatrix} dt + \left[\begin{pmatrix} 0 \\ \delta \end{pmatrix} + \boldsymbol{\gamma}_t \begin{pmatrix} \sigma_P^{-1} \\ -\alpha \end{pmatrix} \right] [dP_t - (m_{t,1} - \sigma_P \alpha m_{t,2})dt]\sigma_P^{-1} \tag{A.1}$$

and

$$\begin{aligned} \frac{d}{dt}\boldsymbol{\gamma}_t = & -\begin{pmatrix} \kappa & 0 \\ 0 & (\alpha\delta+\delta) \end{pmatrix} \boldsymbol{\gamma}_t - \boldsymbol{\gamma}_t \begin{pmatrix} \kappa & 0 \\ 0 & (\alpha\delta+\delta) \end{pmatrix} + \begin{pmatrix} \sigma_\theta^2 & 0 \\ 0 & \delta^2 \end{pmatrix} \\ & - \left[\begin{pmatrix} 0 \\ \delta \end{pmatrix} + \boldsymbol{\gamma}_t \begin{pmatrix} \sigma_P^{-1} \\ -\alpha \end{pmatrix} \right] \left[\begin{pmatrix} 0 \\ \delta \end{pmatrix} + \boldsymbol{\gamma}_t \begin{pmatrix} \sigma_P^{-1} \\ -\alpha \end{pmatrix} \right]^T . \end{aligned} \tag{A.2}$$

To further simplify (A.1) and (A.2), we follow the literature on Kalman filtering and focus on the stationary solution of $\boldsymbol{\gamma}_t$, denoted by $\boldsymbol{\gamma}$. In this case, LSN investors' beliefs are fully specified by equations (12), (13), and (14) in the main text. Equation (A.2) implies that parameters γ_{11}, γ_{12}, and γ_{22} are the solution of

$$\begin{aligned} \begin{pmatrix} 2\kappa\gamma_{11} & (\kappa+\alpha\delta+\delta)\gamma_{12} \\ (\kappa+\alpha\delta+\delta)\gamma_{12} & 2(\alpha\delta+\delta)\gamma_{22} \end{pmatrix} = \begin{pmatrix} \sigma_\theta^2 & 0 \\ 0 & \delta^2 \end{pmatrix} \\ - \begin{pmatrix} (\sigma_P^{-1}\gamma_{11} - \alpha\gamma_{12})^2 & (\sigma_P^{-1}\gamma_{11} - \alpha\gamma_{12})(\delta + \sigma_P^{-1}\gamma_{12} - \alpha\gamma_{22}) \\ (\sigma_P^{-1}\gamma_{11} - \alpha\gamma_{12})(\delta + \sigma_P^{-1}\gamma_{12} - \alpha\gamma_{22}) & (\delta + \sigma_P^{-1}\gamma_{12} - \alpha\gamma_{22})^2 \end{pmatrix}, \end{aligned} \tag{A.3}$$

which is effectively three simultaneous equations. ■

Appendix B. Model Solution

In this section, we discuss the procedure that solves the model described in Section 3. Recall from equations (5) and (6) of the main text that both LSN investors and rational arbitrageurs have instantaneous mean-variance preferences subject to their budget constraints. Substituting (6) into (5) gives

$$N_t^i = \frac{\mathbb{E}_t^i[dP_t]/dt + D_t - rP_t}{\gamma\sigma_P^2}, \qquad i \in \{l, r\}. \tag{B.1}$$

We now solve the model. We start by conjecturing that, as stated in equation (15) of the main text, the equilibrium price of the risky asset is

$$P_t = A + B \cdot m_{t,1} + C \cdot m_{t,2} + \frac{D_t}{r}. \tag{B.2}$$

We solve for the three coefficients, A, B, and C, in three steps. The first step is to solve for LSN investors' share demand. Substituting (12) and (B.2) into (B.1), we obtain

$$N_t^l = \eta_0^l + \eta_1^l m_{t,1} + \eta_2^l m_{t,2}, \tag{B.3}$$

where

$$\eta_0^l = -\frac{rA}{\gamma\sigma_P^2}, \quad \eta_1^l = \frac{1 - rB}{\gamma\sigma_P^2}, \quad \eta_2^l = -\frac{\sigma_P\alpha + rC}{\gamma\sigma_P^2}. \tag{B.4}$$

The next step is to solve for the rational arbitrageurs' share demand. To do so, we take the differential form of (B.2)

$$dP_t = B \cdot dm_{t,1} + C \cdot dm_{t,2} + \frac{dD_t}{r}. \tag{B.5}$$

Substituting equations (12), (13) and (14) into (B.5) yields

$$dD_t = r\left(\begin{array}{l} (m_{t,1} - \sigma_P\alpha m_{t,2}) - \kappa B(\overline{\theta} - m_{t,1}) \\ +C(\alpha\delta + \delta)m_{t,2} \end{array}\right) dt + r(\sigma_P - \sigma_{m1}B - \sigma_{m2}C)d\tilde{\omega}_t^l. \tag{B.6}$$

Comparing (B.6) with (1) leads to

$$d\tilde{\omega}_t^l = d\omega_t^D + (l_0 + l_1 m_{t,1} + l_2 m_{t,2})dt \tag{B.7}$$

and

$$\sigma_P = \frac{\sigma_D}{r} + \sigma_{m1}B + \sigma_{m2}C, \tag{B.8}$$

where $l_0 \equiv \sigma_D^{-1}(g_D + r\kappa B\overline{\theta})$, $l_1 \equiv -\sigma_D^{-1}r(1+\kappa B)$, $l_2 \equiv \sigma_D^{-1}r[\sigma_P\alpha - C(\alpha\delta+\delta)]$, $\sigma_{m1} \equiv \gamma_{11}\sigma_P^{-1} - \gamma_{12}\alpha$, and $\sigma_{m2} \equiv \delta + \gamma_{12}\sigma_P^{-1} - \gamma_{22}\alpha$, as defined in Proposition 1.

Substituting (B.7) into (12) gives (16), which represents the rational arbitrageurs' beliefs about the price evolution. Moreover, substituting (B.7) into (13) and (14) gives (17) and (18), which represents the rational arbitrageurs' beliefs about $m_{t,1}$ and $m_{t,2}$. We combine (B.1), (16), and (B.2) for rational arbitrageurs and obtain

$$N_t^r = \eta_0^r + \eta_1^r m_{t,1} + \eta_2^r m_{t,2}, \tag{B.9}$$

where

$$\begin{aligned} \eta_0^r &= \frac{\sigma_D^{-1}\sigma_P(g_D + r\kappa B\overline{\theta}) - rA}{\gamma\sigma_P^2}, \quad \eta_1^r = \frac{\sigma_D^{-1}\sigma_P[(\sigma_D\sigma_P^{-1} - r) - r\kappa B] - rB}{\gamma\sigma_P^2}, \\ \eta_2^r &= -\frac{\sigma_D^{-1}\sigma_P[(\sigma_D\sigma_P^{-1} - r)\sigma_P\alpha + rC(\alpha\delta + \delta)] + rC}{\gamma\sigma_P^2}. \end{aligned} \tag{B.10}$$

The final step is to substitute the share demands, (B.3) and (B.9), into the market clearing condition in (7). We then obtain

$$\begin{aligned} \mu\eta_0^r + (1-\mu)\eta_0^l &= Q, \\ \mu\eta_1^r + (1-\mu)\eta_1^l &= 0, \\ \mu\eta_2^r + (1-\mu)\eta_2^l &= 0. \end{aligned} \tag{B.11}$$

Substituting (B.4), (B.8), and (B.10) into (B.11) gives three simultaneous equations for three unknowns, A, B, and C. We solve these simultaneous equations using numerical methods. Once coefficients A, B, and C are solved, σ_P is then given by (B.8). ■

Appendix C. Model Extension

In this section, we briefly describe and then solve a more generalized model, one that features three types of investors: LSN investors with $\alpha > 0$, LSN investors with $\alpha = 0$, and rational arbitrageurs. We refer to LSN investors with $\alpha = 0$ as "extrapolators," because their beliefs about the future price change depend positively on past price changes. We then refer to LSN investors with $\alpha > 0$ simply as "LSN investors."

C.1. Model setup

Asset space. As in the baseline model, we consider two assets: a riskless asset with a constant interest rate r, and a risky asset. The risky asset has a fixed per-capita supply of Q, and its dividend payment evolves according to equation (1) in the main text. The price of the risky asset P_t is endogenously determined in equilibrium.

Investor beliefs. Rational arbitrageurs make up a fraction μ_r of the total population; extrapolators make up a fraction μ_e of the total popultion; and LSN investors make up the remaining fraction of $1 - \mu_r - \mu_e$.

LSN investors' perceived price processes are specified by equations (2) to (4). Extrapolators represent a special case of LSN investors. They believe

$$dP_t = \theta_t^e dt + \sigma_P d\tilde{\omega}_t^{P,e}, \tag{C.1}$$

where

$$d\theta_t^e = \kappa^e(\overline{\theta}^e - \theta_t^e)dt + \sigma_\theta^e d\tilde{\omega}_t^{\theta,e}, \tag{C.2}$$

and both $d\tilde{\omega}_t^{P,e}$ and $d\tilde{\omega}_t^{\theta,e}$ are perceived by extrapolators to be i.i.d. shocks that are independent of each other. Rational arbitrageurs hold fully rational beliefs: they understand the dividend process in equation (1); they observe parameters μ_r and μ_e and hence know the population fractions of LSN investors and the extrapolators; and they are aware of the belief structure of LSN investors and the belief structure of the extrapolators. Given this information set, rational arbitrageurs form correct beliefs about the evolution of the risky asset price.

Investor preferences. We assume that all three types of investors have instantaneous mean-variance preferences specified by

$$\max_{N_t^i} \left(\mathbb{E}_t^i[dW_t^i] - \frac{\gamma}{2} \mathbb{V}\mathrm{ar}_t^i[dW_t^i] \right), \tag{C.3}$$

subject to the budget constraint on their wealth W_t^i

$$dW_t^i = rW_t^i dt - rN_t^i P_t dt + N_t^i dP_t + N_t^i D_t dt, \tag{C.4}$$

where N_t^i represents the per-capita share demand on the risky asset from investor i. Here, $i \in \{l, e, r\}$, where superscripts "l," "e," and "r" represent LSN investors, extrapolators, and rational arbitrageurs, respectively.

Market clearing. The share demands from LSN investors, extrapolators, and rational arbitrageurs satisfy the following market clearing condition

$$\mu_r N_t^r + \mu_e N_t^e + (1 - \mu_r - \mu_e) N_t^l = Q \tag{C.5}$$

at each point in time t.

C.2. Model solution

As in the baseline model, applying Kalman filters to equations (2) to (4) yields equations (12) to (14), which specifies the way in which LSN investors update their beliefs based on past prices. For the extrapolators, denote the conditional mean and variance of θ_t^e as

$$S_t = \mathbb{E}^e[\theta_t^e | \mathcal{F}_t^P], \qquad \zeta_t = \mathbb{E}^e[(\theta_t^e - S_t)^2 | \mathcal{F}_t^P]. \tag{C.6}$$

Then we apply Kalman filters (Theorem 12.7 from Lipster and Shiryaev, 2001) to (C.1) and (C.2) and obtain

$$dP_t = S_t dt + \sigma_P d\tilde{\omega}_t^e, \tag{C.7}$$

and

$$dS_t = \kappa^e(\overline{\theta}^e - S_t)dt + (\zeta \sigma_P^{-1}) d\tilde{\omega}_t^e, \tag{C.8}$$

where $d\tilde{\omega}_t^e$ is a Brownian shock perceived by extrapolators, and

$$\zeta = -\kappa^e \sigma_P^2 + \sqrt{(\kappa^e \sigma_P^2)^2 + (\sigma_\theta^e)^2 \sigma_P^2} \tag{C.9}$$

is the stationary solution for ζ_t in (C.8).

To solve the model, we first substitute (C.4) into (C.3) and obtain

$$N_t^i = \frac{\mathbb{E}_t^i[dP_t]/dt + D_t - rP_t}{\gamma \sigma_P^2}, \qquad i \in \{l, e, r\}. \tag{C.10}$$

We conjecture that the equilibrium price of the risky asset is

$$P_t = A + B_1 \cdot m_{t,1} + B_2 \cdot m_{t,2} + C \cdot S_t + \frac{D_t}{r}. \tag{C.11}$$

Substituting (12) and (C.11) into (C.10) for LSN investors, we obtain

$$N_t^l = \eta_0^l + \eta_1^l m_{t,1} + \eta_2^l m_{t,2} + \eta_3^l S_t, \tag{C.12}$$

where

$$\eta_0^l = -\frac{rA}{\gamma\sigma_P^2}, \quad \eta_1^l = \frac{1-rB_1}{\gamma\sigma_P^2}, \quad \eta_2^l = -\frac{\sigma_P\alpha + rB_2}{\gamma\sigma_P^2}, \quad \eta_3^l = -\frac{rC}{\gamma\sigma_P^2}. \tag{C.13}$$

We then substitute (C.7) and (C.11) into (C.10) for extrapolators and obtain

$$N_t^e = \eta_0^e + \eta_1^e m_{t,1} + \eta_2^e m_{t,2} + \eta_3^e S_t, \tag{C.14}$$

where

$$\eta_0^e = -\frac{rA}{\gamma\sigma_P^2}, \quad \eta_1^e = -\frac{rB_1}{\gamma\sigma_P^2}, \quad \eta_2^e = -\frac{rB_2}{\gamma\sigma_P^2}, \quad \eta_3^e = \frac{1-rC}{\gamma\sigma_P^2}. \tag{C.15}$$

Finally, we examine the share demand of the rational arbitrageurs. We take the differential form of (C.11)

$$dP_t = B_1 \cdot dm_{t,1} + B_2 \cdot dm_{t,2} + C \cdot dS_t + \frac{dD_t}{r}. \tag{C.16}$$

Note that $dS_t = \kappa^e(\overline{\theta}^e - S_t)dt + (\zeta\sigma_P^{-2})(dP_t - S_t dt)$. Substituting this equation and equations (12) to (14) into (C.11), we get

$$\begin{aligned} dD_t = \;& r \begin{pmatrix} [1 - C\cdot(\zeta\sigma_P^{-2})](m_{t,1} - \sigma_P\alpha m_{t,2}) - \kappa B_1(\overline{\theta} - m_{t,1}) \\ +B_2(\alpha\delta+\delta)m_{t,2} - C\kappa^e\overline{\theta}^e + C[\kappa^e + (\zeta\sigma_P^{-2})]S_t \end{pmatrix} dt \\ & +r\left([1 - C\cdot(\zeta\sigma_P^{-2})]\sigma_P - B_1\sigma_{m1} - B_2\sigma_{m2}\right) d\tilde{\omega}_t^l. \end{aligned} \tag{C.17}$$

Comparing (C.17) with (1) gives

$$d\tilde{\omega}_t^l = d\omega_t^D + \sigma_D^{-1} r \begin{pmatrix} r^{-1}g_D - [1 - C\cdot(\zeta\sigma_P^{-2})](m_{t,1} - \sigma_P\alpha m_{t,2}) \\ +\kappa B_1(\overline{\theta} - m_{t,1}) - B_2(\alpha\delta+\delta)m_{t,2} \\ +C\kappa^e\overline{\theta}^e - C[\kappa^e + (\zeta\sigma_P^{-2})]S_t \end{pmatrix} dt \tag{C.18}$$

and

$$\sigma_P = \frac{1}{1 - C\cdot(\zeta\sigma_P^{-2})}\left(\frac{\sigma_D}{r} + B_1\sigma_{m1} + B_2\sigma_{m2}\right). \tag{C.19}$$

Substituting (C.18) into (12), we have

$$dP_t = \sigma_P\sigma_D^{-1} r \begin{pmatrix} r^{-1}g_D - [1 - C\cdot(\zeta\sigma_P^{-2})](m_{t,1} - \sigma_P\alpha m_{t,2}) \\ +\kappa B_1(\overline{\theta} - m_{t,1}) - B_2(\alpha\delta+\delta)m_{t,2} \\ +C\kappa^e\overline{\theta}^e - C[\kappa^e + (\zeta\sigma_P^{-2})]S_t \\ +r^{-1}\sigma_D\sigma_P^{-1}(m_{t,1} - \sigma_P\alpha m_{t,2}) \end{pmatrix} dt + \sigma_P d\omega_t^D. \tag{C.20}$$

Then, further substituting (C.20) and (C.11) into (C.10) for the rational arbitrageurs, we get

$$N_t^r = \eta_0^r + \eta_1^r m_{t,1} + \eta_2^r m_{t,2} + \eta_3^r S_t, \tag{C.21}$$

where

$$\begin{aligned}
\eta_0^r &= \frac{\sigma_P \sigma_D^{-1}(g_D + r\kappa B_1 \overline{\theta} + r\kappa^e C \overline{\theta}^e) - rA}{\gamma \sigma_P^2}, \\
\eta_1^r &= \frac{\sigma_P \sigma_D^{-1}[\sigma_D \sigma_P^{-1} - r(1 - C \cdot (\zeta \sigma_P^{-2})) - r\kappa B_1] - rB_1}{\gamma \sigma_P^2}, \\
\eta_2^r &= -\frac{\sigma_P \sigma_D^{-1}[(\sigma_D \sigma_P^{-1} - r(1 - C \cdot (\zeta \sigma_P^{-2})))\sigma_P \alpha + rB_2(\alpha\delta + \delta)] + rB_2}{\gamma \sigma_P^2}, \\
\eta_3^r &= -\frac{\sigma_P \sigma_D^{-1} rC[\kappa^e + (\zeta \sigma_P^{-2})] + rC}{\gamma \sigma_P^2}.
\end{aligned} \tag{C.22}$$

The final step is to substitute the share demands, (C.12), (C.14), and (C.21), into the market clearing condition in (C.5). We obtain

$$\begin{aligned}
\mu_r \eta_0^r + \mu_e \eta_0^e + (1 - \mu_r - \mu_e)\eta_0^l &= Q, \\
\mu_r \eta_1^r + \mu_e \eta_1^e + (1 - \mu_r - \mu_e)\eta_1^l &= 0, \\
\mu_r \eta_2^r + \mu_e \eta_2^e + (1 - \mu_r - \mu_e)\eta_2^l &= 0, \\
\mu_r \eta_3^r + \mu_e \eta_3^e + (1 - \mu_r - \mu_e)\eta_3^l &= 0.
\end{aligned} \tag{C.23}$$

Substituting (C.13), (C.15), (C.19), and (C.22) into (C.23) gives four simultaneous equations for four unknowns, A, B_1, B_2, and C. We solve these simultaneous equations using numerical methods. Once coefficients A, B_1, B_2, and C are solved, σ_P is then given by (C.19). ■

Appendix D. Alternative Specification of LSN Beliefs

The baseline model described in Section 3.1 applies the LSN to the price process; beliefs of LSN investors are specified by equations (2) to (4) in the main text. In this section, we consider an alternative specification in which the LSN is applied to the dividend process. As before, this modified model contains two assets: a risk-free asset and a risky asset. The risk-free asset pays a constant interest rate of r. The stock market has a fixed per-capita supply of Q, and its dividend payment evolves according to

$$dD_t = g_D dt + \sigma_D d\omega_t^D. \tag{D.1}$$

LSN investors are now assumed to perceive the following dividend process

$$\begin{aligned} dD_t = \theta_t dt + \sigma_D d\tilde{\omega}_t^D, \qquad d\theta_t = \kappa(\bar{\theta} - \theta_t)dt + \sigma_\theta d\tilde{\omega}_t^\theta, \\ d\tilde{\omega}_t^D = d\tilde{\omega}_t - \alpha \left(\delta \int_{-\infty}^{t} e^{-\delta(t-s)} d\tilde{\omega}_s^D \right) dt. \end{aligned} \tag{D.2}$$

In words, LSN investors perceive future dividend changes as coming from two components: a persistent yet time-varying quality component, and a transitory noise component that exhibits a negative serial autocorrelation.

An equivalent specification of (D.2) is

$$\begin{aligned} dD_t &= (\theta_t - \sigma_D \alpha \overline{\omega}_t)dt + \sigma_D d\tilde{\omega}, \qquad d\theta_t = \kappa(\bar{\theta} - \theta_t)dt + \sigma_\theta d\tilde{\omega}_t^\theta, \\ d\overline{\omega}_t &= -(\alpha\delta + \delta)\overline{\omega}_t dt + \delta d\tilde{\omega}_t, \end{aligned} \tag{D.3}$$

where $\overline{\omega}_t \equiv \int_{-\infty}^{t} \delta e^{-\delta(t-s)} d\tilde{\omega}_s^D$ and $\mathbb{E}_t^l[d\tilde{\omega}_t \cdot d\tilde{\omega}_t^\theta] = 0$.

LSN investors do not observe θ_t and $\overline{\omega}_t$; they use Bayesian inference to estimate both quantities and then use these estimated quantities to guide trading decisions. Their information set at time t, $\mathcal{F}_t^D$, is defined using past dividends $\{D_s, s \le t\}$—that is, LSN investors update their beliefs about θ_t and $\overline{\omega}_t$ using past dividends as informative signals. The conditional means and variances of $\boldsymbol{\theta}_t \equiv (\theta_t, \overline{\omega}_t)$ are defined by

$$\begin{aligned} \boldsymbol{m}_t &= (m_{t,1}, m_{t,2}) \equiv \mathbb{E}^l[(\theta_t, \overline{\omega}_t) | \mathcal{F}_t^D], \\ \boldsymbol{\gamma}_t &= \begin{pmatrix} \gamma_{t,11} & \gamma_{t,12} \\ \gamma_{t,21} & \gamma_{t,22} \end{pmatrix} \equiv \mathbb{E}^l[(\boldsymbol{\theta}_t - \boldsymbol{m}_t)^T (\boldsymbol{\theta}_t - \boldsymbol{m}_t) | \mathcal{F}_t^D]. \end{aligned} \tag{D.4}$$

We then apply Kalman filtering and obtain

$$dD_t = (m_{t,1} - \sigma_D \alpha m_{t,2})dt + \sigma_D d\tilde{\omega}_t^l \tag{D.5}$$

and

$$dm_{t,1} = \kappa(\overline{\theta} - m_{t,1})dt + \underbrace{(\gamma_{11}\sigma_D^{-1} - \gamma_{12}\alpha)}_{\sigma_{m1}} d\tilde{\omega}_t^l, \tag{D.6}$$

$$dm_{t,2} = -(\alpha\delta + \delta)m_{t,2}dt + \underbrace{(\delta + \gamma_{12}\sigma_D^{-1} - \gamma_{22}\alpha)}_{\sigma_{m2}} d\tilde{\omega}_t^l, \tag{D.7}$$

where $d\tilde{\omega}_t^l$ is a Brownian shock perceived by LSN investors, and γ_{11}, γ_{12}, and γ_{22} are the stationary solutions for $\gamma_{t,11}$, $\gamma_{t,12}$, and $\gamma_{t,22}$, respectively. In these equations, $m_{t,1}$ and $m_{t,2}$ represent the inferred quantities of θ_t and $\overline{\omega}_t$. Moreover, γ_{11}, γ_{12}, and γ_{22} are the solution of

$$\begin{aligned}\begin{pmatrix} 2\kappa\gamma_{11} & (\kappa+\alpha\delta+\delta)\gamma_{12} \\ (\kappa+\alpha\delta+\delta)\gamma_{12} & 2(\alpha\delta+\delta)\gamma_{22} \end{pmatrix} &= \begin{pmatrix} \sigma_\theta^2 & 0 \\ 0 & \delta^2 \end{pmatrix} \\ -\begin{pmatrix} (\sigma_D^{-1}\gamma_{11} - \alpha\gamma_{12})^2 & (\sigma_D^{-1}\gamma_{11} - \alpha\gamma_{12})(\delta + \sigma_D^{-1}\gamma_{12} - \alpha\gamma_{22}) \\ (\sigma_D^{-1}\gamma_{11} - \alpha\gamma_{12})(\delta + \sigma_D^{-1}\gamma_{12} - \alpha\gamma_{22}) & (\delta + \sigma_D^{-1}\gamma_{12} - \alpha\gamma_{22})^2 \end{pmatrix}&.\end{aligned} \tag{D.8}$$

As in the baseline model, we assume there are two types of investors: LSN investors and rational arbitrageurs. Rational arbitrageurs make up μ fraction of the total population; LSN investors make up the remaining $1-\mu$ fraction. Both LSN investors and rational arbitrageurs maximize instantaneous mean-variance preferences, specified by

$$\max_{N_t^i} \left(\mathbb{E}_t^i[dW_t^i] - \frac{\gamma}{2}\mathbb{V}\mathrm{ar}_t^i[dW_t^i]\right), \tag{D.9}$$

subject to the budget constraint on their wealth W_t^i

$$dW_t^i = rW_t^i dt - rN_t^i P_t dt + N_t^i dP_t + N_t^i D_t dt, \tag{D.10}$$

where N_t^i represents the per-capita share demand on the risky asset from investor i and $i \in \{l, r\}$. Substituting (D.10) into (D.9) gives

$$N_t^i = \frac{\mathbb{E}_t^i[dP_t]/dt + D_t - rP_t}{\gamma\sigma_P^2}. \tag{D.11}$$

The conjectured equilibrium price of the stock market is

$$P_t = A + B \cdot m_{t,1} + C \cdot m_{t,2} + \frac{D_t}{r}. \tag{D.12}$$

As before, we solve for the three unknowns, A, B, and C, in three steps. The first step is to solve for LSN investors' share demand. LSN investors differentiate both sides of (D.12) and obtain

$$dP_t = B \cdot dm_{t,1} + C \cdot dm_{t,2} + \frac{dD_t}{r}. \tag{D.13}$$

They then substitute equations (D.5) and (D.6) to the right hand side of (D.12) and obtain

$$\begin{aligned} dP_t &= B\kappa(\overline{\theta} - m_{t,1})dt + B\sigma_{m1}d\omega_t^l - C(\alpha\delta + \delta)m_{t,2}dt + C\sigma_{m2}d\omega_t^l \\ &+ r^{-1}(m_{t,1} - \sigma_D\alpha m_{t,2})dt + r^{-1}\sigma_D d\omega_t^l. \end{aligned} \tag{D.14}$$

LSN investors' expected price change is therefore

$$\mathbb{E}_t^l[dP_t]/dt = B\kappa(\overline{\theta} - m_{t,1}) - C(\alpha\delta + \delta)m_{t,2} + r^{-1}(m_{t,1} - \sigma_D\alpha m_{t,2}). \tag{D.15}$$

Substituting (D.15) and (D.12) into (D.11) gives

$$\begin{aligned} N_t^l &= \frac{B\kappa(\overline{\theta} - m_{t,1}) - C(\alpha\delta + \delta)m_{t,2} + r^{-1}(m_{t,1} - \sigma_D\alpha m_{t,2}) - rA - rB \cdot m_{t,1} - rC \cdot m_{t,2}}{\gamma\sigma_P^2} \\ &\equiv \eta_0^l + \eta_1^l m_{t,1} + \eta_2^l m_{t,2}, \end{aligned} \tag{D.16}$$

where

$$\eta_0^l = \frac{B\kappa\overline{\theta} - rA}{\gamma\sigma_P^2}, \quad \eta_1^l = \frac{r^{-1} - \kappa B - rB}{\gamma\sigma_P^2}, \quad \eta_2^l = -\frac{C(\alpha\delta + \delta) + r^{-1}\sigma_D\alpha + rC}{\gamma\sigma_P^2}. \tag{D.17}$$

The next step is to solve for rational arbitrageurs' share demand. We compare (D.5) with (D.1) and obtain

$$d\omega_t^l = d\omega_t^D + \sigma_D^{-1}(g_D - m_{t,1} + \sigma_D\alpha m_{t,2})dt. \tag{D.18}$$

Substituting (D.18) into (D.14) gives

$$dP_t = \begin{pmatrix} B\kappa(\bar{\theta} - m_{t,1}) - C(\alpha\delta + \delta)m_{t,2} \\ +r^{-1}(m_{t,1} - \sigma_D\alpha m_{t,2}) \\ +\sigma_D^{-1}\sigma_P(g_D - m_{t,1} + \sigma_D\alpha m_{t,2}) \end{pmatrix} dt + \sigma_P d\omega_t^D \tag{D.19}$$

and

$$\sigma_P = \frac{\sigma_D}{r} + \sigma_{m1}B + \sigma_{m2}C. \tag{D.20}$$

Equations (D.19) and (D.20) represent rational arbitrageurs' beliefs about price evolution. We then combine (D.19), (D.11), and (D.12) to obtain

$$N_t^r \equiv \eta_0^r + \eta_1^r m_{t,1} + \eta_2^r m_{t,2}, \tag{D.21}$$

where

$$
\begin{aligned}
\eta_0^r &= \frac{B\kappa\overline{\theta} - rA + \sigma_D^{-1}\sigma_P g_D}{\gamma\sigma_P^2}, \quad \eta_1^r = \frac{r^{-1} - \kappa B - rB - \sigma_D^{-1}\sigma_P}{\gamma\sigma_P^2}, \\
\eta_2^r &= -\frac{C(\alpha\delta + \delta) + r^{-1}\sigma_D\alpha + rC - \sigma_P\alpha}{\gamma\sigma_P^2}.
\end{aligned} \tag{D.22}
$$

The final step is to substitute the share demands (D.16) and (D.21) into the market clearing condition $\mu N_t^r + (1-\mu)N_t^l = Q$. We arrive at three equations

$$
\begin{aligned}
&\mu\eta_0^r + (1-\mu)\eta_0^l = Q, \\
&\mu\eta_1^r + (1-\mu)\eta_1^l = 0, \\
&\mu\eta_2^r + (1-\mu)\eta_2^l = 0.
\end{aligned} \tag{D.23}
$$

Substituting (D.17), (D.20), and (D.22) into (D.23) gives three simultaneous equations for three unknowns, A, B, and C. We solve these equations using numerical methods. ■