

NBER WORKING PAPER SERIES

THE LAW OF SMALL NUMBERS IN FINANCIAL MARKETS:  
THEORY AND EVIDENCE

Lawrence J. Jin  
Cameron Peng

Working Paper 32519  
<http://www.nber.org/papers/w32519>

NATIONAL BUREAU OF ECONOMIC RESEARCH  
1050 Massachusetts Avenue  
Cambridge, MA 02138  
May 2024

We thank Nicholas Barberis, Francesca Bastianello, James Choi, Paul Fontanier, Cary Frydman, David Hirshleifer, 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. Allen Hu and Andrew Wang provided excellent research assistance. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2024 by Lawrence J. Jin and Cameron Peng. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Law of Small Numbers in Financial Markets: Theory and Evidence  
Lawrence J. Jin and Cameron Peng  
NBER Working Paper No. 32519  
May 2024  
JEL No. G0,G02,G11,G12

### **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 our model, a belief in the LSN induces investors to expect short-term price trends to revert and long-term price trends to persist. As a result, asset prices exhibit short-term momentum and long-term reversals. The model can reconcile the coexistence of the disposition effect and return extrapolation. In addition, it makes new predictions about investor behavior, including return patterns before purchases and sales, a weakened disposition effect for long-term holdings, doubling down in buying, a positive correlation between doubling down and the disposition effect, and heterogeneous selling propensities to past returns. By testing these predictions using account-level transaction data, we show that the LSN provides a parsimonious way of understanding a variety of puzzles about investor behavior.

Lawrence J. Jin  
Johnson College of Business  
Cornell University  
Ithaca, NY 14850  
and NBER  
lawrence.jin@cornell.edu

Cameron Peng  
London School of Economics  
Houghton Street  
London  
WC2A 2AE  
UK  
c.peng9@lse.ac.uk

A data appendix is available at <http://www.nber.org/data-appendix/w32519>

# 1. Introduction

When making forecasts about a random outcome, people often succumb to the “gambler’s fallacy.” For instance, after seeing a streak of heads in a series of fair coin tosses, individuals tend to expect a tail on the next toss, despite the constant 50% objective probability (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).<sup>1</sup> More broadly, alongside heuristics like overreaction and base-rate neglect, the LSN illustrates the tendency to draw conclusions too quickly by relying on too little data.

An immediate consequence of the LSN is that people behave as contrarians: when predicting the outcome of a random sequence, they tend to expect an immediate reversal in trends. However, it has also been suggested that the LSN can simultaneously lead to a belief in a “hot hand,” 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 the data-generating process, the LSN results in the gambler’s fallacy; but when they do not, they rely too much on the few data points they have observed to make inferences, leading to a belief in a “hot hand” 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 play an important role, because investors constantly observe past trends in prices and fundamentals, and 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 significantly impact their trading behavior and asset prices. While existing papers have modeled the LSN in general economic settings (e.g., Rabin, 2002; Rabin and Vayanos, 2010), our paper applies this belief structure in a financial setting

---

<sup>1</sup>The same idea has also been labelled “local representativeness” (Bar-Hillel and Wagenaar, 1991).

with equilibrium asset prices. We derive new testable predictions about trading behavior and asset prices. Importantly, we also empirically test these predictions using 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 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. When making portfolio choices, rational arbitrageurs correctly understand the dividend process and have perfect foresight of the price process. However, LSN investors do not directly observe the true price process. Instead, they make forecasts about future price changes based on their information set. We assume that LSN investors use an incorrect yet intuitive mental model to make inferences about the price process: they believe that the risky asset’s price change is determined by a “quality” term—time-varying and unobservable—and a noise term, and they make inferences about the asset’s quality by observing its past prices. In this 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. This assumption intuitively captures the gambler’s fallacy, where LSN investors expect short-term deviations from the mean to quickly revert. Compared to the earlier case without this LSN assumption, LSN investors’ belief structure changes in two significant ways. First, unlike simple return extrapolation where beliefs about future price changes depend *positively* on *all* past price changes, LSN investors’ beliefs now 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, while LSN investors’ beliefs continue to depend positively on price changes in distant periods, this positive dependence is now amplified. After seeing a long streak of positive returns—and perceiving negatively auto-correlated noise terms—LSN investors become more convinced of the asset’s high quality. Thus, the same force that generates short-term contrarian beliefs also reinforces extrapolation over long-term price trends.

The above belief structure, under mean-variance preferences, 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 long-term return extrapolators, buying when asset prices have gone up consistently over a long period of time. In this way, the model can reconcile a longstanding tension between the disposition effect, namely investors selling an asset with good past performance, and extrapolative demand, namely investors buying an asset with good past performance. Moreover, these trading behaviors further feed back into asset price dynamics: in the short run, the disposition effect induces short-term momentum; in the long 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, there is a direct mapping between past price changes and expectations of future price changes. LSN investors form incorrect beliefs about future price changes by looking at past price changes, and they then use these beliefs to decide on their share demand of the risky asset. We consider this thought process as psychologically simple and realistic. For robustness, we also consider an alternative specification in which LSN investors form incorrect beliefs about future *dividend* changes by looking at past dividend changes. In this specification, before making investment decisions, LSN investors take an extra step of deriving beliefs about prices from beliefs about dividends. Due to this extra step, we view such a thought process as less realistic. Nonetheless, this alternative specification similarly produces a 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 predictions about investor behavior using data from a U.S. brokerage firm (Odean, 1998; Barber and Odean, 2000).<sup>2</sup> First, the model makes predictions about the return patterns before purchases and sales. In the model, LSN investors tend to purchase assets with a price that had gone up for many periods but recently went down; they tend to sell assets with the opposite return patterns. We confirm these predictions in the brokerage data: across all buys, the median monthly return remains positive from 36 months prior to the purchase up until around 5 months prior, but experiences a sharp decline in the more recent months. Conversely, the stocks investors tend to sell experience a

---

<sup>2</sup>In Appendix J, we also examine the model's prediction on asset prices regarding the sources of momentum and long-term reversals, although this is not the paper's focus.

dramatic increase in price during the most recent month.

Second, the model predicts a weakened disposition effect for positions held over a long period. Specifically, because contrarian beliefs apply only to recent returns, the model predicts a stronger disposition effect for positions with a short holding period. For positions with a longer holding period, return extrapolation begins to counteract the disposition effect. This prediction is supported by the brokerage data. Among stocks bought within the last month, the probability of selling a winner is almost twice that of selling a loser. In contrast, for positions held for more than two years, the propensities for selling winners and losers are virtually the same.

Third, 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 recently gone down in value rather than shares that have recently gone up in value. We confirm this prediction in the brokerage data: on average, investors are 50% more likely to buy loser stocks than winner stocks, a result that is consistent with [Odean \(1998\)](#).

Fourth, the model predicts not only the coexistence of the disposition effect and doubling down at the *aggregate* level but also a strong association between these two phenomena at the *individual* level. Specifically, investors who are more likely to double down in buying are also expected to exhibit a stronger disposition effect in selling, as the LSN beliefs underlie both trading behaviors. To test this prediction, we first 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, supporting the idea that the LSN drives both buying and selling decisions.

Fifth and finally, the model predicts that the relationship between an individual’s trading propensity and past returns depends on their LSN beliefs. We extend the model to include not only LSN investors, who believe noise is negatively auto-correlated, but also pure extrapolators, who believe noise is i.i.d. This extended model predicts that LSN investors’ selling propensity increases with recent returns, while their buying propensity decreases with recent returns. Conversely, pure extrapolators exhibit the opposite trading patterns. Our empirical findings support these predictions and highlight the importance of investor heterogeneity in studying trading behavior. [Ben-David and Hirshleifer \(2012\)](#)’s “V-shaped” pattern in investors’ selling propensity is

not observed among LSN investors or pure extrapolators individually. Instead, each of the two investor types is responsible for one arm of the V, collectively driving the observed selling propensity. Together, the five empirical tests described above not only provide evidence supporting the model but also document new facts not explained by existing theories of investor behavior. In this regard, the LSN offers a parsimonious way of understanding various puzzles in investor behavior.

One differentiating feature of our model is that it makes realistic predictions about both trading behavior and asset prices, whereas earlier models typically focus on one of the two.<sup>3</sup> To do so, the model builds on previous work that studies belief formation under the LSN in partial equilibrium settings (Rabin, 2002; Rabin and Vayanos, 2010).<sup>4</sup> In particular, it incorporates this belief structure into a general equilibrium setting by specifying investors’ preferences and their portfolio problems, introducing other market participants such as rational arbitrageurs, and analyzing equilibrium asset prices. In this regard, the closest to our model is Tegui (2017), who also develops an equilibrium model that features LSN investors and rational traders. Our paper is different in two important aspects. First, we explore novel predictions that are not considered by Tegui (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 empirically test our model’s predictions using account-level transaction data, and find strong consistency between the data and the model’s predictions.

Our paper has important implications for the study of investor behavior. First, we show that the LSN can bridge the gap between two robust phenomena in investor behavior: the disposition effect and return extrapolation. Because these phenomena entail distinct attitudes towards past trends, their underlying mechanisms are often perceived as different, and generating both simultaneously typically requires invoking multiple psychological forces (Barber and Odean, 2013; Liao, Peng, and Zhu, 2022). However, we show that this does not need to be the case: the LSN alone can simultaneously generate both phenomena without the need for additional forces.

Second, 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

---

<sup>3</sup>For example, earlier representative-agent models such as Campbell and Cochrane (1999), Bansal and Yaron (2004), and Barberis, Huang, and Santos (2001) focus on studying asset prices, whereas earlier models of investor behavior such as Barberis and Xiong (2009, 2012) and Ingersoll and Jin (2013) focus on studying trading patterns.

<sup>4</sup>We follow Rabin and Vayanos (2010) to model LSN beliefs using Bayesian learning.

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 LSN. Consistent with our model, we find that investors are particularly likely to sell stocks 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.

Third, we show that, as the flip side of the disposition effect, there also exists doubling down in buying behavior. Importantly, these two phenomena are tightly linked to each other: investors who tend to double down in buying also display a stronger disposition effect in selling. This link enriches our understanding of investor trading by connecting the buying side with the selling side. Moreover, it raises the bar for explanations of the disposition effect: a unifying explanation should be able to simultaneously explain the disposition effect in selling and doubling down in buying.

Lastly, our results have implications for the well-documented “V-shaped” selling 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: LSN investors mainly sell past winners, while extrapolators mainly sell past losers. Our results call for further understanding of the V-shape through the lens of heterogeneity in investor beliefs.

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 empirically tests the model’s predictions on trading behavior and Section 5 concludes. Additional details and analyses 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. 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 the 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 each other, but it has become clear that they can, in fact, be generated by the *same* psychological 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 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 a belief in a “hot hand.” 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

$$dD_t = g_D dt + \sigma_D d\omega_t^D, \quad (1)$$

where  $\omega_t^D$  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  $\mu$  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. As a result, to make investment decisions, they 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 process. They perceive the price process as

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

where  $\theta_t$  represents the perceived quality of the asset and evolves according to

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

and  $d\tilde{\omega}_t^P$  represents an innovation component. In equation (3),  $\kappa > 0$  is a persistence parameter,  $\bar{\theta}$

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  $\kappa$  measures how quickly the asset's perceived quality  $\theta_t$  changes over time: when  $\kappa$  increases, asset quality is expected to revert back to its long-run mean more quickly. Parameter  $\sigma_\theta$  captures the size of perceived shocks to asset quality: when  $\sigma_\theta$  increases, asset quality is more subject to random shocks and hence exhibits higher variability.

Note 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. Also note that, such a mental model is 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 basis for LSN beliefs to operate: 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

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

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  $\alpha$  and  $\delta$  measure two different aspects of the LSN. Parameter  $\alpha$  measures the *strength* of the gambler's fallacy: a larger  $\alpha$  means a stronger belief in trend reversion. Parameter  $\delta$  measures the relative *weight* put on recent versus distant past realizations of  $d\tilde{\omega}_s^P$ : a larger  $\delta$  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 rather than future price changes; this is to follow a large literature that directly specifies investors' incorrect beliefs about asset fundamentals (e.g., [Barberis, Shleifer, and Vishny, 1998](#)). We show that, under this alternative specification, the model's implications for investor beliefs remain similar.

Next, we turn to rational arbitrageurs, who hold fully rational beliefs: they understand the dividend process in equation (1); they observe parameter  $\mu$  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. Because  $P_t$  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

$$\max_{N_t^i} \left( \mathbb{E}_t^i[dW_t^i] - \frac{\gamma}{2} \text{Var}_t^i[dW_t^i] \right), \quad (5)$$

subject to the budget constraint on their wealth  $W_t^i$

$$dW_t^i = rW_t^i dt - rN_t^i P_t dt + N_t^i dP_t + N_t^i D_t dt, \quad (6)$$

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  $\gamma$  represents risk aversion. For simplicity,  $\gamma$  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 there are overlapping generations of investors (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  $W_t^i - QP_t$  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 end of the period.<sup>5</sup>

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

$$\mu N_t^r + (1 - \mu)N_t^l = Q \quad (7)$$

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

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

and

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

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

where  $\bar{\omega}_t \equiv \int_{-\infty}^t \delta e^{-\delta(t-s)} d\tilde{\omega}_s^P$  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,  $\theta_t$ , but also a contrarian component  $-\sigma_P \alpha \bar{\omega}_t$ . This contrarian term is directly derived from the assumption we made in equation (4) about the gambler's fallacy. Second, in equation (10),  $\bar{\omega}_t$  decays at the rate of  $\alpha\delta + \delta$  rather than  $\delta$ :  $\bar{\omega}_t$  is constructed as a weighted average of past  $d\tilde{\omega}_s^P$ , where the declining weight leads to a baseline decay rate of  $\delta$  in  $\bar{\omega}_t$ ; in addition, the gambler's fallacy implies that LSN investors expect a negative serial autocorrelation in  $d\tilde{\omega}_t^P$ , causing an extra decay rate of  $\alpha\delta$  in  $\bar{\omega}_t$ .

Note that, in the above belief-formation process, LSN investors do not observe  $\theta_t$  and  $\bar{\omega}_t$ : as in [Rabin and Vayanos \(2010\)](#), they use Bayesian inference to estimate both quantities.<sup>6,7</sup> Specific-

<sup>5</sup>Alternatively, investors can be thought of as being infinitely-lived; but they reset their demand to  $Q$  shares every  $dt$  period.

<sup>6</sup>These estimated quantities in turn guide LSN investors' trading decisions.

<sup>7</sup>Our model involves biased learning from equilibrium prices. As shown in equation (4), LSN investors incorrectly

cally, 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  $\theta_t$  and  $\bar{\omega}_t$  using past prices as informative signals. The conditional mean and variance of  $\boldsymbol{\theta}_t \equiv (\theta_t, \bar{\omega}_t)$  are denoted as

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

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

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

and

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

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

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

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 recent 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

---

believe that  $d\tilde{\omega}_t^P$  has a negative serial autocorrelation; based on this incorrect belief, investors engage in Bayesian learning. A recent study by [Bastianello and Fontanier \(2022\)](#) analyzes biased learning from equilibrium prices in a different context. In their model, investors incorrectly learn from prices because they engage in “partial equilibrium thinking”—they fail to recognize that many other investors are also learning from prices.

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  $\gamma_{11}$ ,  $\gamma_{12}$ , and  $\gamma_{22}$  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

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

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

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

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

and

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

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

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

The optimal share demands for the risky asset from LSN investors and from rational arbitrageurs are

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

where  $\eta_0^l$ ,  $\eta_1^l$ ,  $\eta_2^l$ ,  $\eta_0^r$ ,  $\eta_1^r$ , and  $\eta_2^r$  are functions of  $A$ ,  $B$ ,  $C$ , and  $\sigma_P$ . ■

The proof of Proposition 1, the expressions of  $\eta_0^l$ ,  $\eta_1^l$ ,  $\eta_2^l$ ,  $\eta_0^r$ ,  $\eta_1^r$ , and  $\eta_2^r$ , and the numerical procedure that solves for  $A$ ,  $B$ ,  $C$ , and  $\sigma_P$  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  $g_D = 0.05$ ,  $\sigma_D = 0.25$ ,  $r = 0.025$ , and  $Q = 1$ . For risk preferences, we set  $\gamma = 0.01$ . Moreover, we set  $\mu = 0.5$ , so rational arbitrageurs make up 50% 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

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

where  $\gamma_{11} = -\kappa\sigma_P^2 + \sqrt{(\kappa\sigma_P^2)^2 + \sigma_\theta^2\sigma_P^2}$  and is decreasing in  $\kappa$ . For belief parameters, we set  $\bar{\theta} = g_D/r = 2$  and vary the values of  $\kappa$  and  $\sigma_\theta$  for comparative statics. We first discretize the continuous-time model and simulate a time series of 10,000 years at the monthly frequency.<sup>8</sup> 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 values of  $\kappa$  and  $\sigma_\theta$ . Specifically, each line plots the coefficients from regressing 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. In all these plots, beliefs load positively on past price changes, consistent

---

<sup>8</sup>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.



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  $\kappa$  between 0.01 and 1. In these plots, a smaller  $\kappa$  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  $\kappa$ , investors believe that the asset quality  $\theta_t$  can persistently deviate from its long-term mean  $\bar{\theta}$  and hence exhibit high variability. As such, when 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  $\kappa$ , investors believe that  $\theta_t$  tends to quickly mean-revert towards  $\bar{\theta}$  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—term  $\sigma_P d\bar{\omega}_t^l$  in equation (12)—and only infer a small increase in  $m_{t,1}$ . As such, investors do not significantly adjust their forecast of the future price change. In Panel B, we vary parameter  $\sigma_\theta$  between 2.5 and 10. In these plots, a larger  $\sigma_\theta$  is associated with a higher degree of extrapolation. When  $\sigma_\theta$  is high, investors perceive high variability of  $\theta_t$ . Therefore, upon observing a positive price change, 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 into the model. Specifically, we set  $\alpha = 0.5$ , so that investors perceive random errors to be negatively autocorrelated. For the rest of the parameters, we set  $\kappa = 0.05$ ,  $\bar{\theta} = g_D/r = 2$ ,  $\sigma_\theta = 5$ , and  $\delta = 2.77$ , where this value of  $\delta$  indicates a look-back window of about six months; specifically, when forming beliefs about  $\bar{\omega}_t$  defined below equation (10), LSN investors assign a 25% weight on a 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  $\gamma_{11} = 53.90$ ,  $\gamma_{12} = -2.68$ , and  $\gamma_{22} = 0.14$ . For the equilibrium price in equation (15), we obtain  $A = -45.5$ ,  $B = 1.86$ ,  $C = -0.30$ , and  $\sigma_P = 17.45$ . Finally, for the share demands described in equation (20), we obtain  $\eta_0^l = 0.37$ ,  $\eta_1^l = 0.31$ ,  $\eta_2^l = -2.86$ ,  $\eta_0^r = 1.63$ ,  $\eta_1^r = -0.31$ , and  $\eta_2^r = 2.86$ .

[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  $\alpha = 0.5$ . 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  $\alpha = 0.5$  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 half of investors are rational and the remaining half have the LSN beliefs with  $\alpha = 0.5$ . The comparison between the solid line and the dash-dot line shows that, relative to the rational beliefs about the future price change, LSN investors’ beliefs underreact to short-term trends; at the same time, they overreact to longer-term trends.

**[Place Fig. 3 about here]**

Fig. 3 examines how the two belief parameters regulating the LSN,  $\alpha$  and  $\delta$ , affect the dependence of investor beliefs on past price changes. Panel A is concerned with  $\alpha$ , which measures the overall strength of the gambler’s fallacy. When  $\alpha$  increases, not only does short-run mean-reversion increase in magnitude, longer-run extrapolation also increases. The simultaneous increase in both short-term contrarian beliefs and long-term extrapolative beliefs confirms that the LSN is a common driver of both phenomena. Panel B is concerned with  $\delta$ , which measures the relative weight put on recent versus distant past innovation terms. When  $\delta$  increases, investors believe that recent trends tend to mean-revert faster. 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 tend to 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 is an empirically robust pattern 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:  $g_D = 0.05$ ,  $\sigma_D = 0.25$ ,  $r = 0.025$ ,  $Q = 1$ ,  $\gamma = 0.01$ ,  $\mu = 0.5$ ,  $\kappa = 0.05$ ,  $\bar{\theta} = 2$ ,  $\sigma_\theta = 5$ ,  $\alpha = 0.5$ , 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 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 model 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 for positions held over a short period: 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 for positions held over a short period. Taken together, the model implies that return extrapolation and the disposition effect are not necessarily in conflict with each other. Instead, they may operate over different horizons and can be both microfounded by beliefs in 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 the result that LSN investors trade as “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 \(2022\)](#) 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.<sup>9</sup>

---

<sup>9</sup>It remains an open question why retail investors exhibit such contrasting behaviors when trading two different types of assets. One possible explanation, based on the incorrect belief in the LSN, is that investors have a less strong prior about the underlying data-generating process for cryptos than for stocks. As a result, they are more likely to

### 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,  $\alpha$  and  $\delta$ , vary. Table 2 shows that a higher degree of the gambler’s fallacy—measured by an increase in  $\alpha$ —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  $\delta$  is higher—we also find a stronger disposition effect. These findings lead to the following model 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  $\alpha$  or a higher  $\delta$ , 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  $\alpha$  or a higher  $\delta$ . 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 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 trading behavior 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  $\alpha = 0.5$ , LSN investors with 

---

behave as extrapolators. We leave a deeper investigation of this issue to future research.

$\alpha = 0$ , and rational arbitrageurs.<sup>10</sup> 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  $\alpha = 0.5$  simply as “LSN investors.”

[Place Fig. 6 about here]

Fig. 6 Panel A plots, separately for LSN investors and 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 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 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 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 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 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. In equilibrium, market clearing means that rational arbitrageurs must then hold the opposite beliefs—extrapolative over short-term trends and contrarian over longer-term trends, as observed from the dash-dot line in Fig. 2. These beliefs, being fully rational, imply that asset prices will exhibit short-term momentum and long-term reversals.

---

<sup>10</sup>The procedure that solves this more generalized model is given in Appendix C.

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,  $\alpha$  and  $\delta$ , affect asset prices. Earlier in Fig. 3, we showed that a higher  $\alpha$  or  $\delta$  means LSN investors being more contrarian over short-term trends and more extrapolative over longer-term trends. In response, 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. Indeed, in Fig. 8, an increase in  $\alpha$  or  $\delta$  gives rise to stronger patterns of short-term momentum and long-term reversals.

[Place Fig. 9 about here]

Our model also generates excess volatility: with the model parameters specified in Section 3.3, the implied volatility of price change,  $\sigma_P = 17.45$ , is significantly higher than the fundamental volatility of  $\sigma_D/r = 10$ . Fig. 9 further shows that, for a wide range of values of  $\alpha$  and  $\delta$ , our model generates excess volatility:  $\sigma_P$  remains significantly higher than  $\sigma_D/r$ . In Appendix D, we provide a further discussion about the co-existence of the disposition effect and excess volatility in the model.

### 3.6. *Alternative specification of LSN beliefs*

The baseline model described above follows a growing literature in behavioral finance that directly applies investors' belief-formation process to their perceived price process (Barberis and Shleifer, 2003; Barberis, Greenwood, Jin, and Shleifer, 2015, 2018; Jin and Sui, 2022; Liao et al., 2022). At the same time, a separate literature applies investors' belief-formation process to asset fundamentals rather than prices (Barberis et al., 1998; Scheinkman and Xiong, 2003; Basak, 2005; Hirshleifer, Li, and Yu, 2015; Nagel and Xu, 2022). In this section, we follow the latter



literature and consider an alternative specification in which the LSN is applied to the dividend process.

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

$$dD_t = g_D dt + \sigma_D d\omega_t^D. \quad (22)$$

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

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

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 is negatively auto-correlated. This is similar to the perceived price process specified in our baseline model.

The rest of the model can be summarized with the following steps. First, LSN investors update their beliefs about  $\theta_t$  and  $\bar{\omega}_t \equiv \int_{-\infty}^t e^{-\delta(t-s)} d\tilde{\omega}_s^D$  using past dividends as informative signals. Second, they derive beliefs about future price changes from their beliefs about future dividend changes; they then make trading decisions based on these beliefs about future price changes. Third, rational arbitrageurs hold rational beliefs about future price changes and trade according to these rational beliefs. Lastly, equilibrium price is conjectured and solved in a way that allows for market clearing of the risky asset. We leave a detailed description of the model to Appendix E.

**[Place Fig. 10 about here]**

For the alternative model, Fig. 10 plots the dependence of the LSN and rational beliefs about the future price change 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' 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 are similar to those from the baseline model; in both models, trading behavior and asset prices are completely driven by investor beliefs.

While the two models essentially produce the same set of results, we view the baseline model as psychologically more realistic, for the following reason. In that model, the LSN is directly applied to the perceived 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. Therefore, LSN investors apply a belief heuristic to directly guide their trading decisions. By contrast, under the alternative model, LSN investors need to take the extra step of *deriving* beliefs about price changes from their beliefs about dividend changes to make trading decisions. While this extra step of mapping dividend expectations to price expectations is theoretically straightforward, it may not realistically capture the thought process of real-world investors.

## 4. Evidence from investor behavior

### 4.1. Data

Our primary data set is from a large discount brokerage firm and contains individual-level transaction records from 1991 to 1996 (the brokerage data); more details about 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 that 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](#)). Focusing on the brokerage data allows us to benchmark our results to those from previous studies. Our data on stock prices and returns are from the Center for Research in Security Prices (CRSP). In addition to the brokerage data, in [Appendix H](#), we complement our analysis using data from a large brokerage firm in China.

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 period from 1991 to 1996, 2) are not common-share transactions, and 3) have negative commissions. Second, we follow [Hartzmark \(2015\)](#) and 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 only on long positions. This filter also excludes any trading history

that starts with a sell, making it possible to calculate the purchase price for each position. In this filtered sample, the summary statistics of transaction size, price per share, monthly turnover, commission, and spread resemble those reported in [Barber and Odean \(2000\)](#); Appendix F reports the details. Lastly, we focus on active investors by dropping investors whose total numbers of buys and sales are below the medians in the full sample distribution.

## 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, Fig. 5 and Prediction 1 posit 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 purchase or a sale. Panel A 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 purchases.<sup>11</sup> We equal-weight all observations; value-weighting using the transaction amount produces similar patterns. Given that the distribution of stock return is heavily skewed and we would like to characterize the return patterns of a typical trade, we report the median return rather than the average return.

[Place Fig. 11 about here]

Fig. 11 Panel A shows that a stock purchase is associated with the following return pattern: the stock tends to exhibit strong positive 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 Appendix G, we plot the same figure using only initial buys—that is, purchases of stocks that are not currently in the portfolio. Because these purchases happen before any accumulation of holding-period returns, non-belief forces that

---

<sup>11</sup>We use market-adjusted return, not raw return, because the sample period (1991–1996) witnessed a long bull market during which most stocks were winners. When we replace market-adjusted return with raw return, the overall pattern of lagged monthly returns prior to a purchase—in particular, the drop in monthly returns in more recent months—remains robust.

might drive investor behavior, such as realization utility and emotions, have not kicked in, hence allowing for a cleaner test for the role of beliefs. There, we observe a similar drop in returns right before a purchase takes place.

Fig. 11 Panel B concerns selling behavior. It shows that a stock sale 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, for the most recent month prior to the sale, there is a sudden and substantial increase in return; this suggests that investors are more inclined to sell stocks that have recently experienced an increase in price. Such behaviors are consistent with retail investors acting as contrarian traders in response to recent stock returns (Kaniel et al., 2008; Luo et al., 2022). Comparison between Fig. 5 and Fig. 11 suggests that the aggregate trading patterns observed in the brokerage data are generally consistent with our model’s predictions.<sup>12</sup>

The prior literature on investor behavior emphasizes the trend-chasing nature of retail trading (Odean, 1999; Barber, Odean, and Zhu, 2009; Liao et al., 2022). Consistent with this, there is ample evidence of return extrapolation in survey data (Greenwood and Shleifer, 2014; Da, Huang, and Jin, 2021). The patterns presented in Fig. 5 and Fig. 11, however, contain more nuances: while investors do ride on medium and long term trends, they are contrarian to more recent trends. This finding poses a challenge to most existing models: models of return extrapolation or diagnostic expectations generate extrapolative trading over *all* past returns, so they do not produce contrarian trading towards recent returns; models of mean-reverting beliefs do the opposite, generating contrarian trading over all past returns while failing to produce extrapolative trading towards longer-term returns. In contrast, our model of the law of small numbers naturally generates the observed opposing trading patterns towards recent and longer-terms returns.

#### 4.2.2. Stock-level evidence

To provide further evidence in support of Prediction 1, we run stock-level regressions. The idea here is to test, in the cross-section of individual stocks, what type of past return pattern attracts

---

<sup>12</sup>There is one notable discrepancy: the model suggests that investors should sell long-term losers, whereas in the actual data, investors tend to buy *and* sell long-term winners. This discrepancy may arise from two channels. First, in the model, investors continuously adjust their stock holdings, whereas empirically, investors buy a stock first, hold it for a while, and sell it later. Given that investors tend to buy long-term winners to begin with, the stocks being sold tend to also be long-term winners. Second, in the model, selling behavior is completely driven by investor beliefs; however, in the data, non-traditional preferences may induce investors to sell long-term winners (Barberis and Xiong, 2012; Ingersoll and Jin, 2013).

more buying or selling. Specifically, for each date and stock, we aggregate the number of trades for all buys and sells as Buy and Sell. We then consider two measures of trading propensity: the first one is  $(\text{Buy} - \text{Sell}) / (\text{Buy} + \text{Sell})$  and the second one is simply  $\text{Buy} - \text{Sell}$ . We regress the two measures of trading propensity on past stock returns, controlling for date and stock fixed effects. Table 3 reports the regression results, with double-clustered standard errors reported in parentheses.

[Place Table 3 about here]

Column (1) regresses trading propensity, measured by  $(\text{Buy} - \text{Sell}) / (\text{Buy} + \text{Sell})$ , on lagged month returns over the last three months. It shows that heightened selling activity is associated with stocks that have recently experienced price increases. Column (2) then shows 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. Column (3) analyzes a different measure of trading activity and documents consistent evidence that investors tend to buy short-term losers. Column (4) finds that this trading propensity decays over a longer horizon but does not turn positive as in Column (2). Overall, these results replicate the patterns documented in Barber et al. (2009).<sup>13</sup>

### 4.3. *The disposition effect*

#### 4.3.1. *Aggregate evidence*

According to Prediction 2, investors expect short-term trends to reverse. This, on average, leads to a disposition effect: because investors expect current winners to underperform and current losers to outperform in the future, they sell winners and hold on to losers. Prediction 2 further states that the disposition effect is more pronounced for positions associated with a shorter holding period. When the holding period gets longer, extrapolative beliefs begin to have a more significant impact on their trading responses to long-term returns, thereby reducing the disposition effect.

[Place Fig. 12 about here]

---

<sup>13</sup>In Appendix H, we run similar regressions using data from a large Chinese brokerage firm. We find similar evidence of contrarian trading in response to recent returns and extrapolative trading to distant returns. The main difference is that Chinese retail investors have a much shorter look-back window.

In Fig. 12, we test this prediction. Panel A displays the overall propensities of selling winners and losers for daily portfolio holdings, confirming the existence of an overall disposition effect.<sup>14</sup> On an average day, 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, where 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 longer periods. For positions held for more than a year, the propensities of selling winners and losers are virtually the same.<sup>15</sup>

The weakened disposition effect for long-term holdings poses a challenge to existing theories. The current explanations of the disposition effect include prospect theory (Odean, 1998), realization utility (Barberis and Xiong, 2009, 2012; Ingersoll and Jin, 2013; Liao et al., 2022), belief revisions (Ben-David and Hirshleifer, 2012) and other cognitive forces such as cognitive dissonance (Chang et al., 2016; Frydman et al., 2018). In general, these explanations consider gains or losses relative to the purchase price without specifying a mechanism to differentiate gains or losses across different holding periods. As a result, these explanations usually do not make direct predictions about how the disposition effect interacts with the holding period.<sup>16</sup> In contrast, Table 1 shows that the documented horizon-dependent pattern of the disposition effect naturally arises from our model of the LSN. Our results suggest that short-term contrarian beliefs and long-term extrapolation—the natural implications of the LSN—can be an important driver of the disposition effect.

---

<sup>14</sup>In the original study by Odean (1998), the disposition effect is measured based on holdings on days when selling happens. Here, our paper follows Ben-David and Hirshleifer (2012) and measures the disposition effect based on all daily holdings.

<sup>15</sup>In Table 1 of Ben-David and Hirshleifer (2012), the authors show the robustness of the disposition effect across different holding periods. Although their paper’s emphasis is on the V-shaped selling propensity, the same table also shows a weakened disposition effect for positions over longer holding periods.

<sup>16</sup>For example, in models of realization utility, utility is defined over the holding-period return since purchase. It is conceivable that realization utility, in conjunction with a slow-moving reference point that is affected by recent stock prices, might explain why the disposition effect becomes weaker as the holding period increases. However, such a theory is yet to be developed.

#### 4.3.2. Additional buying behavior

Our model predicts that, when investors buy additional shares of stocks they already own in their portfolio, they exhibit a pattern similar to the disposition effect. In particular, 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 stronger for shorter holding periods. It is most pronounced for positions with a holding period less than a quarter. And the degree of doubling down, when measured by the difference between the two propensities to buy, monotonically decreases in the holding horizon.<sup>17</sup>

#### 4.4. Disposition effect and doubling down

Our model not only predicts the disposition effect and doubling down at the aggregate level, but also suggests a direct link between these two phenomena at the *individual* 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 each investor who has made at least ten buys, we compute the stock return in the most recent month before a buy, averaged across all buys. The resulting measure serves as a proxy for an investor’s degree of doubling down.<sup>18</sup> We then use this measure to sort all investors into five groups, with Group 1 being the most extrapolative and Group 5 being the most contrarian. Fig. 14 Panel A validates 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

---

<sup>17</sup>If the degree of doubling down is instead measured by the ratio between these two propensities, then we do not observe a decreasing pattern that is strictly monotonic.

<sup>18</sup>Our main analysis uses the value-weighted average across all buys, but the results remain essentially unchanged if we instead use the equal-weighted average or if we just use initial buys.

LSN beliefs while Group 1 the least. Panel B compares the propensities of selling winners and losers 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 for Group 5. It is worth noting that even in Group 1, we see a positive disposition effect. This implies that other psychological forces such as realization utility, emotions, and cognitive dissonance may jointly drive the disposition effect alongside the LSN.

[Place Fig. 14 about here]

Our model predicts that most of the trading responses are due to investors reacting to recent returns. In Fig. 15, we plot the buying and selling propensities of the five groups of investors; here, we classify gains and losses based on the most recent one-month return, as opposed to the holding-period return. Again, we observe similar patterns across the five groups: both the degree of doubling down and the degree of the disposition effect increase from Group 1 to Group 5.

[Place Fig. 15 about here]

The consistency between buying and selling behavior has additional implications for theories of the disposition effect. Existing models of the disposition effect have focused on the selling side, being able to generate the tendency of selling winners and holding on to losers (Odean, 1998; Barberis and Xiong, 2009, 2012; Ingersoll and Jin, 2013). However, these models do not directly generate doubling down in buying or a consistency between buying and selling behavior. In this regard, documenting a tight relationship between the disposition effect and doubling down adds an additional moment for theories to match. In our LSN model, contrarian beliefs over short-term trends are responsible for both buying and selling decisions.<sup>19</sup>

#### 4.5. *Selling propensity as a function of past returns*

The literature shows that the aggregate selling propensity as a function of the holding-period return is V-shaped (Ben-David and Hirshleifer, 2012). That is, the probability of selling a position increases in the extremeness of the holding-period return. In this section, we revisit the V-shape

---

<sup>19</sup>Frydman and Camerer (2016) provide experimental evidence that connects a particular type of buying decisions—whether investors repurchase a stock that they recently sold—with selling decisions. Their paper argues that regret serves as a driver of both buying and selling behavior.



selling propensity by testing our model’s predictions on the relationship between investors’ selling propensity and their holding-period returns. In particular, according to Prediction 6 from Section 3.4, LSN investors’ selling propensity should depend positively on recent returns, while pure extrapolators’ selling propensity should depend negatively on recent returns.

We first confirm the existence of the V-shape by plotting, in Fig. 16, the probability of selling as a function of holding-period returns, using the daily portfolios constructed from the brokerage data. Here, we consider only positions with a prior holding period of less than one month, because the V-shape is strongest when the holding period is short (e.g., within 20 days according to Ben-David and Hirshleifer (2012)). In addition, to control for the rank effect (Hartzmark, 2015) in subsequent regressions, we require a daily portfolio to contain at least five positions. The binscatter plot in Fig. 16 clearly shows a V-shape: on a given day, the probability of selling the most extreme winner or loser is around 2%; in contrast, the probability of selling a stock with a holding-period return of around zero is about 0.5%.

**[Place Fig. 16 about here]**

For the five investor groups sorted above, we examine the selling propensity as a function of the holding period return. Fig. 17 finds substantial heterogeneity across the five groups. For investors in Group 1, when the holding-period return is negative, the probability of selling decreases monotonically in the holding-period return. However, when the holding-period return is positive, the probability of selling is insensitive to changes in the holding-period return. In contrast, for investors in Group 5, the probability of selling as a function of the holding-period return is flat in the loss region and increasing in the gain region.

**[Place Fig 17 about here]**

To examine the differences across groups more formally, Table 4 studies the relationship between selling and holding-period returns using regressions. In each column, we regress a dummy variable indicating selling on measures of past return patterns while controlling for the rank effect and various fixed effects. In Column (1), we see a V-shape: the selling propensity is decreasing in the holding-period return in the loss region but increasing in the gain region. In Columns (2) to (6), we run the same regression for the five investor groups separately. Consistent with Fig. 17, the

strength and shape of the V-shape vary substantially across groups: Groups 1 and 2 are sensitive to returns only in the loss region, while Groups 4 and 5 are only sensitive in the gain region.

[Place Table 4 about here]

An interesting interpretation for the V-shape emerges from the evidence in Fig. 17 and Table 4: the two arms of the overall V-shape may be driven by two separate investor groups. Extrapolators, represented by Group 1, overwhelmingly sell past losers while LSN investors, represented by Group 5, overwhelmingly sell past winners. Taken together, these results not only provide further support to our LSN model, but also shed light on the underlying mechanisms responsible for the V-shape (Ben-David and Hirshleifer, 2012). Indeed, Ben-David and Hirshleifer (2012) conjecture that the V-shape may stem from investors’ belief revisions. Here, we specify a particular form of belief revision—through the LSN—and show that it indeed can make sense of the puzzling V-shape.

Lastly, it is worth noting that our model makes predictions about the sensitivity of trading to *recent* returns, which overlap but do not coincide exactly with *holding-period* returns. In Appendix I, we reexamine the probability of selling and the probability of buying, each as a function of the past one-month return, and find supportive evidence for Prediction 6.

## 5. Conclusion

A belief in 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 short-term momentum, long-term reversals, and excess volatility. 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.

Our model makes predictions about both asset prices and investor behavior. The current paper tests the model's predictions about investor behavior. In Appendix J, we provide a preliminary empirical test of the model prediction that stocks associated with more pronounced LSN beliefs—stocks with a higher  $\alpha$  or a higher  $\delta$ —should exhibit both stronger short-term momentum and stronger long-term reversals. We find supportive evidence that stocks traded more by LSN investors do experience stronger momentum and reversals. Future research can explore the model's asset pricing predictions more thoroughly, both in time-series and cross-sectional analyses.

## References


- Bansal, R., Yaron, A., 2004. Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles. *Journal of Finance* 59, 1481–1509.
- 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.
- Barber, B. M., Odean, T., 2013. The Behavior of Individual Investors. *Handbook of the Economics of Finance* 2, 1533–1570.
- Barber, B. M., Odean, T., Zhu, N., 2009. Systematic Noise. *Journal of Financial Markets* 12, 547–569.
- Barberis, N., Greenwood, R., Jin, L., Shleifer, A., 2015. X-CAPM: An Extrapolative Capital Asset Pricing Model. *Journal of Financial Economics* 115, 1–24.
- Barberis, N., Greenwood, R., Jin, L., Shleifer, A., 2018. Extrapolation and Bubbles. *Journal of Financial Economics* 129, 203–227.
- Barberis, N., Huang, M., Santos, T., 2001. Prospect Theory and Asset Prices. *Quarterly Journal of Economics* 116, 1–53.
- Barberis, N., Shleifer, A., 2003. Style Investing. *Journal of Financial Economics* 68, 161–199.
- 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.

- Bastianello, F., Fontanier, P., 2022. Expectations and Learning from Prices. Working paper.
- 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.
- Campbell, J. Y., Cochrane, J. H., 1999. By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior. *Journal of Political Economy* 107, 205–251.
- 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.
- Da, Z., Huang, X., Jin, L. J., 2021. Extrapolative beliefs in the cross-section: What can we learn from the crowds? *Journal of Financial Economics* 140, 175–196.
- Frydman, C., Camerer, C., 2016. Neural Evidence of Regret and Its Implications for Investor Behavior. *Review of Financial Studies* 29, 3108–3139.
- 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., Shleifer, A., 2014. Expectations of Returns and Expected Returns. *Review of Financial Studies* 27, 714–746.
- 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.
- Jin, L., Sui, P., 2022. Asset Pricing with Return Extrapolation. *Journal of Financial Economics* 145, 273–295.
- 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.
- Luo, C., Ravina, E., Sammon, M., Viceira, L. M., 2022. Retail Investors’ Contrarian Behavior Around News, Attention, 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.
- Odean, T., 1999. Do Investors Trade Too Much? *American Economic Review* 89, 1279–1298.
- 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.
- Tegua, A., 2017. Law of Small Numbers and Hysteresis in Asset Prices and Portfolio Choices. Working paper.
- 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 Survey and a Field Experiment. Working paper.

Panel A: Effect of  $\kappa$



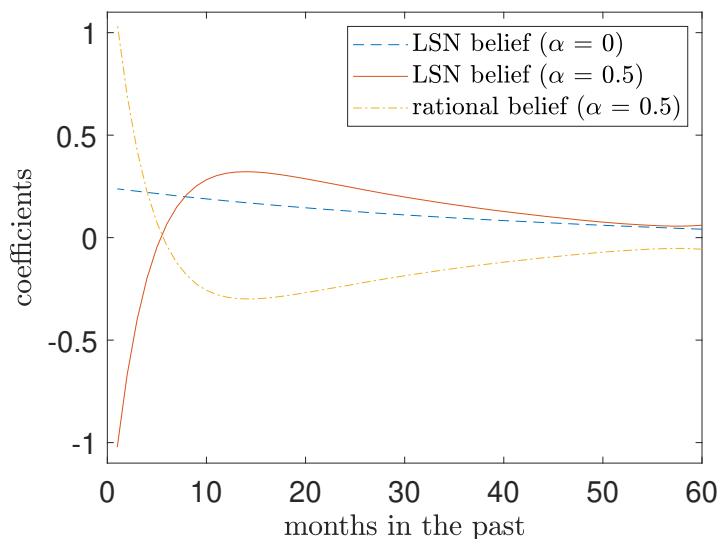
Panel B: Effect of  $\sigma_\theta$



**Fig. 1.** Dependence of the LSN beliefs on past price changes: the  $\alpha = 0$  case.

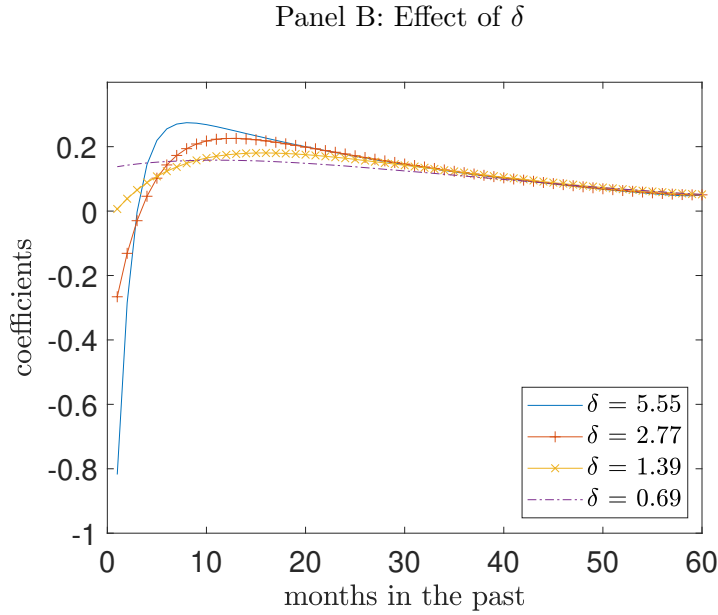

The figure plots, for different values of  $\kappa$  and  $\sigma_\theta$ , the coefficients from regressing LSN investors' beliefs about the future price change,  $\mathbb{E}_t^l(dP_t)/dt = m_{t,1}$ , on price changes over the past 60 months. The default values of  $\kappa$  and  $\sigma_\theta$  are 0.05 and 5, respectively. The other parameters are:  $g_D = 0.05$ ,  $\sigma_D = 0.25$ ,  $r = 0.025$ ,  $Q = 1$ ,  $\delta = 2.77$ ,  $\alpha = 0$ ,  $\bar{\theta} = 2$ ,  $\gamma = 0.01$ , and  $\mu = 0.5$ .






**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  $\mu$  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  $\mu$  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:  $g_D = 0.05$ ,  $\sigma_D = 0.25$ ,  $r = 0.025$ ,  $Q = 1$ ,  $\kappa = 0.05$ ,  $\sigma_\theta = 5$ ,  $\delta = 2.77$ ,  $\bar{\theta} = 2$ ,  $\gamma = 0.01$ , and  $\mu = 0.5$ .




**Fig. 3.** Dependence of the LSN beliefs on past price changes: comparative statics.

The figure plots, for different values of  $\alpha$  and  $\delta$ , 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  $\alpha$  and  $\delta$  are 0.5 and 2.77, respectively. The other parameters are:  $g_D = 0.05$ ,  $\sigma_D = 0.25$ ,  $r = 0.025$ ,  $Q = 1$ ,  $\kappa = 0.05$ ,  $\sigma_\theta = 5$ ,  $\bar{\theta} = 2$ ,  $\gamma = 0.01$ , and  $\mu = 0.5$ .



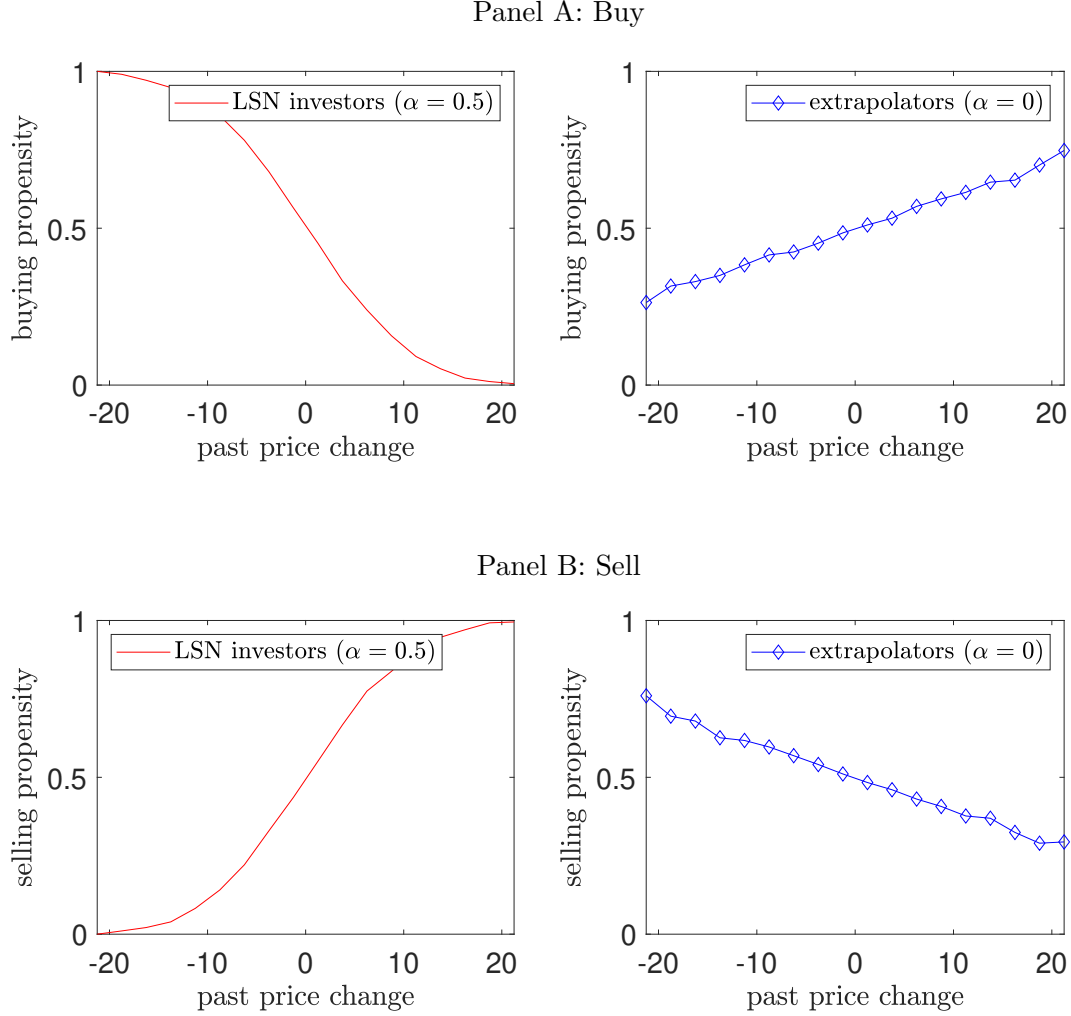
**Fig. 4.** Dependence of the change in LSN investors' demand on past price changes.

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:  $g_D = 0.05$ ,  $\sigma_D = 0.25$ ,  $r = 0.025$ ,  $Q = 1$ ,  $\kappa = 0.05$ ,  $\sigma_\theta = 5$ ,  $\alpha = 0.5$ ,  $\delta = 2.77$ ,  $\bar{\theta} = 2$ ,  $\gamma = 0.01$ , and  $\mu = 0.5$ .




**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:  $g_D = 0.05$ ,  $\sigma_D = 0.25$ ,  $r = 0.025$ ,  $Q = 1$ ,  $\kappa = 0.05$ ,  $\sigma_\theta = 5$ ,  $\alpha = 0.5$ ,  $\delta = 2.77$ ,  $\bar{\theta} = 2$ ,  $\gamma = 0.01$ , and  $\mu = 0.5$ .



**Fig. 6.** Heterogeneous trading responses to past price changes.


We analyze a model with three types of investors: LSN investors with  $\alpha = 0.5$ , 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:  $g_D = 0.05$ ,  $\sigma_D = 0.25$ ,  $r = 0.025$ ,  $Q = 1$ ,  $\kappa = 0.05$ ,  $\sigma_\theta = 5$ ,  $\delta = 2.77$ ,  $\bar{\theta} = 2$ , 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:  $g_D = 0.05$ ,  $\sigma_D = 0.25$ ,  $r = 0.025$ ,  $Q = 1$ ,  $\kappa = 0.05$ ,  $\sigma_\theta = 5$ ,  $\alpha = 0.5$ ,  $\delta = 2.77$ ,  $\bar{\theta} = 2$ ,  $\gamma = 0.01$ , and  $\mu = 0.5$ .

Panel A: Effect of  $\alpha$

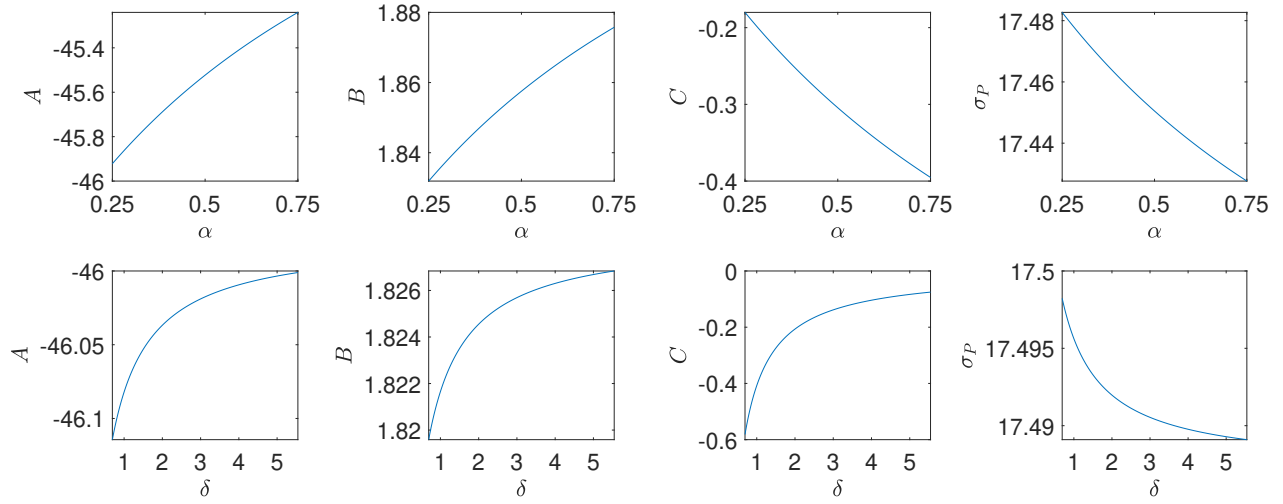


Panel B: Effect of  $\delta$



**Fig. 8.** Autocorrelation of price changes: comparative statics.

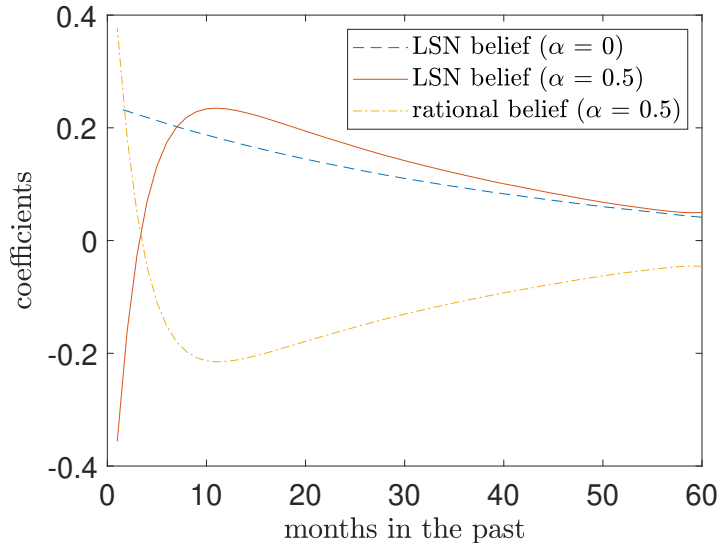
The figures plot, for different values of  $\alpha$  and  $\delta$ , the time-series correlation between the price change over the next  $n$  months and the price change over the past  $n$  months, where  $n$  goes from 1 to 60. The default values of  $\alpha$  and  $\delta$  are 0.5 and 2.77, respectively. The other parameter values are:  $g_D = 0.05$ ,  $\sigma_D = 0.25$ ,  $r = 0.025$ ,  $Q = 1$ ,  $\kappa = 0.05$ ,  $\sigma_\theta = 5$ ,  $\bar{\theta} = 2$ ,  $\gamma = 0.01$ , and  $\mu = 0.5$ .



**Fig. 9.** Model solution as function of  $\alpha$  and  $\delta$ .

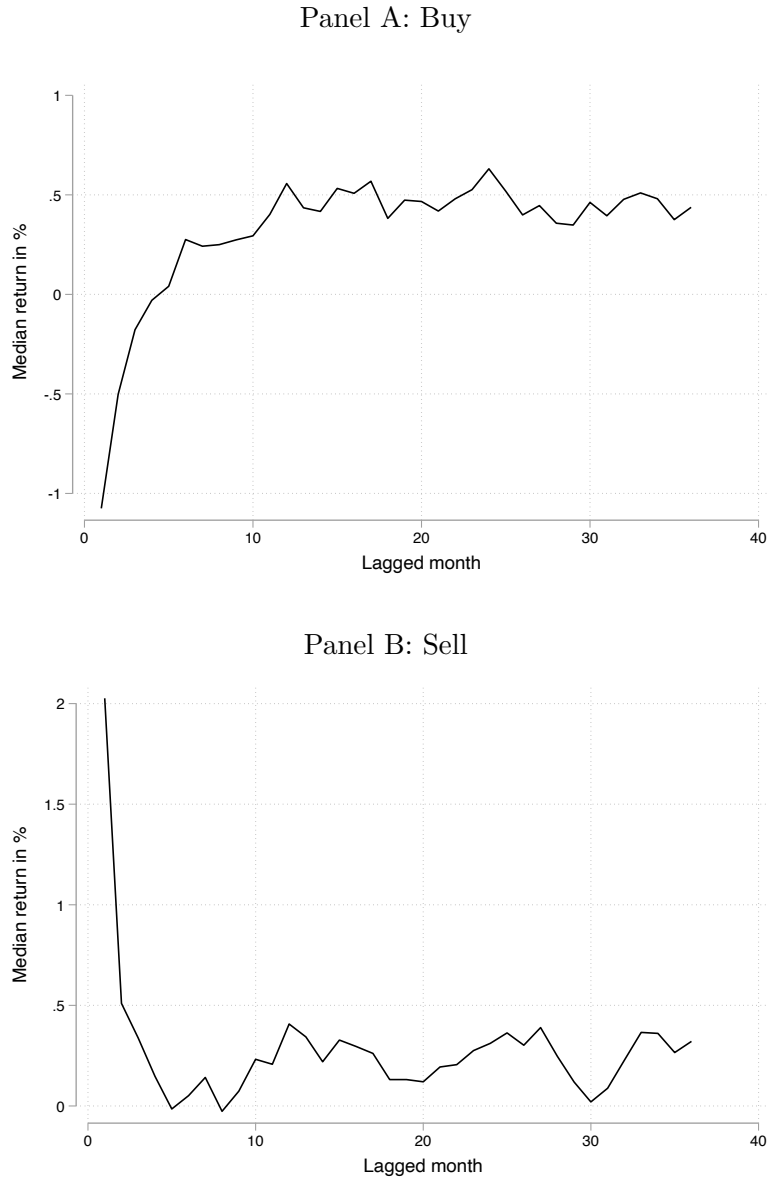
The upper panel plots the model solution—the coefficients  $A$ ,  $B$ , and  $C$ , and the price volatility  $\sigma_P$ —as function of  $\alpha$ . The lower panel plots the same quantities as function of  $\delta$ . The default values of  $\alpha$  and  $\delta$  are 0.5 and 2.77, respectively. The other parameters are:  $g_D = 0.05$ ,  $\sigma_D = 0.25$ ,  $r = 0.025$ ,  $Q = 1$ ,  $\kappa = 0.05$ ,  $\sigma_\theta = 5$ ,  $\bar{\theta} = 2$ ,  $\gamma = 0.01$ , and  $\mu = 0.5$ .





**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 E, on past price changes.


The figure plots the coefficients from regressing either LSN investors' beliefs about the future price change,  $\mathbb{E}_t^i(dP_t)/dt$ , 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  $\mu$  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  $\mu$  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:  $g_D = 0.05$ ,  $\sigma_D = 0.25$ ,  $r = 0.025$ ,  $Q = 1$ ,  $\kappa = 0.05$ ,  $\sigma_\theta = 0.125$ ,  $\delta = 2.77$ ,  $\bar{\theta} = 0.05$ ,  $\gamma = 0.01$ , and  $\mu = 0.5$ .




**Fig. 11.** Return patterns before trading.

This figure plots the return patterns before buys and sells, using transactions observed in the brokerage data. Detailed descriptions about the data and the filters used in constructing the data set can be found in Section 4.1. 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.

This figure examines the propensities of selling winners and losers, using transactions observed in the brokerage data. Detailed descriptions about the data and the filters used in constructing the data set can be found in Section 4.1. Each bar plots, on a random day, the probability of selling a stock conditional on it being at a gain or a loss. Gains and losses are defined based on the purchase price and the most recent closing price. Panel A concerns all positions for all active investors. Panel B concerns five subsamples based on the length of the holding period.

Panel A: Additional buying




Panel B: Additional buying at different holding periods




**Fig. 13.** Additional buying.

This figure examines the propensities of buying winners and losers, using transactions observed in the brokerage data. Detailed descriptions about the data and the filters used in constructing the data set can be found in Section 4.1. Each bar plots, on a random day, the probability of buying a stock conditional on it being at a gain or a loss. Gains and losses are defined based on the purchase price and the most recent closing price. Panel A concerns all positions for all active investors. Panel B concerns five subsamples based on the length of the holding period.

Panel A: Buying



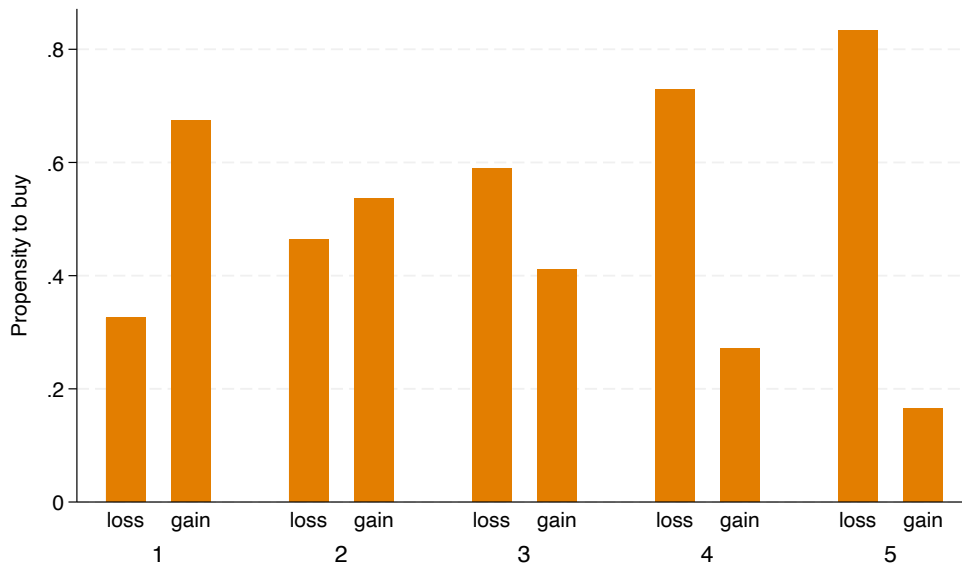
Panel B: Selling




**Fig. 14.** Consistency between buying and selling behavior.

In Panel A, investors are first sorted into five groups based on their degree of doubling down, measured by the stock return in the most recent month before a buy, averaged across all buys. For each group, each bar plots, on a random day, the probability of buying a stock conditional on it being at a gain or a loss. Gains and losses are defined based on the purchase price and the most recent closing price. In Panel B, each bar plots, on a random day, the probability of selling a stock conditional on it being at a gain or a loss.

Panel A: Buying

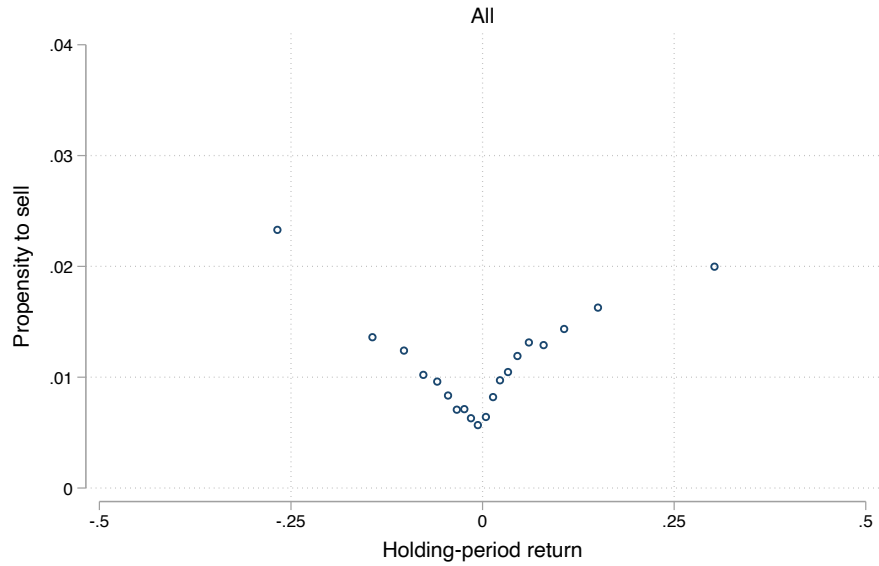


Panel B: Selling



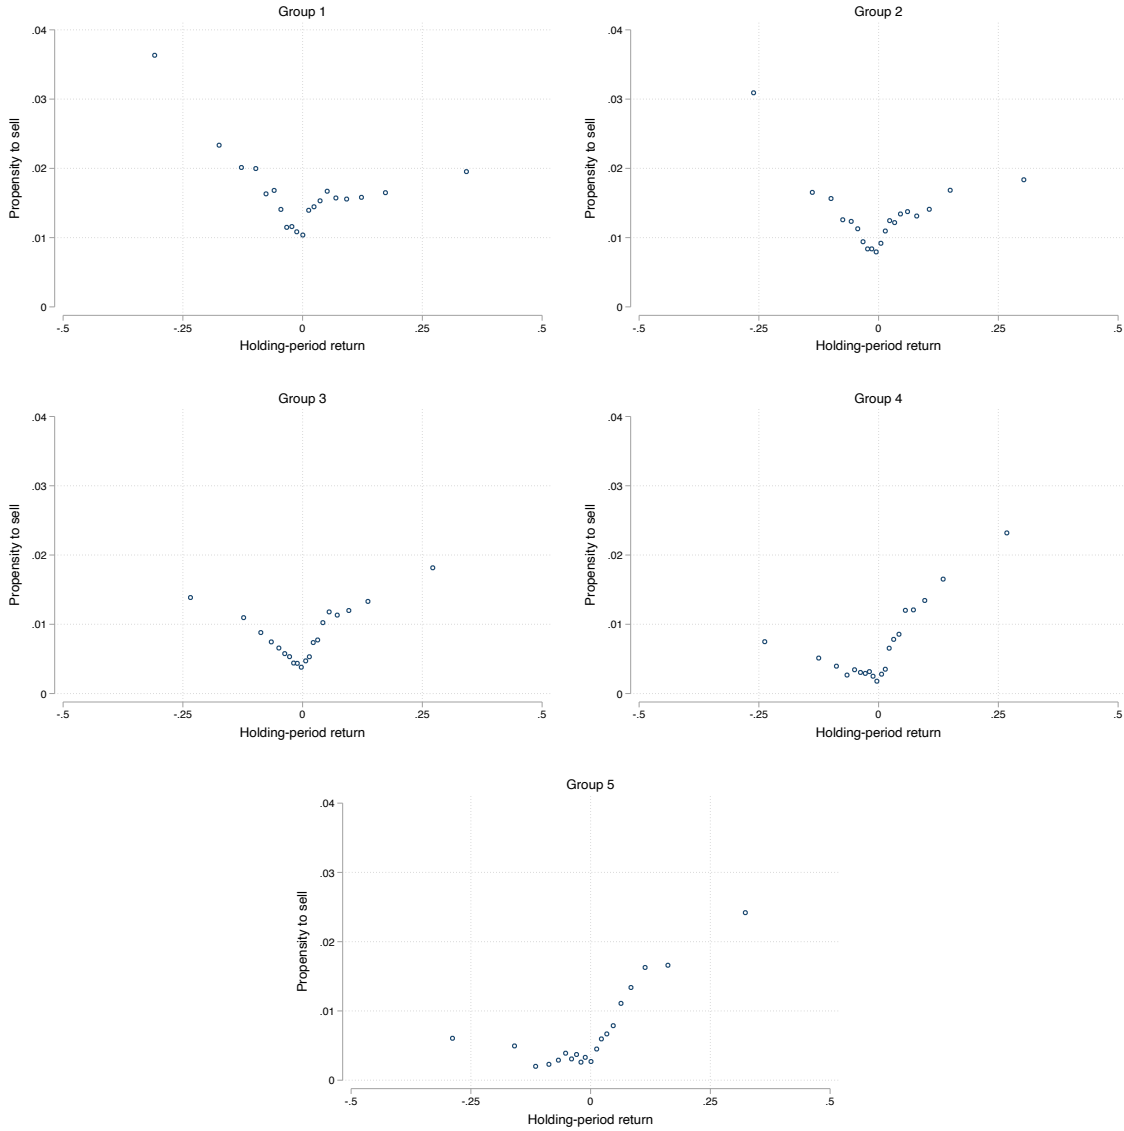
**Fig. 15.** Consistency between buying and selling behavior (one-month return).

All investors are first sorted into five groups based on their degree of doubling down, measured by the stock return in the most recent month before a buy, averaged across all buys. In Panel A, for each group, each bar plots, on a random day, the probability of buying a stock conditional on it being at a gain or a loss over the last month. In Panel B, each bar plots, on a random day, the probability of selling a stock conditional on it being at a gain or a loss over the last month.



**Fig. 16.** V-shape selling behavior in aggregate.

This figure examines the probability of selling as a function of holding-period returns, using the daily portfolios constructed from the brokerage data. Detailed descriptions about the data and the filters used in constructing the daily portfolios can be found in Section 4.1. We only consider positions with a prior holding-period of less than one month. In addition, we require a daily portfolio to contain at least five positions. Each dot plots, on a random day, the probability of selling a stock conditional on the holding-period return, shown in the  $x$ -axis.



**Fig. 17.** V-shape selling behavior by group.

This figure examines the probability of selling as a function of holding-period returns, using the daily portfolios constructed from the brokerage data. Detailed descriptions about the data and the filters used in constructing the daily portfolios can be found in Section 4.1. All investors are first sorted into five groups based on their degree of doubling down, measured by the stock return in the most recent month before a buy, averaged across all buys. Each plot then represents one group, with Group 1 having the lowest degree of doubling down while Group 5 highest. When calculating the probability of selling, we only consider positions with a prior holding-period of less than one month. In addition, we require a daily portfolio to contain at least five positions. In a given panel, each dot plots, on a random day, the probability of selling a stock conditional on the holding-period return, shown in the  $x$ -axis.



	Past horizon			
	1M	1M to 1Q	1Q to 1Y	1Y to 5Y
Buy at gain	18,510	20,073	33,151	46,083
Sell at gain	43,083	42,107	30,727	25,706
<i>Propensity of selling at gain</i>	<i>69.9%</i>	<i>67.7%</i>	<i>48.1%</i>	<i>35.8%</i>
Buy at loss	41,696	40,133	27,055	14,123
Sell at loss	16,651	17,627	29,007	34,028
<i>Propensity of selling at loss</i>	<i>28.5%</i>	<i>30.5%</i>	<i>51.7%</i>	<i>70.7%</i>
<i>Disposition effect</i>	<i>2.45</i>	<i>2.22</i>	<i>0.93</i>	<i>0.51</i>

**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:  $g_D = 0.05$ ,  $\sigma_D = 0.25$ ,  $r = 0.025$ ,  $Q = 1$ ,  $\gamma = 0.01$ ,  $\mu = 0.5$ ,  $\kappa = 0.05$ ,  $\bar{\theta} = 2$ ,  $\sigma_\theta = 5$ ,  $\alpha = 0.5$ , 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: $\alpha = 0.5, \delta = 2.77$	2.45	2.22	0.93	0.51
Low $\alpha$ : $\alpha = 0.25, \delta = 2.77$	1.81	1.55	0.69	0.44
High $\alpha$ : $\alpha = 0.75, \delta = 2.77$	2.85	2.75	1.11	0.54
Low $\delta$ : $\alpha = 0.5, \delta = 1.39$	1.65	1.68	1.04	0.34
High $\delta$ : $\alpha = 0.5, \delta = 5.55$	4.27	2.19	0.70	0.70

**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:  $g_D = 0.05$ ,  $\sigma_D = 0.25$ ,  $r = 0.025$ ,  $Q = 1$ ,  $\gamma = 0.01$ ,  $\mu = 0.5$ ,  $\kappa = 0.05$ ,  $\bar{\theta} = 2$ ,  $\sigma_\theta = 5$ ,  $\alpha = 0.5$ , 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 return, 1M	-0.304*** (0.021)		-0.429* (0.167)	
Lagged return, 2M	-0.235*** (0.014)		-0.439*** (0.044)	
Lagged return, 3M	-0.125*** (0.012)		-0.219*** (0.046)	
Lagged return, 1Q		-0.170*** (0.013)		-0.314*** (0.054)
Lagged return, 2Q		-0.0293*** (0.008)		-0.143*** (0.025)
Lagged return, 3Q		0.0212** (0.008)		-0.0594* (0.029)
Lagged return, 4Q		0.0193* (0.008)		-0.0238 (0.033)
Lagged return, 5Q		0.00464 (0.008)		-0.0349 (0.029)
Lagged return, 6Q		0.0154 (0.009)		-0.0693* (0.028)
Lagged return, 7Q		0.0290*** (0.009)		-0.0220 (0.032)
Lagged return, 8Q		0.0114 (0.009)		-0.00765 (0.029)
Lagged return, 9Q		0.00990 (0.009)		-0.0382 (0.023)
Lagged return, 10Q		0.0269** (0.009)		0.0104 (0.027)
Lagged return, 11Q		0.0214* (0.009)		-0.0128 (0.025)
Lagged return, 12Q		0.0178* (0.008)		-0.0025 (0.026)
Observations	508,723	508,723	508,723	508,723
R-squared	0.0604	0.0588	0.0387	0.0385

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

**Table 3.** Stock-level regressions results, the brokerage data.

On each date, we aggregate all the transactions for each stock to get the total volume (in thousand shares) 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)	(5)	(6)
	By group					
	All	1	2	3	4	5
Loss $\times$ HoldRet	-0.054*** (0.011)	-0.063*** (0.017)	-0.083*** (0.029)	-0.039*** (0.010)	-0.012 (0.009)	-0.011 (0.008)
Gain $\times$ HoldRet	0.029*** (0.004)	0.016** (0.007)	0.018* (0.010)	0.027*** (0.008)	0.056*** (0.011)	0.066*** (0.017)
Gain	0.004*** (0.001)	0.003 (0.002)	0.005 (0.004)	0.004*** (0.001)	0.004*** (0.002)	0.002* (0.001)
Rank effect	Yes	Yes	Yes	Yes	Yes	Yes
Date FE	Yes	Yes	Yes	Yes	Yes	Yes
Account FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,905,500	431,536	512,277	438,974	322,773	199,940
<i>R</i> -squared	0.047	0.057	0.056	0.046	0.046	0.050

Clustered standard errors in parentheses; \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$

**Table 4.** V-shape selling behavior by group.

This table examines the probability of selling as a function of holding-period returns, based on the daily holdings of investors observed in the brokerage data. Detailed descriptions about the data and the filters used in constructing the data set can be found in Section 4.1. We only consider positions with a holding period less than one month, because it has been documented that the V-shape is strongest for short holding-periods. In addition, to control for the rank effect, we require an investor to hold at least five positions to be included in the analysis. All investors are first sorted into five groups based on their degree of doubling down, measured by the stock return in the most recent month before a buy, averaged across all buys; Group 1 has the lowest degree of doubling down, while Group 5 has the highest. The dependent variable is a dummy variable that equals to one if a stock is sold. Gain (Loss) is a dummy variable indicating a positive (negative) return. HoldRet is the return since purchase. Rank effect is controlled by the inclusion of four dummy variables: Best and Worst, dummy variables that equal to one if the stock has the highest and lowest return in the portfolio; and 2nd Best and 2nd Worst, dummy variables for the second highest and second lowest return. Standard errors are clustered by date and account.