

Ownership, Learning, and Beliefs

Samuel M. Hartzmark* Samuel Hirshman†
Alex Imas‡

October 2019

Abstract

We examine how owning a good affects learning and beliefs about its quality. In a series of studies, we show that after receiving a positive or negative signal about a good, ownership causes more optimistic or pessimistic beliefs, respectively, compared to receiving the same signal about a good that is not owned. This effect on beliefs impacts the valuation gap between owners and non-owners (i.e. the endowment effect), which we show to increase with positive signals and disappear with negative signals. Moreover, we demonstrate that people *overreact* to signals about goods that they own relative to normative benchmarks, but learning is close to Bayesian for goods that are not owned. In exploring the mechanism for this effect, we find that ownership increases attention to recent signals about owned goods, exacerbating over-extrapolation. We demonstrate a similar relationship between ownership and over-extrapolation in survey data about stock market expectations. Our findings provide a microfoundation for models of disagreement that generate volume in asset markets and have implications for any setting with trade and scope for learning.

KEYWORDS: biased beliefs, endowment effect, ownership, attention, behavioral economics, learning, extrapolation

JEL Classifications: D9, D12, C93

*Booth School of Business, University of Chicago: samuel.hartzmark@chicagobooth.edu

†Booth School of Business, University of Chicago: shirshma@chicagobooth.edu

‡Social and Decision Sciences, Carnegie Mellon University: aimas@andrew.cmu.edu

We are grateful to Devin Pope, Collin Raymond, Andrei Shleifer, Kelly Shue, David Solomon, Richard Thaler, George Wu as well as seminar participants at Boston University, Chicago Booth, Carnegie Mellon, Max Planck Institute, University of British Columbia, Washington University in St. Louis, UCSB Theory and Experiments Conference, European Behavioral Economics Meeting, for their helpful suggestions.

I. Introduction

Does ownership affect learning and beliefs about the owned good? While prior work has studied how ownership influences attitudes towards a good, the focus has largely been on how the initial stage of merely owning a good affects valuation and preferences. For example, in the influential demonstration of the endowment effect, Kahneman, Knetsch, and Thaler (1990) show that randomly assigning someone ownership of a good increases their valuation of it. However, many important economic contexts involve periods of learning about both goods that are owned and those that are not, with people making decisions after receiving information and updating their beliefs.

In this paper, we examine how owning a good affects learning about its value and quality. In a series of experimental studies, we show that ownership leads to stronger reactions to both positive and negative signals about a good. Compared to goods that they do not own, people become more optimistic after seeing positive signals and more pessimistic after seeing negative signals. This pattern has implications for how ownership affects valuations: if owners are more pessimistic than non-owners after observing negative information about a good, and more optimistic after observing positive information, then the valuation gap between owners and non-owners (i.e. the endowment effect) should expand after good news and shrink after bad news. We demonstrate that this is indeed the case. In exploring the mechanism for the differential in learning, we show that belief responses to new information are more likely to deviate from the Bayesian benchmark for owned than non-owned goods: belief-updating is closer to Bayesian when learning about non-owned goods. This effect appears to operate through the channel of attention: ownership leads people to focus on, and as a result, overweight the informativeness of recent signals about a good’s fundamental quality. Finally, we document a similar relationship between ownership and over-extrapolation in a large field survey on stock market expectations.

We begin our investigation by constructing a setting where ownership can be as-if exogenously assigned, beliefs can be cleanly elicited and a normative benchmark for “correct” learning can be established as well as reasonably attained. To do this, we employ a controlled laboratory experiment where people chose to buy any three of six ex-ante identical goods and report beliefs about their quality. Participants know that each good has a good-specific

probability of a price increase in each period, which we refer to as its fundamental quality. Specifically, in each period t a good i has a constant probability s^i of increasing in price and a constant probability $1 - s^i$ of decreasing in price.¹ Because s^i does not change across periods, a price increase (decrease) is a positive (negative) signal about good i 's quality. Participants observe 15 periods of price movements and are paid based on the final price of the goods they own. In each period t , we elicit beliefs \hat{s}_t^i about the probability of a price increase s^i for each good i – both those that they own and those that they do not – with truthful reporting incentivized. Since participants are not given information about the goods' quality before making their allocation decisions, the choice of which goods to own is as-if random.² Thus, ownership can be thought of as exogenous to any omitted variable related to differences in preferences, skill or knowledge. Additionally, this is a fairly simple learning environment since in each period the total number of price increases and decreases is a sufficient statistic for forming a Bayesian posterior.

Different models of learning make distinct predictions in this setting. A rational Bayesian agent would report beliefs \hat{s}_t^i that depend only on the observed price paths and do not vary by ownership because price movements are equally informative signals for goods that are owned and not owned.³ Behavioral models of learning make predictions that systematically deviate from this Bayesian benchmark. Models of motivated beliefs such as the optimal expectations framework of Brunnermeier and Parker (2005) predict that people should be more optimistic about goods they own relative to goods they do not, because they derive utility from believing that a good they own has a higher fundamental quality. The misattribution model of Bushong and Gagnon-Bartsch (2019) predicts over-reaction to signals for owned goods, but with a stronger effect for negative signals which leads to greater overall pessimism for owned versus

¹Prior work has used an asset market with a similar structure to study the disposition effect in a controlled environment (Fischbacher, Hoffmann, and Schudy, 2017). We also run a version of the study where people chose to buy one of two ex-ante identical goods. The same pattern of results is obtained.

²In a separate treatment, we show that our results do not depend on whether participants actively choose the goods or are randomly endowed with them.

³Note that in our environment, models of rational inattention similarly predict no differences in learning about owned versus non-owned goods (Martin, 2017; Caplin and Dean, 2015; Mackowiak, Matejka, and Wiederholt, 2018). Because beliefs are equally incentivized for both types of goods, there are no instrumental motives to pay more attention to one type over the other. Also, findings about heterogeneity in learning based on fixed characteristics, such as IQ (DAcunto et al., 2019), life experience (Malmendier and Nagel, 2015), or socioeconomic status (Das, Kuhnen, and Nagel, 2017) predict no difference as these characteristics are balanced across ownership conditions and fixed differences can be controlled for.

non-owned goods.⁴

Models of over-extrapolation predict that a person becomes too pessimistic relative to a normative benchmark after observing negative signals and too optimistic after observing positive signals. In theories of diagnostic expectations, the representativeness heuristic (Kahneman and Tversky, 1972) leads individuals to overweight the possibility of states that become objectively more likely after observing a given signal (Bordalo, Gennaioli, and Shleifer, 2018, 2017).⁵ Though straightforward applications of these models do not predict an ownership effect, if ownership channels attention to make recent signals more salient for those goods, then this will exacerbate over-extrapolation for owned goods.

We find that positive signals (price increases) lead to greater optimism, higher \hat{s}^i , about goods that are owned relative to goods that are not. The opposite pattern emerges in response to negative signals (price decreases), we find that people are more pessimistic about goods that they own compared to those that they do not. The ordering holds under any prior that does not condition on ownership and cannot be explained by fixed subject characteristics. These results are not consistent with rational models, which predicts no ownership effects, nor are they consistent with models of motivated beliefs (Brunnermeier and Parker, 2005; Kunda, 1990), which predicts uniform optimism, or models of misattribution (Bushong and Gagnon-Bartsch, 2019), which predict average pessimism.

This relationship between ownership and beliefs has predictable implications for valuation and behavior when owners and non-owners receive information about the good. Specifically, the initial valuation gap between owners and non-owners – the endowment effect – should increase with positive signals as owners become more optimistic than non-owners, and shrink with negative signals as they become more pessimistic. To test these predictions, we endowed participants with one of two goods – battery powerbanks A or B. The endowment effect predicts an initial valuation gap between owners and non-owners: the minimum that a person

⁴In this model, loss averse individuals with reference-dependent preferences misattribute sensations from gain-loss utility to the quality of the good, resulting in a relatively greater overreaction to negative signals about goods that are owned compared to those that are not.

⁵This mechanism is in the spirit of Griffin and Tversky (1992) and Massey and Wu (2005), who predict a similar pattern of over-extrapolation as a result of individuals placing more weight on recent signals relative to accumulated evidence. Griffin and Tversky (1992) reconcile the evidence on conservatism (Edwards, 1968) and representativeness (Kahneman and Tversky, 1972) in belief updating. They argue that people overweight the strength of evidence (e.g. magnitude of recent signal) relative to the weight of evidence (e.g. the prior). Massey and Wu (2005) find support for this framework in a judgment task where participants are asked to identify the state of a system based on the diagnosticity of a signal and their prior.

endowed with powerbank A is willing to accept for it is larger than the maximum a person not endowed with the powerbank is willing to pay for it. We examine how this gap varies with signals about the goods by providing owners and non-owners with real Amazon ratings about each. After establishing an initial endowment effect without the signals, we show that seeing the same positive ratings leads to a substantial increase in the valuation gap between owners and non-owners. In contrast, identical negative ratings lead to a disappearance of the gap.

Next, we examine how ownership influences learning relative to a normative benchmark. For this, we return to the first paradigm where elicited beliefs can be compared to Bayesian learning. We find that belief errors relative to a Bayesian benchmark for non-owned goods are not significantly correlated with signals, suggesting near-Bayesian learning from information about goods that are not owned. In contrast, we find that belief errors have a strong, positive correlation with signals about owned goods. This indicates that, relative to a Bayesian benchmark, individuals over-react to information associated with owned goods. While Bayesian updating predicts that beliefs should be independent of how the signals are ordered, we find that this over-reaction is associated with greater extrapolation from recent signals about goods that are owned. The increased over-extrapolation for owned goods is robust to a host of possible normative benchmarks – including priors that vary with ownership and cumulative signals – as well as benchmarks that do not require distributional assumptions. Recent signals appear to play a larger role in explaining beliefs for owned good than non-owned goods, providing further evidence that the results reflect an over-reaction to information for owned goods.

In exploring the mechanism behind this ownership effect, we find evidence consistent with ownership channeling attention to recent signals, which exacerbates over-extrapolation via the mechanism proposed by Bordalo, Gennaioli, and Shleifer (2017). Specifically, we run an additional treatment to exogenously shift attention to goods that are not owned and compare learning to our baseline condition. In the treatment, participants' beliefs are elicited only for goods that they *do not* own. We find that exogenously manipulating attention in this manner leads to similar over-extrapolation for non-owned goods as we observed for owned goods in the baseline condition. This implicates attention as the driver of the over-extrapolation associated with ownership.

Lastly, we use the Michigan Survey of Consumers to examine the relationship between asset ownership and extrapolation in field data. The survey elicits both forecasts of future stock market returns and whether the individual owns equities. Combing these responses with data on prior market performance allows us to test whether those who own stocks extrapolate more from past market returns compared to those who do not own stocks. We find that even after controlling for baseline optimism and a host of demographic characteristics, those who own assets extrapolate about *twice* as much from prior returns compared to those who do not. These findings are consistent with the pattern observed in the experimental data, but should be viewed as complimentary to those results given the potential for omitted variable bias and selection concerns present in the field setting.

Since ownership is a fundamental aspect of economic interactions, particularly those involving trade, its influence on learning and beliefs likely has a profound influence in a wide-range of settings. Almost any economic decision relating to durable goods with the possibility of future resale has an aspect of learning from new information when forming beliefs about its underlying value or forecasting its future price. Most business settings involve buying certain goods and not others, followed by evaluating and making business decisions based on these expectations. In settings where agents can decide what signals consumers receive, such as in advertising, the strategy and structure of a marketing campaign will differ depending on whether consumers already own a product.

Our results have implications for the measurement of psychological frictions from data in settings that involve ownership. Economic analysis generally assumes that owners and non-owners form beliefs using the same process, explaining differences in behavior through preferences. Our results suggest that ignoring the influence of ownership on the learning process can lead to erroneous conclusions. For example, one of the most well-documented behavioral anomalies is the disposition effect, where people are more prone to sell a good after a gain than a loss.⁶ Our results suggest that studies of the effect likely understate the psychological frictions stemming from preferences by ignoring the influence of ownership on

⁶Beginning with the initial discussion of the phenomenon (Shefrin and Statman, 1985) and demonstration in a large brokerage data set (Odean, 1998), the disposition effect has been replicated in a variety of settings (e.g. equity and housing markets) and with different types of market participants (e.g. retail and day traders) (Kaustia, 2010).

beliefs.⁷ More generally, ascribing a result to preferences rather than beliefs may lead to different conclusions about the underlying mechanism, which in turn can lead to different policy prescriptions.⁸

Our results also have significant implications for understanding behavior in financial markets. A well-known puzzle in finance is that standard models predict only a small fraction of the trade volume observed in real financial markets. Models of disagreement, where agents disagree about the value of an asset given the same information, are the dominant explanation for this puzzle (see Hong and Stein (2007) for a survey of this literature), but the mechanism for *why* such agents have different beliefs is not well understood. Our findings provide a potential microfoundation for such heterogeneity in beliefs: if owning an asset systematically changes the way that an agent updates to information compared to an agent who does own the asset, then the two will disagree about its value despite seeing the same signals.

Finally, our results contribute to the large literature on the influence of attention on expectation formation. While it is often assumed that more attention improves decision quality through more accurate beliefs (see Gabaix, 2017, for review), theoretically this need not be the case. Notably, Dawes (1979) and Dawes, Faust, and Meehl (1989) argue that attention can have a ‘more is less’ effect on judgment if greater attention is combined with an incorrect mental model of the problem. There, the authors conjecture that more attention leads forecasters to overweight features of the decision-problem relative to the normative benchmark.⁹ While argued on theoretical grounds, we believe that ours is the first piece of direct empirical evidence for such a ‘more is less’ effect of attention.

The paper proceeds as follows. Section II describes the experimental paradigm used to explore the influence of ownership on learning and documents the basic effect. Section III demonstrates the implications of differential learning on the valuation gap between owners and non-owners. Section IV presents result on learning relative to a normative benchmark

⁷A belief in mean reversion has been offered as a potential explanation for this effect: people hold on to losers and sell winners because they believe that the former will go back up and the latter will go back down. While many studies have proposed preference-based mechanisms such as realization utility for the disposition effect (Barberis and Xiong, 2012), belief in mean reversion has not been ruled out – largely due to a lack of data on beliefs in trading contexts (Barber and Odean, 2013).

⁸For example, Bohren et al. (2019) argue that wrongly ascribing discrimination to preferences rather than beliefs can lead to vastly different implications for policy.

⁹Importantly, in settings where attention leads people to shift away from heuristics towards more deliberative processing, or where individuals have the correct mental models but cognitive costs lead them to form noisy expectations, greater attention is likely to lead to learning closer to normative benchmarks.

and explores mechanism. Section V provides complimentary evidence from survey field data. Section VI concludes.

II. Ownership and Beliefs

In order to examine the impact of ownership on beliefs we designed a simple experiment where we can measure learning and beliefs while varying ownership status. We employ an experimental asset market on a large crowdsourcing platform called Amazon Mechanical Turk where participants purchase goods, view prices of those goods over a sequence of periods, and report their beliefs about the fundamental quality of the goods. The market consists of six goods with equal starting prices of 100 experimental points per share. Participants were endowed with 2000 experimental points (500 points = 50 cents) and asked to spend the entire sum on shares of three of the six goods.

In each round t , a good $i \in \{1, \dots, 6\}$, has a fixed probability of a price increase, s^i , which represents its fundamental quality. This good-specific quality remains constant throughout the experiment. In each round, the price of the good either increases or decreases by a constant amount; a price increase is always 6% and a price decrease is always 5%. The number of prior price increases and decreases is provided to the participants in every round. Since a price increase is more likely to be observed if a good has a higher fundamental quality—a good with $s^i = .7$ has a higher probability of experiencing a price increase in any period t than a good with $s^i = .4$ —price changes can be thought of as signals about a good’s fundamentals. Throughout the analyses, we use percent returns as our measure of prior signals since they are isomorphic to price changes in our context; a good that had an initial price of 100 and a current price of 120 is classified as having a return of 20. At the end of the study, participants earned a bonus based on what their portfolio (i.e. owned goods) was worth in addition to a base fee of \$1.20.

While participants were told that each good had a fixed s^i , they were not informed of the actual quality for any of the goods: their task was to infer this quality from the price signals. The key component of our study is the elicitation of beliefs about each good’s fundamentals in each round. We refer to these elicited beliefs as \hat{s}_t^i . When participants decided which three goods to buy, all six of the goods were identical with a price of 100; as a result, ownership

as a function of quality was as-if random. Participants observed price changes and reported their beliefs about each good’s quality value over the course of 15 rounds.¹⁰ Participants were compensated based on the accuracy of their forecasts, receiving a bonus of \$1 if a randomly selected estimate was within plus or minus 5% of the true probability s^i . We chose to use this elicitation procedure as opposed to more complex mechanisms such as versions of the Binarized Scoring Rule (e.g. quadratic scoring rule) due to recent evidence showing that the BSR systematically biases truthful reporting. Danz, Vesterlund, and Wilson (2019) demonstrate that the BSR mechanism leads to systematic conservatism in elicited beliefs, resulting in substantially greater error rates relative to a simpler mechanism that offers little to no information about the specific incentives. The authors argue that simpler mechanisms that incentivize reporting of belief quantiles — such as the one used here — will result in more truthful reporting while imposing fewer cognitive burdens on participants.

This setting represents a simple learning environment for a rational Bayesian agent. Since subjects are evaluating likelihoods, beliefs can be represented using β distributions, which are distributions over probabilities. Define u_t^i as the total number of price increases and d_t^i as the total number of price decreases observed in round t for good i . A Bayesian with an initial prior of $\beta(a, b)$ would update to $(\frac{a+u_t^i}{a+u_t^i+b+d_t^i})$ after receiving signals u_t^i and d_t^i . In turn, the number of prior increases u_t^i and decreases d_t^i — which is available to participants in every round — is a sufficient statistic for calculating the posterior: the order does not matter.¹¹

We follow convention (e.g. Fischbacher, Hoffmann, and Schudy, 2017) in randomly generating the price paths before the experiment. This facilitates between-subject analyses since it allows for comparisons of beliefs by ownership status conditional on seeing the same price paths. We examined two sets of fundamentals s_i ; the probabilities ranged from .1 to .9 with a median of .43 across both.¹² We randomly reorder the price paths to counterbalance the relationship between the order goods are presented to participants and the price paths. There were four different orderings for each set of price paths.

¹⁰We ran a treatment with 20 rounds to ensure that there was not an end of game effect. The number of rounds had no significant effect on our measures of interest. In turn, we present results for both 15 and 20 round treatments together up through round 15. We show that the results are robust to including rounds 16-20 as part of our robustness checks.

¹¹As evidence for the transparency of the learning environment, Section IV demonstrates that participants’ beliefs match this Bayesian benchmark fairly well for non-owned goods.

¹²Price paths are reported in the Appendix.

Because of the structure of the market, a decrease in price is a negative signal about quality while an increase in price is a positive signal. In order to restrict our sample to those who understood the structure of the market, we include participants whose beliefs were positively correlated with prices and significant at the 10% level. Participants also answered control questions to assure they understood that the probabilities of each good going up or down in price were independent in each round, that their reported beliefs did not influence these prices, and that they would purchase and hold three of the six goods. This results in a final sample of 410 out of 598 subjects who completed the survey.¹³

A. Results

Our experimental design allows us to examine the impact of owning a good on learning about its fundamental quality. Since goods are ex-ante identical to the participants, different elicited beliefs \hat{s}^i between purchased and non-purchased goods will be driven purely by the effect of ownership.

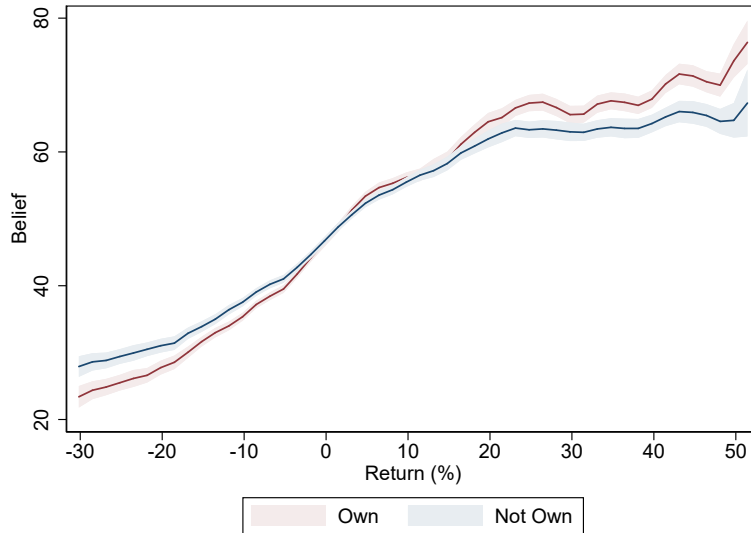


Figure 1. Beliefs by Return. This graph shows a local linear plot of beliefs, \hat{s}_t^i , on returns separately for goods that are owned and not owned. Data include observations with returns from the 5th to the 95th percentile. Shaded area represents the 95% confidence interval.

We begin by comparing the beliefs about various goods at different return levels in Figure 1. The red line shows the average beliefs \hat{s}_t^i associated with goods that are owned for each

¹³The qualitative findings are robust to removing the comprehension filter, but are noisier.

return level. The blue line shows the average beliefs for goods that were not purchased. The shaded areas indicate 95% confidence intervals. The red line has a steeper slope than the blue, consistent with a greater response to a return level for goods that are owned. For lower returns on the left of the graph, the red line is consistently below the blue line. This indicates that, for a given return level, participants are more pessimistic about goods they chose to buy – believing them to be worse than goods they decided not to buy. For higher returns on the right side of the graph the general pattern is reversed. The red line is consistently above the blue line indicating that participants are more optimistic about goods that they purchased compared to goods that they did not.¹⁴

Table I examines this pattern in greater detail. Panel A examines beliefs based on the return level and how they differ depending on ownership. Beliefs are regressed on the return, an *Own* dummy variable equal to one if the good is owned and an interaction of the two variables. The coefficient on *Return* in Column 1 shows that there is a strong positive relationship between good i 's performance and the respective belief \hat{s}_t^i for non-owned goods, which is expected given the structure of the experiment. The coefficient of interest is on the interaction of *Own*Return*, which is positive and significant. This indicates that beliefs about goods that are owned respond to cumulative price signals more than beliefs about goods that are not owned, consistent with the red line being steeper in the prior figure.

In our setting, a rational Bayesian would need to know only the return level and the round to form their expectations. Thus in Column 2 we include return by round fixed effects. This column shows how beliefs about owned positions differ from those that are not owned given any Bayesian benchmark that does not condition on ownership.¹⁵ This also controls for any non-Bayesian benchmark that takes price paths as its input and does not condition on ownership when forming beliefs. If anything, the results are slightly stronger, with a point estimate of 0.116 and a t-statistic of 5.29. To put this number in context, in response to seeing a return increase of 26% (one standard deviation in our sample) a coefficient of 0.116

¹⁴One potential concern with incentivizing both ownership and beliefs is the potential for hedging, which would lead participants to bias their reported beliefs to be more pessimistic about owned goods with positive signals and more optimistic about owned good with negative signals. This is the opposite of the pattern we observe. These concerns are further assuaged with additional treatments reported in Section IV which replicate the general pattern of results while incentivizing only one set of outcomes.

¹⁵This is the same as fixed effects for the number of price increases and decreases in our setting, which are the inputs to posterior beliefs for a Bayesian in our setting. See Section IV for a discussion.

implies that a subject would believe the fundamental quality to be 3% higher for a good that they owned compared to one that she did not own.

People may update their beliefs differently depending on their individual characteristics, for example due to differences in IQ (DAcunto et al., 2019), differences in life experience (Malmendier and Nagel, 2015), or differences in socioeconomic status (Das, Kuhnen, and Nagel, 2017). Column 3 adds subject fixed effects to control for such differences. Results are similar, suggesting that the effect is not driven by individual differences. After removing individual averages, the *same* person is more optimistic for owned goods after receiving positive signals and more pessimistic for owned goods after receiving negative signals, compared to receiving similar signals about goods she does not own.

Next, we explore three intervals of return levels based on Figure 1 - returns below 0, between 0 and 20 and above 20. Beliefs are regressed on dummy variables for a return below 0 and a dummy variable for a return above 20, omitting the range from 0 to 20. The coefficients of interest are dummy variables for these ranges interacted with *Own*, a variable equal to one if the good was purchased. The coefficient on $Own * Return < 0$ is -1.9% and significant, which indicates that owned goods that experienced a negative return are viewed with greater pessimism – expected to have a 1.9% lower fundamental quality \hat{s}_t^i – compared to goods that experienced the same returns but were not owned. The insignificant coefficient on the $Own * (Return[0,20])$ variable indicates that there was not a difference between goods that were owned and not owned in this range of returns. Finally, the significant 5.2% coefficient indicates optimism for goods with returns above 20 relative to those that experienced the same return levels but were not owned. Columns 2 and 3 add round by price fixed effects and subject fixed effects and find similar results after controlling for subject heterogeneity and any Bayesian prior.

Finally, we sought to test whether this pattern was robust to decreasing the number of goods that participants had to keep track of. To do so, we ran a version of the experiment where participants chose to own one of two ex-ante identical goods. All other features of the experiment were kept the same. Table IA.2 in the Appendix presents the results, which follow the same pattern as in the six-good case.

These findings suggest a robust difference in learning in response to cumulative signals

about owned goods compared to those that are not owned. Participants react more to information about goods that they own relative to those that they do not. They are more pessimistic about owned goods that experienced negative signals and more optimistic about owned goods that experienced positive signals, relative to goods that are not owned. This ordering holds using beliefs and belief errors under any rational Bayesian prior.

III. Ownership, learning, and valuation

The learning results documented in Section II have implications for contexts where owners and non-owners have the opportunity to update their valuations of a good in response to information. After observing negative signals, owners will be more pessimistic and assign a lower worth to the good than non-owners (conditional on the initial valuation). In contrast, after observing positive signals, owners will be more optimistic and assign a higher worth to the good than non-owners. Prior work has documented an initial valuation gap between owners and non-owners termed the endowment effect. Kahneman, Knetsch, and Thaler (1990) showed that ownership increases people’s minimum willingness to accept (WTA) to part with the good relative to non-owners’ maximum willingness to pay (WTP) for the same good (see Marzilli Ericson and Fuster (2014a) for review). In this context, we predict that the valence of information will have an asymmetric effect on this initial WTA-WTP gap: the gap will shrink in reaction to negative signals, as owners become more pessimistic than non-owners about the good, and expand in reaction to positive signals, as owners become more optimistic than non-owners.

To test this, we endowed participants with power banks, which are auxiliary batteries designed for charging cell phones remotely.¹⁶ After being endowed with one of two power banks, each participant observed signals about the quality of the power bank they owned and the one that they did not own over the course of five rounds. Signals came in the form of ratings (1 to 5 star ratings) of the power banks taken from Amazon. After observing a rating, we elicited a participant’s WTA for the owned power bank and WTP for the non-owned power bank on a \$0 to \$100 scale. This was done in each of the five rounds. We drew multiple sets

¹⁶We chose powerbanks as they are generic products with substantial heterogeneity in quality. Thus there is scope for significant learning about product quality from product ratings. They are also reasonably priced goods, making it practical to purchase a large number of them to give to participants.

of ratings such that the cumulative signals were better for one power bank than the other in some sets, and vice versa in the other sets. Both endowment and the set of ratings drawn was counterbalanced.

We implemented two version of the paradigm to ensure the results are robust to different methodologies that have been used to in previous endowment effect experiments. The elicitation procedure was identical between the versions; the main difference was how a potential transaction would be carried out. In the “Purchase” paradigm, transactions were carried out in the following manner. After submitting their valuations, participants from Amazon Mechanical Turk (N=100) were randomly selected to be an owner or a non-owner, and one random round of valuations was selected to be the relevant round. If a non-owner’s WTP exceeded the true price of the power bank, which was unknown to the participants, they purchased the good at this true price. The participant received the good and was additionally paid the difference between their endowment and the true price. If their WTP did not exceed the true price, they received their endowment.¹⁷ If an owner’s WTA was below the true price, they received the true price and did not end up with the good. If their WTA exceeded the true price, then they retained the good.¹⁸ Participants were also paid a base fee of \$1 for completing the study.¹⁹

In the “Trade” paradigm, participants (N=199) were asked to imagine that they are in a similar scenario to the one above. Each was told that there would be a possibility to trade between owners and non-owners. Specifically, participants selected to be owners would be paired with participants selected to be non-owners. Their respective WTA and WTP valuations would be randomly chosen as in the “Purchase” paradigm; trade would occur at the non-owner’s WTP if it exceeded the owner’s WTA. If the WTA exceeded the WTP, no trade would occur. Participants were paid \$1 to complete the study.²⁰

¹⁷The amount of the endowment was not specified to participants ex-ante in order to avoid anchoring their valuations; they knew that earnings could never be negative. In practice the endowment was \$25 dollars, which exceeds the true price of the most expensive power bank. Thus if a participant stated a WTP of \$25 or above, they would receive the powerbank and the difference between the true price and \$25.

¹⁸The first study, reported in Section 2, allowed us to generate directional hypotheses for the endowment effect setting. We pre-registered the hypotheses, the methods and the analyses at AsPredicted (<http://aspredicted.org/blind.php?x=yh53yd>).

¹⁹All transactions involving the power banks were done through Amazon.com.

²⁰Numerous demonstrations of the endowment effect have used hypothetical scenarios (e.g. Carmon and Ariely (2000)), including one of the very first demonstrations of the effect (Thaler, 1980). We employed both incentivized and hypothetical versions of the paradigm to demonstrate the robustness of the effect on this dimension as well.

To ensure that our paradigm replicated the standard endowment effect without information, we ran a separate treatment ($N = 99$) that used the same format as the trade setting but did not provide participants with ratings. In turn, the WTA and WTP measures were elicited once. We found a sizable and significant endowment effect. Non-owners had an average WTP of \$28.93 while owners had an average WTA of \$34.47 ($p < .01$). Endowing participants with a good in our setting increased their valuation of it by 19%, which is well within the range of prior demonstrations of the effect (Marzilli Ericson and Fuster, 2014b).

A. Results

Figure 2 graphs the average valuation based on the cumulative average rating in that round. The red line graphs the WTA for the good that is owned and the blue line graphs the WTP for the good that is not owned. The gray bars throughout the graph represent the endowment effect in the absence of any information (\$5.54). Panel A reports pooled results from the Purchase and Trade studies, Panel B reports results from the Purchase study only, and Panel C reports results from the Trade study only. The pattern is similar across both studies, so we focus our discussion on on the pooled results in Panel A.

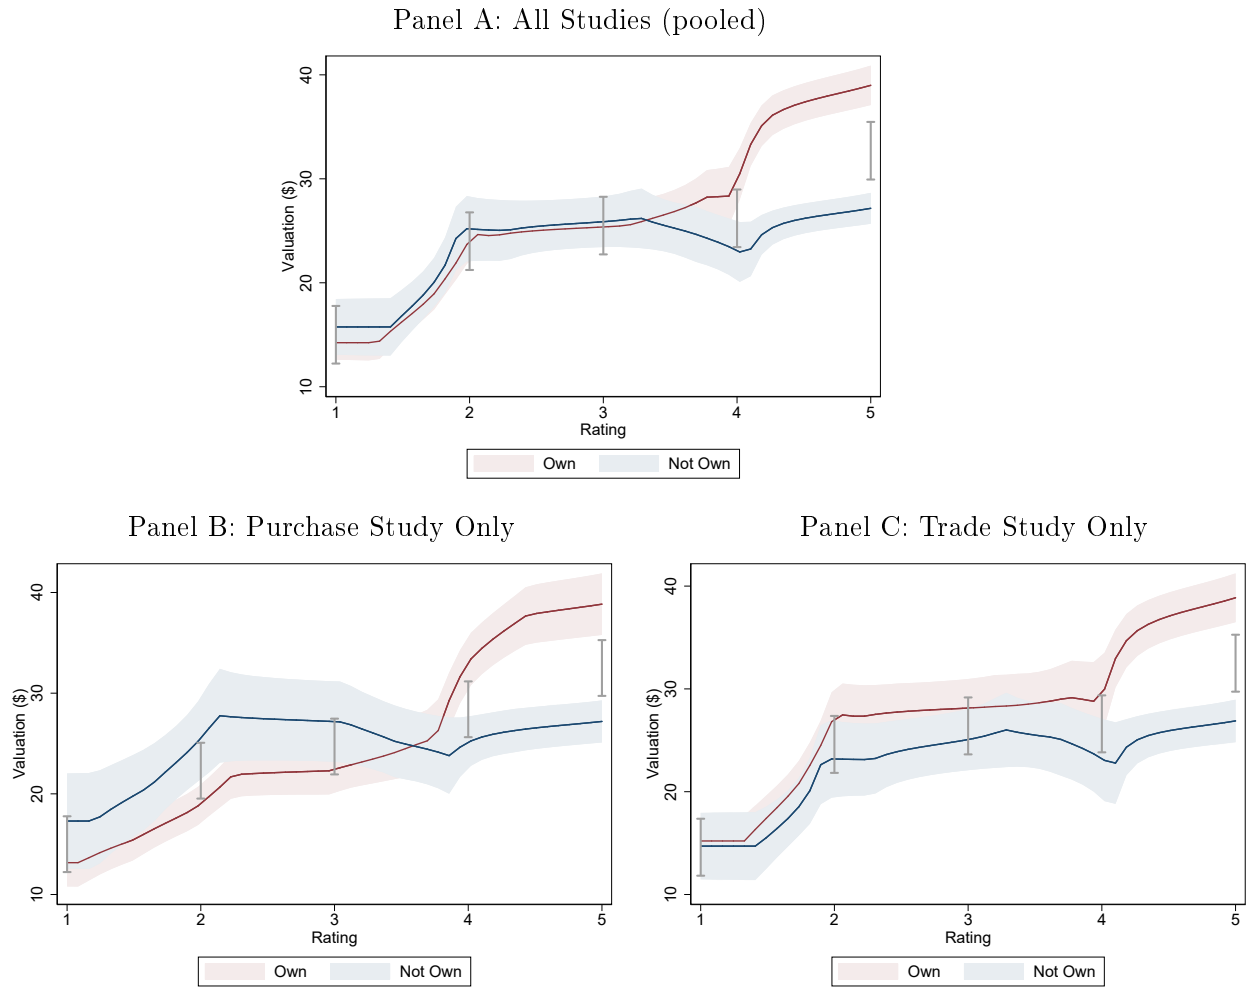


Figure 2. Valuation by Rating in Endowment Effect. The figure shows a local linear plot of willingness to accept for owned goods and willingness to pay for not owned goods based on its average cumulative rating. Panel A includes both treatments, Panel B includes only the Purchase study and Panel C includes only the Trade study. The gray bars represent \$5.54, the magnitude of the endowment effect without information. Shaded area represents the 95% confidence intervals.

The figures show that ownership influences valuations in line with its influence on beliefs. Specifically, the valuation gap increases with positive signals to the right of the figure, and decreases with negative signals to the left of the figure. A reasonable initial benchmark for quality is likely around 4 stars given a participant's experience on Amazon and the average rating in our experiment (3.7 stars). Consistent with this, when participants saw a rating in line with their expectations (roughly 4 out of 5 stars), the gap between the red and blue line is similar to the valuation gap in the absence of information (the gray bars). Five star

ratings can be considered as positive signals, while ratings below 4 stars can be considered as negative signals. When participants observed positive signals of around 5 stars to the right of the figure, the valuation gap increases substantially, roughly *doubling* in magnitude. The opposite pattern is observed to the left of the figure in response to negative signals: the gap between the valuations disappears, and even directionally reverses.²¹

Table II examines the pattern more formally. It reports coefficients from the following regression:

$$Value_{it} = \alpha + \beta_1 Own * Rating_{it} + \beta_2 Rating_{it} + \beta_3 Own_{it} \quad (1)$$

Rating is measured as the average rating observed for good i by round t in Panel A and the most recent rating in Panel B. To make the coefficients easier to interpret, $Rating_{it}$ is normalized to 3 stars.²² Columns 1 through 3 report the results with no controls for the version of the study. Examining Column 3, the coefficient on Own is 3.62 with a t-statistic of 2.98. This corresponds to a \$3.62 increase in WTP when the average rating is three-stars.

The coefficient of interest for our experiment is the interaction term, which we find to be robustly positive, consistent with ownership influencing valuations in line with our predictions. The interaction term in Column 3 is \$3.78 with a t-statistic of 3.39. This implies that a one-star decrease in the good's rating decreases the valuation gap by \$3.78. Column 4 adds subject fixed effects to control for heterogeneity in valuation of power banks. The coefficient on the interaction term decreases, though remains a significant \$1.41 with a t-statistic of 2.03. The next column adds round fixed effects and finds a coefficient on the interaction term of \$3.79 with a t-statistic of 3.39. While the analysis examines responses to the average rating, subjects may weight ratings differently than a simple arithmetic average. Column 6 includes a round by rating dummy, which removes the average value for any sequence of ratings. After doing so the interaction coefficient is \$3.77 with a t-statistic of 3.37. The final column includes

²¹The elimination of the valuation gap in response to negative signals is related to the findings of Lerner, Small, and Loewenstein (2004), who show that inducing negative emotions prior to trade similarly eliminates the endowment effect. If the induced emotions spillover to the valuation process, as Lerner, Small, and Loewenstein (2004) argue, then they can be interpreted as generating a negative signal about the good. In turn, these results can be interpreted as a complimentary demonstration of the effect presented here.

²²Thus a five-star rating has a value of $Rating_{it} = 2$ and a one-star rating has a value of $Rating_{it} = -2$. Centering at 3-stars does not change the coefficients on $Own * Rating_{it}$ or $Rating_{it}$. It leads Own_{it} to represent the difference in value between owned and non-owned positions at a 3-star rating.

individual fixed effects and round by rating fixed effects; the interaction term is similar to the case of subject fixed effects alone. Panel B repeats the analysis using the most recent rating rather than the cumulative average rating. The interaction term is positive and significant across all specifications, with t-statistics above 3.

In addition to exploring valuation effects, this setting also demonstrates the robustness of the results obtained using the paradigm described in the previous section. While we attempted to make that experiment as transparent as possible, one may be concerned that participants were confused about trading financial assets, concepts of probabilities or the abstract nature of the setting. The endowment setting involves physical goods rather than abstract assets and participants reported valuations rather than probabilities. The setting has been utilized in so many experiments in part because it is viewed as intuitive and straightforward. Thus, the fact that we find analogous results in a classic endowment effect setting should assuage concerns that the learning results documented in our initial experiment were driven by artifacts of the design. Additionally, while the first experiment allowed us to directly demonstrate the effect of ownership on learning and beliefs, the second experiment shows how this differential learning affects behavior. This suggests that while preferences are likely at the root of the initial valuation gap between owners and non-owners (as has been argued by others (Kahneman, Knetsch, and Thaler, 1991)), beliefs play a significant role in how the gap evolves in response to new information.

IV. Exploring mechanism

The previous two sections demonstrate differential learning as a function of ownership. People who own a good are more optimistic (pessimistic) about its quality after seeing positive (negative) signals about it compared to people who do not own it, and this is reflected in their reported valuations. In this next section, we explore whether owners or non-owners are less ‘correct’ when learning about the quality of a good (compared to a Bayesian benchmark), and further study the mechanism behind the effect. To do this, we return to the stylized setting used in Section II. We find that individuals appear to *overreact* to information about goods that they own compared to goods that they do not, with owners being more likely to over-extrapolate from recent signals. Further, our evidence points to channeled attention as

the driver for this difference in over-extrapolation.

A. Ownership and Belief Errors

A.1. Establishing a Prior

To construct a Bayesian benchmark it is necessary to establish a reasonable prior for participants in our setting. To do so we ran a separate experiment where subjects were shown the instructions in the study, but did not participate in it. Instead, they were asked to rank bins from 0% to 20%, 20% to 40%, 40% to 60%, 60% to 80% and 80% to 100% based on which they thought was more likely to contain the true s_i of a randomly chosen good. This ranking was incentivized based on the true distributions of our goods.




Figure 3. Prior Beliefs about s^i . These figures show the frequency of beliefs about s^i made after seeing the main experiment’s instructions. Each bar represents the frequency of beliefs made in the indicated 20% bin. To the left is the proportion of subjects ranking a given bin as the most likely. To the right is the proportion of subjects ranking a bin least likely or 2nd least likely.

Figure 3 illustrates the results. To the left is the histogram of the bin ranked most likely; we find that the middle 40% to 60% bin is most commonly rated the most likely. To the right is the fraction ranked worst or 2nd worst; here, the most extreme positive and negative bins are most likely to receive these ratings. This suggests that the average participant in our experiment had a prior centered around 50% with more mass in the center and the lowest mass in the tails. Consistent with this, in round 1 of the baseline condition the average belief was 48%. Based on this evidence, we consider symmetric β distributions as they represent distributions over probabilities. A $\beta(a, b)$ distribution is symmetric when $a = b$. For succinctness, we refer

to such a symmetric distribution using a single parameter, e.g. a $\beta(3, 3)$ is denoted as $\beta(3)$.

While these results speak to the shape of the distribution, they do not capture the mass concentrated at 50% relative to the tails. We hone in on this feature by examining how participants update their beliefs to new information in the first few rounds of the experiment. We look at the first two rounds since this is when priors should have the strongest influence. For a given belief, we calculate the implied parameters for a Bayesian agent who had seen a given price signal and had that belief. For example, the average belief in round 1 after viewing a price increase was 57.5%. A Bayesian with a $\beta(2.83)$ prior observing a price increase would report a belief of 57.5%.²³ The implied parameters in the first two rounds are 2.83, 2.24, 2.96 and 2.46, which have an average of 2.62.²⁴ To provide further evidence supporting this parameterization, we ran an additional treatment where we told participants that s^i was drawn from a $\beta(2.62)$ distribution. The results (reported in the following section) are similar with or without this information, which suggests that participants in the baseline treatment have a prior similar to a $\beta(2.62)$ distribution.

Based on this evidence, we use a Bayesian benchmark based on a $\beta(2.62)$ prior.²⁵ This distribution is centered at 50% with more mass in the middle, though it is relatively diffuse. Before viewing any price movements, such an agent believes there is about a 34% chance s^i is between 40% and 60%, 25.8% that it is in the 20% to 40% or 40% to 80% bin and roughly 7.2% for each of the two extreme bins.

A.2. Results

Figure 4 graphs the belief errors relative to a Bayesian with a $\beta(2.62)$ prior by return level. The blue line, representing goods that are not owned, is relatively flat. This indicates that the learning of a rational Bayesian with a $\beta(2.62)$ prior is similar to the learning of an average subject in our experiment for goods that she *does not own*. On the other hand, the

²³Starting with a symmetric prior of $\beta(2.83)$, the posterior mean belief is $\frac{2.83+u_i^i}{2.83*2+u_i^i+d_i^i}$. If such an agent observes one round with one price increase, the mean Bayesian posterior is equal to $\frac{2.83+1}{2.83*2+1+0} = 0.575$.

²⁴Based on mean beliefs \hat{s}_t^i after a single decrease in round 1 of 40.1, after 2 increases in round 2 of 62.6, and of two decreases in round 2 of 35.6. After a price increase and decrease in round 2 the average belief is 49.99, consistent with any symmetric β prior centered at 50%.

²⁵In the Appendix we explore alternative benchmarks including using priors $\beta(2)$, $\beta(2.5)$, $\beta(3)$, $\beta(3.5)$, a simulation and ex-post forecast errors and Section IV.B.1 explores overreaction relative to a variety of alternative priors. While different benchmarks yield slightly different results, the main patterns discussed in this section are robust to any of these specifications.

red line – which represents belief errors associated with goods that participants own – has a positive slope. This indicates that participants update to a greater extent than a Bayesian agent for goods that they own, consistent with an *overreaction* to signals about owned goods.




Figure 4. Belief Error by Price. This graph shows the belief error relative to a benchmark of a Bayesian agent with a β (2.62) prior based on whether a good is owned as a function of its return. Data include observations with returns from the 5th to the 95th percentile. Shaded area represents the 95% confidence interval.

Table I repeats the regression analysis from the previous subsection using the belief error – the participant’s reported belief minus the belief of a Bayesian agent with a β (2.62) prior – as the dependent variable. Column 1 of Panel A shows the regression without controls. The coefficient on *Return* is roughly 0, which indicates little difference in learning between participants in our experiment and a Bayesian agent for goods that are not owned. In contrast, the coefficient on *Own*Return* is nearly 0.1 and significant at the 1% level. This implies that in response to a positive signal about an owned good, participants increase their beliefs by 20% more than both a rational Bayesian and the response to the same signal about a non-owned good.²⁶ Column 2 includes an individual fixed effect and shows similar results.²⁷

²⁶In response to a positive signal (6% return) about a non-owned good, participants increase their stated quality by 3% (based on the coefficient on *Return* of 0.5), which is consistent with rational Bayesian updating. In contrast, the 0.1 coefficient on *Own*Return* implies that in response to the same positive signal, participants increase their beliefs about quality by 3.6% – 20% higher than a Bayesian observing the same information.

²⁷We do not add a round by return fixed effect as this controls for any Bayesian prior and thus does not add information when explicitly including a benchmark.

While the data suggests that participants enter the experiment with an average prior of $\beta(2.62)$, it is possible that this is not the case. We test the robustness of our results to a host of different priors in the Internet appendix and demonstrate the specific choice of prior does not drive the results. We also conducted an additional treatment which repeated the experiment while providing participants information on the underlying distribution of s_i 's – corresponding approximately to a $\beta(2.62)$.²⁸ If participants in this treatment behave similarly to agents who are not given this information, it provides further support for using the Bayesian benchmark of $\beta(2.62)$ for participants in the baseline treatment.

Columns 3 through 5 of Table III shows that this is indeed the case. These columns include data from the baseline experiment as well as data from the new treatment. The regressions introduce a variable *Treat* which is equal to one if the data are from the new treatment with prior information. This variable is included in the regression along with interactions of *Treat* with the original regressors.²⁹ The coefficients on the interaction variable are small and do not approach significance in any of the specifications. These results indicate that after being provided with information that the the distribution of s_i was drawn from a $\beta(2.62)$, participants behave similarly as when this information is not provided.

The evidence suggests that learning about goods that are not owned is closer to the Bayesian benchmark. Thus, the observed pessimism after negative signals and optimism after positive signals about owned goods can be interpreted as an *overreaction* to the signals.

B. Mechanism

These results demonstrate that ownership leads to incorrect learning and beliefs about a good's fundamental quality. This pattern is not consistent with rational Bayesian learning, which predicts no differences by ownership status. It is also not consistent with behavioral models of motivated beliefs (Brunnermeier and Parker, 2005; Kunda, 1990), which predict overall optimism, nor models of misattribution (Bushong and Gagnon-Bartsch, 2019), which

²⁸Specifically, we told participants that: “Each investment’s underlying probability is drawn from the same distribution with the following property. If you were to randomly draw 100 probabilities from this distribution you would expect to find 7 investments with a likelihood of a price increase between 0 and 20% per round, 26 investments with a likelihood between 20 and 40%, 34 Investments between 40 and 60%, 26 Investments between 60 and 80%, and 7 Investments between 80 and 100%.”

²⁹The Internet Appendix repeats the baseline analysis using only data from the new treatment. Results are qualitatively unchanged.

predict overall pessimism, rather than the symmetric over-reaction we observe in the data.³⁰ Moreover, models of rational inattention are also unable to explain our findings because reported beliefs are incentivized in the same way for owned and non-owned goods. Finally, since our results are robust to the inclusion of subject fixed effects, the learning pattern cannot be explained by heterogeneity based on fixed participant characteristics.

We now consider a potential mechanism which posits that, rather than affecting how information is interpreted (as models of motivated beliefs and misattribution predict), ownership channels attention towards signals associated with owned goods. Work in cognitive psychology has shown that attention has an intimate relationship with value-relevant information (Smith and Krajbich, 2019, 2018; Enax, Krajbich, and Weber, 2016); in turn, a greater share of attention is likely to be allocated towards signals associated with payoff-relevant assets – owned goods. Learning models of diagnostic expectations (Bordalo, Gennaioli, and Shleifer, 2017, 2018) predict over-extrapolation due to individuals’ use of the representativeness heuristic, which leads them to overweight the degree to which a recent positive (negative) signal is diagnostic of a high (low) underlying value.³¹ Ownership can potentially exacerbate this effect by directing attention towards recent signals, increasing their salience and weight in the judgment process. This would lead to greater belief distortions and over-extrapolation about owned goods relative to non-owned goods.

Note that greater attention exacerbating belief errors runs counter to the standard predictions of the attention literature, which implicitly assumes that more attention will improve decision quality (see Gabaix (2017), for review). Rather, it is consistent with the ‘more is less’ effect discussed in Dawes (1979) and Dawes, Faust, and Meehl (1989). There, the authors conjecture that greater attention may lead forecasters overweight features of the decision-problem relative to the normative benchmark. In our setting this would correspond to ownership directing greater attention to the arrival of new information, which leads to this information being overweighted in the judgment process.

³⁰In our setting, ownership would lead to overreaction if people derive gain-loss utility from signals, as opposed to payoffs, associated with goods they own versus those they do not. Loss aversion would generate greater pessimism for owned goods relative to non-owned goods.

³¹Griffin and Tversky (1992) predict a similar over-extrapolation from recent signals without specifying the underlying heuristic process.

B.1. Ownership and Extrapolation

We first provide evidence for the prediction that there should be greater extrapolation of recent signals for positions that are owned compared to those that are not owned. To do so, we regress beliefs on *Price Increase*, a dummy variable equal to one if there was a positive price signal in round t and zero if there was a negative price signal, the *Own* dummy variable and an interaction between the two. The coefficient on *Own*Price Increase* corresponds to how much more or less agents respond to a price increase for positions they own compared to positions they do not. The coefficient on *Own* represents how much more or less agents respond to a price decrease for positions they own compared to those they do not.

Table IV Panel A presents the results which show that agents appear to extrapolate more from both recent price increases and decreases for positions that they own. Column 1 examines raw beliefs without controls. The interaction of *Price Increase* with *Own* has a coefficient of 4.98 and is significant at the 1% level, which indicates that individuals update their beliefs by 5% more after seeing a positive signal about an owned good compared to non-owned good. The coefficient on *Own* is negative, which indicates there is also a larger negative reaction to price decreases for positions that are owned.

Without further controls it is unclear whether the difference in updating based on ownership is due to differences in extrapolation, or whether it simply reflects differential updating based on a given information set. As noted in Section II, the posterior mean belief \hat{s}_t^i in round t associated with good i for a Bayesian agent with symmetric prior $\beta(a)$ is $(\frac{a+u_t^i}{2a+u_t^i+d_t^i})$, where u_t^i is the total number of price increases and d_t^i is the total number of price decreases observed by round t . This underscores that in our simple learning setting, the ordering of signals does not matter for a rational Bayesian as the constant prior parameter (a) and the number of positive and negative signals is sufficient to calculate the posterior.

In contrast, the order of price signals does matter for an agent who over-extrapolates recent signals. We use the following expression of the mean posterior belief \hat{s}_t^i to estimate the degree of extrapolation for good i in round t :

$$\hat{s}_t^i = \frac{a + u_t^i}{2a + u_t^i + d_t^i} + \nu * Z_t^i \quad (2)$$

where $Z_t^i = 1$ if good i experienced a price increase and $Z_t^i = -1$ if it experienced a price decrease in round t . The first term captures the mean posterior belief of a Bayesian agent with prior $\beta(a)$. The parameter ν in the second term captures the extent of over or under-extrapolation from the recent price signal. A $\nu > 0$ corresponds to over-extrapolation while a $\nu < 0$ corresponds to under-extrapolation. The expression reduces to Bayesian updating when $\nu = 0$.

To estimate the value of ν we measure belief errors relative to a Bayesian benchmark and use them as the dependent variable in the extrapolation regression. If an agent updates as a Bayesian, the difference between \hat{s}_t^i and the Bayesian benchmark should not be influenced by recent price changes as the benchmark accounts for updating with regard to that information. If an agent over- or under-extrapolates from recent signals, we expect recent signals to have significant explanatory power for \hat{s}_t^i even after controlling for the benchmark.

Column 2 presents belief errors relative to such benchmarks, capturing the degree of over- or under-extrapolation relative to a Bayesian with an initial prior of β (2.62). The coefficient on *Price Increase* is an insignificant -1.145 which indicates there is no significant over- or under-extrapolation from price increases for non-owned positions.³² The point estimate on *Own*Price Increase* decreases slightly to 3.77 and the point estimate on *Own* becomes slightly less negative moving to -2.35, but both are significant at the 1% level. This suggests that the majority of the effect captured in Column 1 represents *over-extrapolation* from recent signals for owned positions rather than Bayesian updating. Equation 1 above imposes a uniform ν to price increases and decreases, which means the degree of extrapolation from positive signals and negative signals is uniform. The regression specification used in Panel A allows for differential extrapolation from positive and negative signals. The coefficient on *Own*Price Increase* can be interpreted as the ν in response to price increases and the coefficient on *Own* can be interpreted as the ν in response to price decreases.³³

These results illustrate over-extrapolation relative to any prior that does not condition on ownership. However, the regressions may be capturing the general difference in beliefs for owned versus not owned positions rather than differential extrapolation from the most

³²The constant of the regression is -0.988 with a t-statistic of -1.42, which indicates there is also no extrapolation of price decreases for non-owned positions.

³³The analysis imposing a symmetric ν is conducted in the Internet Appendix. The results are materially similar.

recent signal. To test for such a possibility, we allow for benchmarks where a varies by ownership and by round. We do so in two ways. First, we repeat the technique used to calibrate the β (2.62) (discussed in Section IV.A.1), but do so separately for owned and not owned positions. We refer to the belief relative to this benchmark as β (*Own*) *Error*. Second, we use the average belief reported for a given price signal and ownership status to identify the implied value of a in a given round $t - 1$.³⁴ Using this estimate of a_{t-1} , we can calculate the posterior for a Bayesian who observes the realized price signal in the next round. We term this benchmark β (*Own Round*) *Error*.

Columns 3 and 4 present belief errors relative to these benchmarks and provide further evidence of over-extrapolation for owned goods. In Panel A, the coefficients on *Own*Price Increase* are positive and significant and the coefficient on *Own* is negative and significant. In Panel B, the coefficients on *Own*Price Change* are positive and significant at the 1% level. This indicates that even after allowing for different prior beliefs based on ownership and ownership interacted with price, participants exhibit greater over-extrapolation for owned goods.

As with any benchmark, there is a concern that it is misspecified. We address this by presenting a series of results which do not rely on distributional assumptions. Panel B presents a number of alternative specifications where we control for return levels. The order of returns should not matter for a Bayesian. Thus, if dummy variables for the direction of recent price movements are significant after controlling for the effect of returns, this is evidence that these investors are over-extrapolating from recent signals.

Column 1 of Panel B includes a linear control for returns. The coefficient on *Own * (Price Increase)* indicates that participants extrapolate 3.4 more from a positive signal about owned goods than they do from the same signal about non-owned goods. The coefficient on *Own* indicates that participants extrapolate 2.45 more from negative signals about owned goods than they do from the same signals about non-owned goods. Linear controls may obfuscate interesting dynamics of the return response pattern, so in Column 2 we include dummy variables for levels of return in 10% increments. Including these controls yields similar

³⁴We drop observations where an equal number of positive and negative signals have been observed. For such observations, a response of $\hat{s}_t^i = 50$ is consistent with any symmetric β prior, and hence any other response is inconsistent with any symmetric β prior.

results. It may also be the case that the extrapolation coefficients are capturing differential updating to return levels based on ownership rather than extrapolation. Column 3 includes a linear control for returns and also an interaction of return with *Own* to capture such a differential reaction. Again, results are similar suggesting, 2.13 greater extrapolation from positive signals and -1.57 greater extrapolation from negative signals about owned positions. Column 4 includes dummy variable for returns along with an interaction of those dummy variables. This flexibly controls for level of returns separately for owned and non-owned positions. The pattern of results is unchanged.

Relative to Bayesian benchmarks or to flexible controls for return levels, these findings imply that people over-extrapolate from recent signals to a substantially greater extent when learning about owned goods, both relative to non-owned goods and a variety of normative benchmarks.

B.2. Ownership and Attention

We designed an additional treatment to examine whether the proposed attentional mechanism could explain our results. The treatment was designed to direct attention to signals about goods that are not owned. If the effect observed in the baseline condition is driven by ownership channeling attention and exacerbating a heuristic process that over-extrapolates from recent signals, then beliefs about these non-owned goods should resemble those of owned goods in the baseline condition.

In this treatment, we asked participants to buy three of six goods as in the baseline condition. However, after the purchase decision they reported beliefs only for the goods they did not own. By only incentivizing beliefs about non-owned goods we sought to shift attention towards signals associated with them. If the effect reported in the preceding sections was driven by ownership channeling greater attention towards associated signals, then beliefs about non-owned goods in this treatment should resemble those about owned goods in the baseline condition.

Table V Panel A repeats the regressions from Table I and Table V Panel B repeats the regressions from Table IV, adding the data from the new treatment. The *Own* dummy is equal to one for goods owned by participants in the baseline condition. *No Own Treat* is a

dummy variable that is equal to one for observations in the attentional treatment. Thus, the *Own* variables can be interpreted similar to the prior regressions: the difference in updating from signals about an owned good relative to a non-owned good in the baseline condition. The *No Own Treat* coefficient represent the difference in beliefs about non-owned goods in the attentional treatment compared to non-owned goods in the baseline condition.

Table V shows that beliefs about non-owned goods in the attentional treatment resemble beliefs about owned goods in the baseline condition. For example, looking at Column 2 of Panel A which includes price fixed effects, the coefficient on $(No\ Own\ Treat)*Return$ is 0.185 and is significant at the 1% level, which is similar to the point estimate in the baseline condition. Examining extrapolation in Panel B we again see a positive and significant coefficient on the $(No\ Own\ Treat)*(Price\ Increase)$ variable, consistent with over-extrapolation of recent price increases in the attentional treatment condition. The beliefs about goods in the attentional treatment are generally closer to owned goods than to non-owned goods in the baseline condition. These results provide evidence for ownership channeling attention to recent signals and exacerbating over-extrapolation in the belief-updating process.³⁵

C. *Exploring Ownership: Active Choice or Random Endowment?*

While the belief elicitation study involved an active choice to purchase goods, the experiment eliciting valuations used random endowment. Here, we examine whether this difference in how ownership was attained matters for the learning process. We ran a treatment of the baseline condition in the belief elicitation study, but instead of having people purchase three of the six goods, we randomly allocated three goods to the participants. All other aspects of the treatment were identical to the baseline condition. If the differential learning results are due to active choice, then beliefs about owned goods in the random allocation treatment should not differ from those about non-owned goods. On the other hand, if the results are driven by ownership – whether randomly endowed or not – then the learning pattern in this

³⁵Findings from the two-good treatment provide further suggestive evidence for the attentional channel. Under the assumption of limited attention, relatively more attention can be devoted to both the owned and non-owned goods in the two-good treatment which would increase the scope for over-reaction. Looking at Table IA.2 in the Appendix, the coefficient on *Return* is .03 greater than that in Table I and the coefficient on $Own*Return$ is .04 higher. This suggests a greater *level* of over-reaction in the two-good experiment compared to the six-good experiment for owned and non-owned goods. These results are merely suggestive though, while the coefficient on $Own*Return$ is 50% bigger, the difference is not statistically significant.

treatment should resemble the baseline condition.

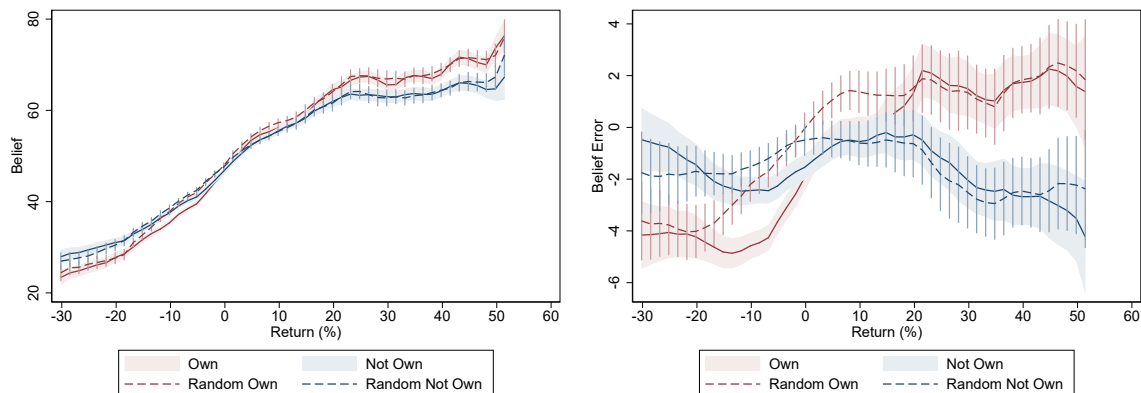


Figure 5. Beliefs in Baseline and Random Allocation Condition. The left figure shows a local linear plot of beliefs, \hat{s}_t^i , on returns separately for goods that are owned and not owned. The Figure repeats Figure 1, but overlays the same graph based on the random allocation condition which is graphed as the dashed line. The right Figure shows the belief error relative to a benchmark of a Bayesian agent with a $\beta(2.62)$ prior based on whether a good is owned as a function of its return. The Figure repeats Figure 4, but overlays the same graph based on the random allocation condition which is graphed as the dashed line. Data include observations with returns from the 5th to the 95th percentile. Shaded area represents the 95% confidence interval for the baseline condition and vertical lines represent the 95% confidence interval for the random allocation condition.

Figure 5 graphs the results of this random allocation condition and illustrates that the learning pattern is similar to the baseline condition; active choice does not seem to be necessary to generate the observed ownership effect on learning. The left Figure repeats Figure 1, but overlays the same analysis using data from the random allocation condition. The random allocation is graphed as a dashed line with the confidence interval shown by vertical lines. Examining the figure, the lines generally overlap. The right Figure repeats the analysis relative to the benchmark of a Bayesian with a $\beta(2.62)$ prior. The pattern is generally the same.

Table VI repeats the regression analysis using only data from the random allocation condition. The coefficients on the *Own * Return* dummy variable are 0.075, 0.079 and 0.75 (with t-statistics of 2.82, 2.99 and 2.95) in the specifications with no controls, with round by price fixed effects and for belief errors relative to a $\beta(2.62)$. The first coefficient is similar to the one in Table I, while the second two are slightly smaller. This again is consistent with the random allocation condition yielding similar results to the baseline condition where participants chose to actively buy the goods.

While a full exploration of the psychology of ownership is beyond the scope of the current

paper, it is clearly a multi-faceted experience. There are many different ways that one can experience ownership, such as owning a physical good, receiving a payment stream from a good (as in our experiment), feeling responsible about a project due to explicit career concerns, or because an organization has been incorporated into one’s identity (see Pierce, Kostova, and Dirks (2001, 2003) for reviews of the literature). In our setting, the random allocation condition suggests that owning an abstract asset with a revenue stream is sufficient to produce differential learning, but it is possible, and even likely, that the manifestations of ownership will be context-dependent and broader than the conditions we consider here.

V. Stock Market Expectations and Ownership

To examine the generalizability of our laboratory findings, we explore the impact of ownership on learning and beliefs in field data. Studying this question requires information on signals and beliefs, and a setting where it is plausible that agents who hold and do not hold a given good are reasonably aware of the signals when forming beliefs. For this reason, we examine beliefs about aggregate stock market performance.

We study whether the belief response to recent market performance – the signal analogue to our experiment – is different depending on whether the individual owns stocks or not. The data comes from the University of Michigan Survey of Consumers. The survey asks whether a respondent owns stocks as well as “What do you think is the percent chance that a one thousand dollar investment in a diversified stock mutual fund will increase in value in the year ahead, so that it is worth more than one thousand dollars one year from now?” We interpret stated beliefs about expectations for the stock market similarly to beliefs about the fundamentals \hat{s}^i in our experiment. The data set contains the relevant data for 187 months, covering the years 2002 until 2019.³⁶

In this setting, investors select into owning stocks and may be differentially aware of information relating to those investments (unlike in our experiment where goods are identical ex-ante). These concerns are somewhat mitigated by examining aggregate market performance since recent market performance (i.e. the signals) is widely reported and discussed in

³⁶In the Appendix we repeat the analysis using the RAND American Life Panel survey on the effects of the financial crisis. This sample has relevant data for the shorter period of 2011 through 2016 so we focus on the Michigan survey. Results are materially similar using the RAND survey.

the media. In turn, it is likely that people are aware of it regardless of whether they own stocks. Additionally, to the extent that owners and non-owners differ based on observable characteristics, we employ a rich data set to control for these factors. However, concerns about systematic differences based on non-observable factors remain, so these results should be viewed as complimentary to the experimental findings where such issues are mitigated.

To begin, we examine how belief expectations vary with horizon of past market performance. Greenwood and Shleifer (2014) show that in six different surveys, investors extrapolate past market performance to form expectations about the future. To document a similar pattern in our sample, we examine two different left-hand side variable. The first is the percent ranking from 0 to 100 on whether the market will be higher. The second is whether a participant thinks the chance of a market increase is greater than 50%, a proxy of the bearish versus bullish measure used in Greenwood and Shleifer (2014). For past stock market returns we use the CRSP value weighted index over the prior quarter, six months, year and two years.³⁷ Market returns are from the period ending the month prior to when the survey was conducted. For example, if a participant took the survey in June of 2014, the lagged quarterly return would be the cumulative return from March 2014 through May of 2014.

A. Results

Table VII shows coefficients from the regression of beliefs about market performance on lagged measures of market performance. Panel A examines the percent measure, while Panel B examines the expectations above 50% dummy variable. In all specifications there is a positive coefficient that is significant at the 1% level. This indicates that respondents are extrapolating – upon receiving a high signal of past market performance, they believe that this is indicative of a positive state such that high performance will persist. In Panel C we regress the returns over the next twelve months on lagged market performance for every month that we have data for. Consistent with Greenwood and Shleifer (2014) all of the point estimates are negative, though at the one quarter level the coefficient is insignificant. This suggests that a rational investor should predict an inverse relationship between recent past performance and future market performance, but that respondents mistakenly over-extrapolate from past

³⁷Hartzmark and Solomon (2017) argue that investors actually pay attention to market indices such as the S&P 500 or the Dow Jones. The Appendix shows similar results using these measures.

signals which leads to incorrect beliefs about the future.

We now explore how this over-extrapolation varies with stock ownership. To do so, we use the same left-hand side variables and regress them on past performance and interactions of ownership and past market performance. Specifically we examine:

$$Probability\ Increase_{it} = Market_{[-m,-1]} + Market_{[-m,-1]} * Own_{it} + Own_{it} \quad (3)$$

where $Market_{[-m,-1]}$ is the previous market return during the relevant horizon and Own_{it} is equal to one if the participant states that they own the assets. Thus, the coefficient on $Market_{[-m,-1]}$ is the degree of extrapolation of past performance by participants who indicate that they do not own the assets. The coefficient on Own_{it} controls for the average difference in expectation between those who own and do not own the assets. The coefficient of interest is $Market_{t-1} * Own_{it}$. This corresponds to the difference in extrapolation between those who own and do not own the assets.

Table VIII shows that owners of stocks extrapolate significantly more than those that do not own stocks. Panel A examines the percent measure, while Panel B examines the expectations above 50% dummy variable. The first two columns in Panel A present the probability of a market increase regressed on lagged quarterly market returns for participants in the University of Michigan survey. In Column 1, the *Own* dummy has a coefficient of 13 and is significant at the 1% level, indicating that asset owners are about 13% more optimistic than non-asset owners. This is consistent with more optimistic people selecting into owning stocks. The coefficient on lagged market returns is 14.22 and significant at the 1% level. This indicates that those who do not own assets extrapolate based on past market performance. Most important for our investigation, the coefficient on the interaction term with ownership is 17.85 and is significant at the 1% level. This indicates that those who own assets extrapolate from recent signals at roughly *twice* the level of those who do not.

The decision to own stocks is correlated with other demographic variables, so it could be that the ownership effects reported in Column 1 is capturing differences in demographic attributes. Column 2 presents the analysis including a large number of controls; specifically, dummy variables for sex, race, age, geographic region, education, and income. Interestingly, the coefficient on *Own* nearly halves, which indicates that a significant amount of the base

level of optimism between owners and non-owners can be accounted for with demographic variables. That being said, the estimates of extrapolation are robust to demographic controls; if anything, the difference in extrapolation between owners and non-owners becomes larger upon their inclusion. The coefficient for those who do not own the assets is 13.89 while the interaction term has a coefficient of 18.99 – both significant at the 1% level. Even after adjusting for differences in observables, asset owners extrapolate about twice as much as non-asset owners, consistent with the results we observed in the experiment.

Lastly, we explore a variety of different lags of market performance, using both measures of expectations, and find similar results across the board. The next six columns of Panel A repeat the regressions using different lags of past market performance and find the same pattern. Panel B repeats the analysis using the binary measure of a market increase as the dependent variable. Results are qualitatively the same. In the 16 interactions using various lags of past market performance, two measures of future expectations, and with various demographic controls, we find that owners of assets extrapolate more than non-owners, with each specification significant at the 1% level.

VI. Conclusion

In this paper, we examine how owning a good affects learning and beliefs about its underlying quality. We find that upon receiving a negative signal about a good they own, people become systematically more pessimistic. They underestimate the good’s quality both compared to receiving the same signal about a good they do not own and a normative benchmark. We observe that the effect reverses after observing positive signals: people overestimate the good’s quality relative to seeing the same signals for goods they do not own and the normative benchmark. These differences in learning translate to differences in valuations: the valuation gap between owners and non-owners (i.e. the endowment effect) becomes larger when both observe positive signals about the good, and disappears when observing negative signals about the good. In our simple setting, we demonstrate that ownership channels attention, leading to overreaction and exacerbating the over-extrapolation from recent signals for goods that are owned – both in an absolute sense and compared to goods that are not owned. Our paper provides empirical evidence of the the “more is less” effect posited by Dawes (1979) and Dawes,

Faust, and Meehl (1989), whereby more attention leads to less accurate judgments. While it demonstrates the existence of the effect, follow up work should explore questions surrounding its generality and portability to other contexts.

The main contribution of the paper is to cleanly demonstrate that ownership influences learning and beliefs in simple, transparent settings. This leaves open a number of interesting questions regarding how this effect interacts with contextual factors and orthogonal psychological mechanisms which could be related to ownership. Our evidence on the attentional mechanism suggests that the documented relationship between ownership and learning is perceptual.³⁸ In settings where cognitive factors like wishful thinking play a larger role, such as when ownership is linked to identity, a level effect of greater optimism may indeed arise. Conditional on this level effect, however, we would still anticipate an interaction between ownership status and the valence of incoming signals. Moreover, it is important to explore the boundaries and moderators of psychological ownership: would our results extend to settings with multiple owners of the same good, or identity-based goods with no extrinsic payoffs? Lastly, research should examine the ownership effect in settings where attention has an unambiguously positive effect on accuracy and to study how it interacts with other factors that influence learning. For example, Kuhnen (2015) finds that learning is more biased when information is framed in the negative domain than when the *same* information is framed in the positive domain.³⁹ In our setting, this framing effect would likely exacerbate the errors that owners are already making in response to negative signals.

Many market settings involve learning from signals about goods that are owned and non-owned. Almost any setting with durable goods and subsequent resale such as real estate and financial markets involve this aspect. Our paper suggests that ignoring differences in learning and beliefs caused by ownership misses an important component of the decision-making process. The results also have significant implications for the dynamics of trade volume in response to public signals. As demonstrated in our second experiment, the valuation gap between owners and non-owners shrinks in response to bad news and expands in response to good news. This should increase the potential for trade in the former case and decrease

³⁸In a similar vein, Bordalo, Gennaioli, and Shleifer (2012) present evidence that a perceptual mechanism may be responsible for anomalies in choice under risk.

³⁹Note that this is distinct from our setting where the valence of signals is informative about the underlying state (fundamental quality).

it in the latter case. Future research should explore these dynamics in observational and experimental data.

REFERENCES

- Barber, Brad M and Terrance Odean, 2013, The behavior of individual investors, in *Handbook of the Economics of Finance*, vol. 2, 1533–1570 (Elsevier).
- Barberis, Nicholas and Wei Xiong, 2012, Realization utility, *Journal of Financial Economics* 104, 251–271.
- Bohren, J Aislinn, Kareem Haggag, Alex Imas, and Devin G Pope, 2019, Inaccurate statistical discrimination, Tech. rep., National Bureau of Economic Research.
- Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer, 2012, Salience theory of choice under risk, *The Quarterly journal of economics* qjs018.
- Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer, 2017, Memory, attention, and choice, Tech. rep., National Bureau of Economic Research.
- Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer, 2018, Diagnostic expectations and credit cycles, *The Journal of Finance* 73, 199–227.
- Brunnermeier, Markus K. and Jonathan A. Parker, 2005, Optimal expectations, *American Economic Review* 95, 1092–1118.
- Bushong, Benjamin and Tristan Gagnon-Bartsch, 2019, Learning with misattribution of reference dependence, Tech. rep., working paper.
- Caplin, Andrew and Mark Dean, 2015, Revealed preference, rational inattention, and costly information acquisition, *American Economic Review* 105, 2183–2203.
- Carmon, Ziv and Dan Ariely, 2000, Focusing on the forgone: How value can appear so different to buyers and sellers, *Journal of consumer research* 27, 360–370.
- DAcunto, Francesco, Daniel Hoang, Maritta Paloviita, and Michael Weber, 2019, Iq, expectations, and choice, Tech. rep., National Bureau of Economic Research.
- Danz, David N., Lise Vesterlund, and Alistair J Wilson, 2019, Belief elicitation: Limiting truth telling with information on incentives, *Mimeo* .
- Das, Sreyoshi, Camelia M Kuhnen, and Stefan Nagel, 2017, Socioeconomic status and macroeconomic expectations, Tech. rep., National Bureau of Economic Research.
- Dawes, Robyn M, 1979, The robust beauty of improper linear models in decision making., *American psychologist* 34, 571.
- Dawes, Robyn M, David Faust, and Paul E Meehl, 1989, Clinical versus actuarial judgment, *Science* 243, 1668–1674.
- Edwards, Ward, 1968, Conservatism in human information processing, *Formal representation of human judgment* .
- Enax, Laura, Ian Krajbich, and Bernd Weber, 2016, Salient nutrition labels increase the integration of health attributes in food decision-making, *Judgment and Decision Making* 11, 460.
- Fischbacher, Urs, Gerson Hoffmann, and Simeon Schudy, 2017, The causal effect of stop-loss and take-gain orders on the disposition effect, *The Review of Financial Studies* 30, 2110–2129.
- Gabaix, Xavier, 2017, Behavioral inattention, Tech. rep., National Bureau of Economic Research.

- Greenwood, Robin and Andrei Shleifer, 2014, Expectations of returns and expected returns, *Review of Financial Studies* 27, 714–746.
- Griffin, Dale and Amos Tversky, 1992, The weighing of evidence and the determinants of confidence, *Cognitive psychology* 24, 411–435.
- Hartzmark, Samuel M and David H Solomon, 2017, Reconsidering returns .
- Hong, Harrison and Jeremy C Stein, 2007, Disagreement and the stock market, *Journal of Economic perspectives* 21, 109–128.
- Kahneman, Daniel, Jack L Knetsch, and Richard H Thaler, 1990, Experimental tests of the endowment effect and the coase theorem, *Journal of political Economy* 98, 1325–1348.
- Kahneman, Daniel, Jack L Knetsch, and Richard H Thaler, 1991, Anomalies: The endowment effect, loss aversion, and status quo bias, *Journal of Economic perspectives* 5, 193–206.
- Kahneman, Daniel and Amos Tversky, 1972, Subjective probability: A judgment of representativeness, *Cognitive psychology* 3, 430–454.
- Kaustia, Markku, 2010, Disposition effect, *Behavioral finance* 169–189.
- Kuhnen, Camelia M, 2015, Asymmetric learning from financial information, *The Journal of Finance* 70, 2029–2062.
- Kunda, Ziva, 1990, The case for motivated reasoning., *Psychological bulletin* 108, 480.
- Lerner, Jennifer S, Deborah A Small, and George Loewenstein, 2004, Heart strings and purse strings: Carryover effects of emotions on economic decisions, *Psychological science* 15, 337–341.
- Mackowiak, Bartosz, Filip Matejka, and Mirko Wiederholt, 2018, Rational inattention: A disciplined behavioral model, Tech. rep., mimeo.
- Malmendier, Ulrike and Stefan Nagel, 2015, Learning from inflation experiences, *The Quarterly Journal of Economics* 131, 53–87.
- Martin, Daniel, 2017, Strategic pricing with rational inattention to quality, *Games and Economic Behavior* 104, 131–145.
- Marzilli Ericson, Keith M and Andreas Fuster, 2014a, The endowment effect, *Annu. Rev. Econ.* 6, 555–579.
- Marzilli Ericson, Keith M and Andreas Fuster, 2014b, The endowment effect, *Annu. Rev. Econ.* 6, 555–579.
- Massey, Cade and George Wu, 2005, Detecting regime shifts: The causes of under- and overreaction, *Management Science* 51, 932–947.
- Odean, Terrance, 1998, Are investors reluctant to realize their losses?, *The Journal of finance* 53, 1775–1798.
- Pierce, Jon L, Tatiana Kostova, and Kurt T Dirks, 2001, Toward a theory of psychological ownership in organizations, *Academy of management review* 26, 298–310.
- Pierce, Jon L, Tatiana Kostova, and Kurt T Dirks, 2003, The state of psychological ownership: Integrating and extending a century of research, *Review of general psychology* 7, 84–107.

- Shefrin, Hersh and Meir Statman, 1985, The disposition to sell winners too early and ride losers too long: Theory and evidence, *The Journal of finance* 40, 777–790.
- Smith, Stephanie M and Ian Krajbich, 2018, Attention and choice across domains., *Journal of Experimental Psychology: General* .
- Smith, Stephanie M and Ian Krajbich, 2019, Gaze amplifies value in decision making, *Psychological science* 30, 116–128.
- Thaler, Richard, 1980, Toward a positive theory of consumer choice, *Journal of Economic Behavior & Organization* 1, 39–60.

Table I

Beliefs by Returns and Ownership

This table shows how beliefs vary with ownership based on the price level. *Own* is a dummy variable equal to one if the good was purchased by the subject. *Return* is the level of returns. Fixed effects are indicated below the regression results. Standard errors are clustered by subject, and t-statistics are in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Panel A: Continuous Return Regressions			
	(1)	(2)	(3)
Own*Return	0.0715*** (2.95)	0.116*** (5.29)	0.0958*** (5.21)
Return	0.505*** (22.22)		
Own	-0.748 (-1.38)	-0.386 (-0.77)	-0.326 (-0.65)
Ret x Round FE	No	Yes	Yes
Subject FE	No	No	Yes
R ²	0.335	0.396	0.555
Observations	36900	36900	36900
Panel B: Discrete Ranges of Returns Regressions			
	(1)	(2)	(3)
Own*(Return<0)	-1.896** (-2.51)	-2.288*** (-3.32)	-2.192*** (-3.39)
Own*(Ret[0,20])	0.715 (1.09)	0.544 (0.88)	1.229** (2.26)
Own*(Return>20)	5.202*** (4.21)	5.231*** (4.26)	3.996*** (3.72)
Return<0	-19.33*** (-22.99)		
Return>20	10.94*** (11.67)		
Ret x Round FE	No	Yes	Yes
Subject FE	No	No	Yes
R ²	0.326	0.395	0.555
Observations	36900	36900	36900

Table II

Endowment Effect Updating based on Ownership

This table shows how the value of a good vary with ownership based on the ratings for a product. *Own* is a dummy variable equal to one if the subject was endowed with the good. *Rating* is the average star rating for a product in that round. *Last Rating* is the most recent rating. Fixed effects are indicated below the regression results. Regressions in the all studies column contain a fixed effect for the hypothetical treatment. Standard errors are clustered by subject, and t-statistics are in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Panel A: Cummulative Rating							
	Purchase	Trade	All Studies				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Own*Rating	4.767** (2.60)	3.017** (2.19)	3.783*** (3.39)	1.405** (2.03)	3.786*** (3.39)	3.771*** (3.37)	1.334* (1.95)
Rating	1.928* (1.70)	2.680*** (3.15)	2.391*** (3.49)	3.591*** (6.68)	2.323*** (3.34)		
Own	1.318 (0.73)	5.445*** (3.31)	3.620*** (2.98)	4.688*** (3.93)	3.607*** (2.97)	3.548*** (2.92)	4.640*** (3.90)
Subject FE	No	No	No	Yes	No	No	Yes
Round FE	No	No	No	No	Yes	No	No
Review x Round FE	No	No	No	No	No	Yes	Yes
R ²	0.121	0.125	0.121	0.635	0.124	0.130	0.644
Observations	1170	1480	2650	2650	2650	2650	2650
Panel B: Last Rating							
	Purchase	Trade	All Studies				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Own*Last Rating	1.665*** (3.29)	1.835*** (4.35)	1.757*** (5.41)	0.890*** (3.96)	1.714*** (5.28)	1.723*** (5.31)	0.849*** (3.88)
Last Rating	2.148*** (5.16)	2.369*** (6.25)	2.286*** (8.17)	2.723*** (10.59)	2.535*** (8.08)		
Own	0.979 (0.56)	5.513*** (3.29)	3.525*** (2.89)	4.132*** (3.36)	3.569*** (2.96)	4.032*** (3.54)	4.644*** (4.06)
Subject FE	No	No	No	Yes	No	No	Yes
Round FE	No	No	No	No	Yes	No	No
Review x Round FE	No	No	No	No	No	Yes	Yes
R ²	0.0652	0.104	0.0837	0.611	0.0862	0.117	0.644
Observations	1170	1480	2650	2650	2650	2650	2650

Table III

Beliefs Errors by Returns and Ownership

This table shows how belief errors vary with ownership based on the price level. Columns labeled *Error* use belief error relative to a Bayesian with a β (2.62) prior as the dependent variable while columns labeled *Belief* use the belief. Columns labeled *Baseline* include data from the baseline treatment, and those labeled *Including Treatment* include data from the baseline and the treatment with information about the distribution of s_i . *Own* is a dummy variable equal to one if the good was purchased by the subject. *Return* is the level of returns. *Treat* is a dummy variable equal to one if the data is from the treatment which included information about the distribution of s_i . Fixed effects are indicated below the regression results. Standard errors are clustered by subject, and t-statistics are in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

	Baseline		Including Treatment		
	(1) Error	(2) Error	(3) Belief	(4) Error	(5) Error
Own*Return	0.0988*** (4.68)	0.0769*** (4.34)	0.0715*** (2.95)	0.0988*** (4.68)	0.0769*** (4.34)
Return	-0.0233 (-1.10)	-0.0170 (-0.87)	0.505*** (22.23)	-0.0233 (-1.10)	-0.0170 (-0.87)
Own	-0.720 (-1.45)	-0.631 (-1.27)	-0.748 (-1.38)	-0.720 (-1.45)	-0.631 (-1.27)
Treat*Own*Return			0.0289 (0.47)	0.00777 (0.14)	0.0399 (0.97)
Treat*Return			-0.0810 (-1.59)	-0.0451 (-0.95)	-0.0565 (-1.28)
Treat*Own			2.796 (1.55)	2.421 (1.51)	2.216 (1.40)
Treat			0.413 (0.24)	0.550 (0.33)	
Subject FE	No	Yes	No	No	Yes
R ²	0.00631	0.266	0.337	0.00815	0.269
Observations	36900	36900	43830	43830	43830

Table IV

Extrapolation of Signals

This table shows how beliefs vary with recent price changes based on ownership. Price Increase is a dummy variable which is equal to one if the good experienced a price increase in the prior round. In Panel A, column 1 the dependent variable is the raw belief. In Column 2 it is the belief error for a Bayesian with a $\beta(2.62)$ prior. In Column 3 it is the belief error relative to a priors calibrated separately for owned and non-owned positions, indicated by $\beta(Own)$ Error. Column 4 uses priors based on the average parameter from subjects from the prior round by price by ownership condition, indicated by $\beta(Round\ Own)$ Error. Panel B examines raw belief as the dependent variable. *Ret* indicates a linear control for return. *RetDummy* indicates a dummy variable for intervals of 10% returns. Below are indications for those variables with the *Own* dummy variable. Other variables are described in Table I. Fixed effects are indicated in the bottom row. Standard errors are clustered by subject, and t-statistics are in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Panel A: Extrapolation Relative to Benchmarks

	Belief	$\beta(2.62)$ Error	$\beta(Own)$ Error	$\beta(Round\ Own)$ Error
	(1)	(2)	(3)	(4)
Own*(Price Increase)	4.980*** (6.41)	3.781*** (6.01)	2.480*** (3.96)	2.874*** (4.26)
Price Increase	12.64*** (18.82)	-1.157** (-1.99)	-0.449 (-0.77)	3.263*** (5.33)
Own	-2.461*** (-2.72)	-2.356*** (-3.90)	-1.572*** (-2.60)	-1.834*** (-2.70)
R ²	0.0957	0.00298	0.00157	0.0155
Observations	34440	34440	34440	30524

Panel B: Extrapolation Relative to Return Controls

	(1)	(2)	(3)	(4)
Own*(Price Increase)	3.424*** (5.05)	4.203*** (6.34)	2.133*** (4.52)	2.004*** (4.22)
Price Increase	1.510*** (3.04)	-0.0901 (-0.18)	2.163*** (5.06)	0.995** (2.24)
Own	-2.450*** (-3.71)	-2.312*** (-3.66)	-2.034*** (-3.22)	-1.686** (-2.20)
Ret	Yes	No	Yes	No
Ret Dummy	No	Yes	No	Yes
Own x Ret	No	No	Yes	No
Own x Ret Dummy	No	No	No	Yes
R ²	0.350	0.379	0.351	0.382
Observations	34440	34440	34440	34440

Table V

Difference Across No Ownership Treatment and the Baseline Experiment

This table shows how beliefs and extrapolation vary across a no ownership experimental treatment. Panel A explores beliefs and belief errors based on returns while Panel B explores the degree of extrapolation based on a positive price signal the prior period. Regressions include the baseline experiment and the data from the no ownership condition. *No Own Treat* is equal to one if the data is from the treatment condition. *Own* is equal to one if the good is owned and the observation is from the base experiment. Regressions also include *No Own Treat* and *Own* dummy variables. Columns labeled “Forecast” examine raw beliefs while columns labeled $\beta(2.62)$ examine belief errors relative to a Bayesian with a $\beta(2.62)$ prior. Fixed Effects are indicated in the bottom row. Standard errors are clustered by subject, and t-statistics are in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Panel A: Returns			
	Belief		$\beta(2.62)$
	(1)	(2)	(3)
(No Own Treat)*Return	0.170*** (3.63)	0.185*** (4.01)	0.178*** (3.98)
Own*Return	0.0715*** (2.95)	0.116*** (5.28)	0.0986*** (4.68)
Return	0.505*** (22.23)		-0.0230 (-1.09)
Ret x Round FE	No	Yes	No
R ²	0.343	0.403	0.00905
Observations	40275	40275	40275
Panel B: Extrapolation			
	Belief		$\beta(2.62)$
	(1)	(2)	(3)
(No Own Treat)*(Price Increase)	3.570* (1.85)	4.811*** (3.46)	4.707*** (3.42)
Own*(Price Increase)	4.980*** (6.41)	4.208*** (6.37)	3.781*** (6.02)
Price Increase	12.64*** (18.82)	-3.219*** (-5.50)	-1.157** (-1.99)
Ret x Round FE	No	Yes	No
R ²	0.0967	0.410	0.00356
Observations	37590	37590	37590

Table VI

Random Allocation Condition

This table shows how belief errors vary with ownership based on the price level when ownership is randomly assigned. The dependent variable is raw beliefs in columns 1, 2 and belief errors in column 3. These are regressed on *Own*Return*, *Return* and *Own*. Fixed Effects are indicated in the bottom row. Standard errors are clustered by subject, and t-statistics are in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

	Belief		$\beta(2.62)$ Error
	(1)	(2)	(3)
Own*Return	0.0750*** (3.08)	0.0793*** (3.33)	0.0750*** (3.28)
Return	0.488*** (20.16)		-0.0123 (-0.52)
Own	-0.165 (-0.26)	-0.0392 (-0.07)	0.0108 (0.02)
Ret x Round FE	No	Yes	No
Subject FE	0.320	0.381	0.00395
R ²	22680	22674	22680

Table VII
Field Data Expectations

This table shows how market expectations vary with past market forecasts. Panel A regresses probability of a stock market increase over the next 12 months on prior market return from month -m to -1, with m indicated in each column. Panel B repeats the regression with the left hand side variable equal to one if the probability of increase is above 50%. Panel C regresses future 12 month return on past returns for each month with a survey forecast. Standard errors are clustered by month, and t-statistics are in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Panel A: Extrapolation of Probability of Increase				
	3 Month=m	6 Month=m	1 Year=m	2 Year=m
	(1)	(2)	(3)	(4)
Market[-m,-1]	26.62*** (4.10)	22.61*** (5.71)	19.62*** (8.84)	13.43*** (8.52)
Observations	98828	98828	98828	98828
Panel B: Extrapolation of Increase Probability >50				
	3 Month=m	6 Month=m	1 Year=m	2 Year=m
	(1)	(2)	(3)	(4)
Market[-m,-1]	0.376*** (3.97)	0.320*** (5.60)	0.282*** (8.98)	0.192*** (8.65)
Observations	98828	98828	98828	98828
Panel C: Future Returns on Past Returns				
	3 Month=m	6 Month=m	1 Year=m	2 Year=m
	(1)	(2)	(3)	(4)
Market[-m,-1]	-0.176 (-1.18)	-0.194** (-2.01)	-0.167** (-2.52)	-0.204*** (-4.96)
R ²	187	187	187	187

Table VIII

Field Data Extrapolation by Ownership

This table shows how extrapolation of prior market performance varies with ownership. Panel A examines the probability of a stock market increase over the next 12 months and Panel B examines a dummy variable equal to one if this is greater than 50. Prior market return is from month -m to -1, with m indicated in each column. *Own* is a dummy variable equal to one if the subject owns stocks. Demographics indicate fixed effects for income, age, race, marital status and education. Standard errors are clustered by month, and t-statistics are in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Panel A: Probability of Increase

	3 Month=m		6 Month=m		1 Year=m		2 Year=m	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Own*Mkt[-m,-1]	17.85*** (4.28)	18.99*** (4.41)	18.22*** (8.49)	18.86*** (8.90)	14.16*** (8.47)	14.61*** (8.89)	8.427*** (7.63)	8.516*** (7.65)
Mkt[-m,-1]	14.22*** (2.73)	13.89*** (3.42)	9.857*** (2.90)	9.775*** (3.56)	9.734*** (4.78)	9.217*** (5.74)	7.366*** (5.04)	7.140*** (6.25)
Own	13.26*** (47.87)	8.074*** (30.48)	12.75*** (48.50)	7.562*** (29.45)	12.20*** (41.02)	7.010*** (24.60)	11.97*** (33.82)	6.833*** (20.32)
Demographics	No	Yes	No	Yes	No	Yes	No	Yes
R ²	0.0514	0.112	0.0546	0.116	0.0591	0.120	0.0601	0.121
Observations	98828	92264	98828	92264	98828	92264	98828	92264

Panel B: Increase Probability >50

	3 Month=m		6 Month=m		1 Year=m		2 Year=m	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Own*Mkt[-m,-1]	0.240*** (3.98)	0.253*** (4.01)	0.243*** (7.29)	0.247*** (7.16)	0.182*** (6.69)	0.183*** (6.80)	0.105*** (5.68)	0.104*** (5.64)
Mkt[-m,-1]	0.209** (2.44)	0.204*** (2.88)	0.149*** (2.77)	0.151*** (3.37)	0.154*** (4.66)	0.151*** (5.58)	0.116*** (5.05)	0.114*** (6.20)
Own	0.190*** (43.23)	0.108*** (23.89)	0.183*** (42.31)	0.101*** (22.12)	0.176*** (34.62)	0.0950*** (18.29)	0.174*** (29.00)	0.0931*** (15.92)
Demographics	No	Yes	No	Yes	No	Yes	No	Yes
R ²	0.0395	0.0910	0.0419	0.0935	0.0454	0.0968	0.0462	0.0973
Observations	98828	92264	98828	92264	98828	92264	98828	92264