Can Agents Add and Subtract When Forming Beliefs? Evidence from the Lab and Field*

Pascal Kieren, Jan Müller-Dethard, and Martin Weber†

November 2021

Abstract

We study an intrinsic property of Bayesian information processing which does not rely on individuals having rational *absolute* beliefs: two equally-diagnostic signals of opposite direction should cancel out. Using evidence from both the lab and field, we show that individuals not always follow this counting-based principle. Systematic violations occur whenever a sequence of identical evidence is interrupted by a signal of opposite direction, which produces strong and robust overreactions. Conversely, individuals correctly follow this counting-based principle whenever signals alternate while they underreact to sequences of same-directed evidence. Next, we empirically analyze announcement and post-announcement stock return reactions in financial markets. Consistent with our experimental evidence, we find that initial stock reactions are significantly stronger and subsequent price drifts weaker for opposite-directed earnings surprises than for same-directed earnings surprises. Our results provide novel insights to the paradoxical co-existence of over- and underreaction to new information at the individual and market level.

Keywords: *Belief Formation, Bayes Theorem, Information Processing, Overreaction*

JEL Classification: *D81, D83, D84, G41*

---

* For valuable comments, we thank Alex Imas, Oliver Spalt, as well as participants of the Annual Meeting of the SJDM 2020, the Maastricht Behavioral Experimental Economics Symposium 2021, the ERIC 2021, the Experimental Finance 2021, the ESA 2021, the EEA 2021, and seminar participants at the University of Mannheim.

† Pascal Kieren is at the University of Mannheim, Finance Department, L9, 1-2, 68161 Mannheim, pascal.kieren@gess.uni-mannheim.de (corresponding author), Jan Müller-Dethard is at the University of Mannheim, Finance Department, L9, 1-2, 68161 Mannheim, mueller-dethard@bank.bwl.uni-mannheim.de, Martin Weber is at the University of Mannheim, Finance Department, L9, 1-2, 68161 Mannheim, and CEPR, London, weber@bank.bwl.uni-mannheim.de.

1. Introduction

Modeling human information processing as Bayesian is standard in cognitive science, neuroscience, and economics. However, there is ample evidence that individuals' actual beliefs often deviate in systematic ways from Bayesian beliefs. Some of the documented biases suggest that individuals generally underinfer from new information (e.g. Edwards, 1982), while others indicate that people update too much (e.g. Kahneman & Tversky, 1972). What is less clear, is the question *when* we may expect one versus the other (Benjamin, 2019). To judge the rationality of beliefs, studies frequently compare individuals' *subjective* assessment of probabilities with the *normative* Bayesian benchmark.

In this paper, we study a fundamental property of information processing which does not rely on individuals having the "right" or "rational" *absolute* beliefs: two equally-diagnostic signals of opposite direction should cancel out. Thus, taken together they should not affect prior beliefs. Relying on this intrinsic property of Bayesian updating allows us to cleanly investigate *belief changes*, while remaining "agnostic" about the rationality of *absolute beliefs*. To illustrate this idea, consider a doctor who wants to learn about the medical condition of a patient. To assess whether the patient suffers from a particular disease or not, the doctor launches a series of equally-diagnostic tests. The first test signals that the disease is not present, making the doctor rather optimistic about the patient's condition. The second and third tests yield balanced results, one shows again that the disease it not present, whereas the other signals that the patient suffers from the disease. Would the doctor be just as optimistic as she was after the first test? In other words, assuming all tests are equally informative, is *one* negative test result just as good as *two* negative results and *one* positive, as prescribed by principles of rational information processing?

Using evidence from both the lab and field, we show that this is not always the case. In a series of experimental studies, we document consistent and systematic patterns of both- over and underreaction to new information that can be tightly associated with the intrinsic properties of Bayesian updating. Instead of balancing out two equally-diagnostic but opposite pieces of information, individuals strongly overinfer from evidence which is inconsistent with prior information. Conversely, we find that individuals generally underinfer in situations in which they cannot rely on this "counting-based" principle and in which they would need to know Bayes' rule (e.g. consecutive pieces of same-directed information). As potential mechanism, we identify that individuals update their expectations as if they must selectively allocate their cognitive resources. Successive same-directed information signals are increasingly underweighted as individuals deem their beliefs as *good enough*. Information which disconfirms prior beliefs is – however – treated as more valuable to process as it indicates that their beliefs might not be in line with reality. Based on these insights, we turn to financial markets and apply our well-identified findings to market reactions to firm earnings announcements. Consistent with our experimental data and the identified mechanism, we show that stock market reactions are considerably stronger for earnings surprises which are of opposite direction compared to earnings surprises which are of same direction relative to a firm's prior earnings announcements. This differential reaction to earnings surprises constitutes a mispricing which suggests that prices react not only to the absolute content of released news, but also to a bias induced by relative content.

In order to identify whether individuals correctly apply this counting-based principle when forming expectations, we designed a setting in which the received information is as-if exogenously assigned, beliefs can be clearly elicited, and a normative benchmark for learning can be established. In a series of four experimental studies with more than 2,000 participants, we create environments in which participants repeatedly observe binary signals to learn about one of two possible underlying states of the world. While such a binary decision-making problem appears to present a specific, commonly used and simplified setting in experimental research, it applies to many every-day decision problems (e.g. are we in a good or bad stock market regime, is a defendant guilty or not). In the experiment, the fundamental quality of a risky asset could either be determined by a *good distribution* or a *bad distribution*. Participants knew that the quality is fixed at the start of the experiment and could learn about the quality by observing price changes. The probability of observing a price increase or decrease depends on the underlying distribution of the risky asset such that price increases signal that the good distribution is more likely, whereas price decreases signal that the bad distribution is more likely. In each period in which new information is revealed, we elicit beliefs about the probability that the risky asset's quality is determined by the good distribution – with truthful reporting incentivized. To test the information processing principle described earlier, we exogenously manipulate the number of subsequent same-directed signals that gets interrupted by a single signal which disconfirms prior evidence. Overall, the learning environment in our experiment is fairly simple and transparent, since the price change in each period is a sufficient statistic to infer the underlying quality of the risky asset. This setup allows us (i) to test how individuals revise their beliefs both after a sequence of same-directed information as well as after observing evidence that is inconsistent with prior information; and (ii) to cleanly identify *when* individuals over- and underreact to new information without relying on strong assumptions about individuals' rationality.

Our experimental findings can be summarized as follows. First, we consistently find that subjects systematically and strongly overreact after observing a signal which disconfirms a prior streak of same-directed information. In relative terms, subjects overweight such *disconfirming* evidence by about four times the magnitude compared to predictions of rational information processing. Referring to our opening example, this suggests that individuals do not rely on basic principles of information processing by counting the difference between positive and negative signals. Conditional on the same informational value, subjective beliefs are substantially more extreme after short streaks that only consist of identical information (e.g. three good signals) than after longer streaks that also consist of mixed information (e.g. four good and one bad signal). Importantly, the overreaction we document is independent of whether subjects' prior beliefs are in line with normative Bayesian benchmarks or not, since the *belief changes* we investigate do not require the same distributional assumptions than *absolute beliefs*.

Next, we seek to establish conditions *when* individuals overreact and *when* they underreact to new information. We find that overreactions in the formation of expectations strongly correlate with violations of the principle that two equally-diagnostic signals of opposite direction should cancel out. In other words, individuals generally overinfer as soon as a sequence of same-directed information is interrupted by a single (equally-diagnostic) signal which disconfirms prior evidence. Additionally, the overinference is independent of subjects having extreme prior beliefs and only requires a prior sequence of at least two signals that go in the same direction. Conversely, we find that individuals correctly follow this counting-based principle whenever signals alternate while they underreact to sequences of same-directed evidence.

Finally, we investigate two potential mechanisms: (1) a false belief in trends or reversals (e.g. price reversals in financial markets), and (2) a misallocation of cognitive resources. For example, the paradoxical co-existence of both over- and underreaction to new information in financial markets is frequently associated with a combination of multiple behavioral biases, which cause investors' beliefs to shift between a regime where prices either trend or revert (e.g. Barberis, Shleifer, & Vishny, 1998; Rabin, 2002; Rabin & Vayanos, 2010). To test whether a false belief in trends or reversals can account for the pattern we find, we conduct a variation of our baseline experiment in which participants have perfect foresight of future outcomes. Yet, even in a setting in which false beliefs about future outcomes can be confidently ruled out, we find stable and pronounced patterns of both over- and underreaction.

Next, we turn to explore the second link, which is a misallocation of cognitive resources. Recent studies show that the introduction of mental updating costs can jointly rationalize prominent biases that predict over- and underinference and thus point into different directions (Kominers, Mu, & Peysakhovich, 2017). In this framework, individuals face a mental updating cost which they pay every time they update their beliefs. In the light of such an updating cost, individuals need to selectively allocate their cognitive resources such that they only fully update when they deem a signal "valuable" and ignore "invaluable" signals. To test this link, we used tools from cognitive psychology to incorporate a change detection task into the baseline experiment, allowing us to track participants' response time and their allocation of cognitive resources between same-directed and disconfirming evidence. We find that individuals take significantly more time to process disconfirming evidence relative to confirming evidence. Importantly, we find that the longer the streak of confirming signals, the faster individuals tend to process subsequent confirming signals. Conversely, the longer the streak of confirming signals before individuals observe a signal which disconfirms prior evidence, the longer they take to process the disconfirming signal. Consistent with the hypothesis that individuals suffer from a misallocation of cognitive resources, repeated signals which confirm prior evidence become increasingly "invaluable" for the decision maker as they consider their beliefs as *good enough*. In contrast to that, once a signal disconfirms a prior sequence of evidence, individuals realize that their beliefs might not be in line with reality and devote significantly more time to process it, thereby overreacting to the signal. This finding is also consistent with the 'more is less' hypothesis of Dawes (1979), who conjectures that greater attention may lead forecasters to overweight properties of the decision problem.

In a next step, we apply our experimental insights to a real-case application in financial markets, that is, stock return reactions to earnings announcements. Quarterly earnings announcements represent one of the main recurring sources of news released by publicly-traded firms. Prior to the actual earnings announcement, both investors and financial analysts form beliefs about what the actual earnings will be. After the announcement, earnings surprises, i.e. the degree to which the announced earnings exceed or fall short of those expectations, lead to stock price movements, as they represent new information that shifts investors' expectations of firm prospects. Consistent with our prior conjecture, we examine how the sign of prior earnings surprises relative to the current surprise affects announcement and post-announcement stock returns. We argue that the presence of a large subset of investors who display similar over- and underreaction patterns as observed in our well-identified experimental results can generate stock price over- and underreactions to news, respectively, and, in turn, return predictability and post-announcement price drift. If the earnings surprise is of opposite sign relative to previous earnings surprises, we predict a stronger announcement stock return reaction than if it is of the same sign. This follows from the experimentally observed updating behavior that shows a stronger absolute reaction after an opposite-directed signal relative to a prior same-directed signal. Moreover, and given the stronger initial reaction after opposite-directed earnings surprises, we would predict a weaker post-announcement drift following opposite-directed earnings surprises.

To test this hypothesis, we examine stock returns and earnings announcements of more than 1400 US firms over the period from 2009 to 2020. Consistent with our experimental data, we find that initial stock reactions are significantly stronger for opposite-directed earnings surprises than for same-directed earnings surprises. That is, after positive earnings surprises, we observe on average stronger return reactions if the surprise is following a prior negative surprise, compared to when it is following a prior positive surprise. The same holds true for negative earnings surprises that follow a prior positive earnings surprise. This larger initial reaction also affects the strength of subsequent price drifts. For positive surprises, we observe an upward price drift, suggesting an initial underreaction. However, the drift is stronger for same-directed positive surprises than for opposite-directed positive surprises resulting in a convergence of prices in the long run, consistent with mispricing that is eventually corrected. For negative surprises, we cannot observe a clear drift pattern. Prices of stocks with same- or opposite-directed negative earnings surprises remain on average relatively stable. Consequently, the initial difference carries forward and persists even 60 days after the earnings announcement.

Our paper contributes to several strands of literature. First, we contribute to the various studies that document biases and heuristics in probabilistic reasoning (for an overview see Camerer, 1987, 1995, Benjamin, 2019). A common finding, by and large is that people update too little (Peterson et al., 1965; Dave & Wolfe, 2003; Sasaki & Kawagoe, 2007), with three exceptions as noted by Benjamin (2019): (i) People overinfer from signals if the diagnosticity is low (Peterson & Miller, 1965; Grether, 1992; Griffin & Tversky, 1992), (ii) people may overinfer when signals go in the same direction of the priors reinforcing an agent's current beliefs (i.e. prior-biased updating, Geller & Pitz, 1968; Pitz, 1969; Pitz et al., 1976), and (iii) people may overinfer when priors are extreme and signals go in the opposite direction of the priors moving an agent's beliefs away from certainty (due to base-rate neglect, DuCharme &

Peterson, 1968; Charness & Dave, 2017). Especially, (ii) and (iii) push in opposite directions which makes it important to understand when one or the other dominates. Our study first documents a common link between both phenomena: individuals on average overinfer whenever they violate basic counting-based principles of information processing. Conversely, we document a persistent underinference whenever individuals cannot or do not violate such a counting heuristic. This is either because there are (i) only signals of same direction, or (ii) positive and negative signals alternate. Additionally, we highlight a channel through which both over- and underreaction to new information might operate and confirm the information processing cost hypothesis by Kominers, Mu, and Peysakhovich (2017). Finally, our study furthers the understanding of how cognitive biases found in the lab manifest in the real world (Levitt and List, 2007a, 2007b; Hartzmark and Shue, 2018) by showing that the bias we document persists in financial markets where many sophisticated investors interact.

Second, we also contribute to the literature on the paradoxical co-existence of both over- and underreaction of prices to information in financial markets (De Bond & Thaler, 1985). In order to generate both over- and underreaction, the most prominent models rely on one (or an interaction) of the following building blocks: multiple behavioral biases, information asymmetry and a long-lasting exogenous belief distortion (i.e. there is limited learning even after a long history of signals). Barberis et al. (1998) propose a model in which agents suffer from two types of biases – the representativeness heuristic and the conservatism bias to generate both over- and underreaction. Hong and Stein (1999) build a model with two types of investors – news watchers and momentum traders – and with asymmetric information. Daniel et al. (1998) also rely on the interaction of two types of behavioral biases: overconfidence and the self-attribution bias. Our results indicate that both over- and underreaction to information can be generated by a single paradigm which does not rely on the usual building blocks. In both our lab and field data, we show that individuals generally underinfer when they observe information signals that confirm a history of prior signals (e.g. observing a positive earnings surprise when prior earnings announcements of the same company were also associated with positive surprises). However, whenever individuals observe information that is at odds with the history of prior signals, they generally overinfer. As underlying mechanism, we identified that agents must selectively allocate their resources. Having limited cognitive capacities, individuals devote increasingly less attention to signals that confirm their priors as they deem their beliefs as "good enough", thereby underreacting to same-directed information signals. Conversely, information which disconfirms prior evidence signals that individuals' beliefs might no longer be an accurate representation of reality. In response, individuals devote significantly more time to process the signal, thereby overweighting its informational value, consistent with a ‘more is less’ effect of attention (Dawes, 1979).

Finally, we contribute to the literature on stock return reactions to earnings announcements. One of the most puzzling anomalies in financial economics is the underreaction to earnings announcements resulting in the so-called post-earnings announcement drift (Bernard & Thomas, 1989, 1990). While an extensive body of literature consensually reports that stock prices appear to drift after major news announcements, the reasons for its occurrence are controversial and challenge market efficiency. Even though recent studies find an attenuated drift for large-cap US stocks (Martineau, 2019; Richardson et al., 2010), it remains puzzling why the effect has remained or even still remains so robust across many countries and over time (Hung et al., 2015). Motivated by the updating behavior observed in various experiments, we propose an interesting new dimension which affects the initial price sensitivity towards an announcement and consequently the intensity of a price drift. This is, how the sign of a firm’s earnings announcement surprise relates to the sign of the firm’s prior earnings announcement surprises. Per se, this information should not have any impact on prices, which are determined through the interaction of many investors. Thus, even if a subset of investors would suffer from such a cognitive bias, their impact on market prices would be limited by the disciplining presence of arbitrage. Yet, we show that investors indeed tend to incorporate new positive information in a more comprehensive manner when it is contrary to the prior information, resulting in less underreaction and a smaller post-announcement drift. This form of mispricing suggests that prices react not only to the absolute content of released news, but also to a bias induced by relative content. Recently, Hartzmark and Shue (2018) document a similar error in which investors perceive earnings news as either more or less impressive depending on other earnings surprises released beforehand. We show that such a relative valuation not only exists in the cross-section around a particular earnings announcement but also in the time-series of announcements of a given firm.

The rest of the paper proceeds as follows. In Section II, we present the empirical framework which we use to explore individuals’ information processing behavior and derive hypotheses. In Section III, we describe the experimental design, present our experimental results, and explore the underlying mechanism. In Section IV, we turn to an application in financial markets, describe our data, empirical methodology and present results. Finally, in Section IV we conclude.

2. Empirical Framework and Hypotheses

In this section, we describe the framework which serves as a basis for our hypotheses as well as the later empirical analyses. Suppose there is an agent who wants to learn about the quality of a risky asset. The risky asset can generate outcomes from one of two distributions, denoted *G* and *B*. To learn about the true underlying distribution of the risky asset, the agent can observe two possible information signals, s_+ and s_-. The data generating process of both distributions is as follows:

$$G\colon \begin{cases} s_+ & w.prob. \quad \theta \\ s_- & w.prob. \quad 1-\theta \end{cases} \qquad\qquad B\colon \begin{cases} s_+ & w.prob. \quad 1-\delta \\ s_- & w.prob. \quad \delta \end{cases},$$

with both $\theta > 0.5$ and $\delta > 0.5$. As such, observing s_+ signals that distribution G is more likely to be the underlying true distribution, whereas observing s_- signals that distribution B is more likely to be the underlying true distribution. Let *n* denote the number of s_+ signals that the agent observes until some period *t* and *m* denote the number of s_- signals that the agent observed. Then, the objective Bayesian posterior probability that *G* is the true underlying state of the world conditional on the observed history of signals *S* can be calculated as:

$$P_t^{Bayes}(G \mid S) = \frac{\theta^n(1-\theta)^m}{\theta^n(1-\theta)^m + (1-\delta)^n\delta^m} \qquad (1)$$

Next, assume that both signals are equally informative.[3] In other words, signal s_+ is just as diagnostic about state G as signal s_- is diagnostic about state B, which requires that $\theta = \delta$. When $\theta = \delta$, the above equation simplifies to:

$$P_t^{Bayes}(G \mid S) = \frac{\theta^{n-m}}{\theta^{n-m}(1-\theta)^{n-m}} \qquad (2)$$

Equation 2 provides several implications on how individuals should update their expectations when observing binary information signals. First, note that a Bayesian agent in this setting is indifferent regarding the order of the signals, since only the difference *n-m* observed until point *t* is of relevance. Next, given the fact that the equation only depends on θ (i.e. that both signals are equally diagnostic about one state of the world), individuals should update their prior expectations with the same magnitude when observing two unequal signals. For example, after observing signal s_+ followed by s_-, a Bayesian would increase his priors that state G is the

[3] While simplifying the analysis, this assumption is a commonly used paradigm to approximate uncertain underlying states of the world where probabilities cannot be determined with any numeric precision. One could still investigate a similar counting rule by relaxing the assumption. Instead of investigating *two* signals, one would require *multiple* signals (in proportion to the diagnosticity).

objective state of the world with the same magnitude as he would subsequently decrease his prior after observing s_- such that his final posterior belief is effectively unchanged. Formally, this requires that the following relation holds:

$$|\, P_{t+1}(G|S_{t-1}, s_+, s_-) - P_t(G|S_{t-1}, s_+) \,| = |\, P_t(G|S_{t-1}, s_+) - P_{t-1}(G|S_{t-1}) \,| \qquad (3)$$

Equation (3) dictates that the absolute change in posterior beliefs from any period $t-1$ to period t conditional on observing signal s_+ should be equal to the absolute change in posterior beliefs from period t to period $t+1$ conditional on observing signal s_-. In other words, two equally diagnostic signals of opposite direction should simply cancel out. This relation is also visualized in Figure 1.

[INSERT FIGURE 1 ABOUT HERE]

Figure 1 plots Equation 2 with an exemplary sequence of three signals in the order s_+, s_+, s_-, for $\theta = 0.7$.[4] After a sequence of *same-directed* signals, an agent cannot yet rely on the simple counting metric in determining the precise probability estimate. In other words, to state the correct *Bayesian posterior probability*, the agent would need to know Bayes' Rule. However, once a signal occurs which *disconfirms* prior evidence (i.e. s_- in this case), the agent should reduce his prior belief about the underlying state by the same magnitude than he increased it after the previous signal (s_+). In other words, a Bayesian agent would report the same probability estimate than he did two signals ago, which requires that $|P_3 - P_2| = |P_2 - P_1|$ holds, as determined by equation 3. Importantly, this approach of investigating *belief updating* does not require individuals to have the "right" beliefs, as we can still judge belief *changes* based solely on the intrinsic properties of Bayesian updating. As such, while we remain "agnostic" about the rationality of *absolute* beliefs, we can nonetheless investigate a simple principle of rational information processing: two equally-diagnostic signals of opposite direction should cancel out.

Based on the established framework, we can differentiate three possible cases of how individuals process new information which disconfirms prior evidence:

1) $\boldsymbol{|\, P_{t+1}(G|S_{t-1}, s_+, s_-) - P_t(G|S_{t-1}, s_+) \,| - |\, P_t(G|S_{t-1}, s_+) - P_{t-1}(G|S_{t-1}) \,| = 0}$

[4] Note that the displayed relationship does not depend on the value of θ. The diagnosticity only governs how fast beliefs converge after each signal.

This case corresponds to rational information processing based on the intrinsic properties of Bayesian learning.

2) $\boldsymbol{| P_{t+1}(G|S_{t-1}, s_+, s_-) - P_t(G|S_{t-1}, s_+) | - | P_t(G|S_{t-1}, s_+) - P_{t-1}(G|S_{t-1}) | > 0}$

This case corresponds to decision makers overreacting to information which disconfirms prior information.

3) $\boldsymbol{| P_{t+1}(G|S_{t-1}, s_+, s_-) - P_t(G|S_{t-1}, s_+) | - | P_t(G|S_{t-1}, s_+) - P_{t-1}(G|S_{t-1}) | < 0}$

This case corresponds to decision makers underreacting to information which disconfirms prior information.

3. Experimental Evidence

3.1 Baseline Experimental Design

To study whether subjects revise their beliefs in a manner consistent with a counting-based principle of information processing, we need an environment which requires: (i) a sequential set-up with room for subjective belief formation, (ii) control over Bayesian beliefs, (iii) streaks of signals with variation in the number of confirming signals prior to an (equally-diagnostic) disconfirming signal, and (iv) an incentive-compatible belief elicitation. Our design accommodates all of these features.

Over the course of six periods, we provide subjects with information about a risky asset (similar to Grether, 1980). In all of our experiments, the risky asset has an initial value of 50 which either increases or decreases over the course of six periods depending on the asset's price development. Price changes are either drawn from a "good distribution" or from a "bad distribution". Both distributions are binary with a possible price increase of +5 and a price decrease of –5. In the good distribution, the price increase occurs with 70 % probability while the decrease occurs with 30 % probability. In the bad distribution, the probabilities are reversed, i.e. an increase occurs with 70 % probability while a decrease occurs with 30 % probability.

To keep our design simple and tractable, we focus on a single disconfirming signal within six periods. In our setting, a disconfirming signal is any signal which interrupts a streak of same-directed signals. In other words, a disconfirming signal can either be a price increase following a streak of prior decreases or a price decrease following a streak of prior increases. Similarly, a confirming signal is any signal which *confirms* a prior streak of same-directed signals. This results in six possible price paths per distribution (the disconfirming signal is equally distributed from rounds 1 to 6). Table 1 provides an overview of all possible price paths.

[INSERT TABLE 1 ABOUT HERE]

A key feature of our design is that we shift the single disconfirming signal between a sequence of six signals. That allows us to test how subjects update their beliefs after observing a single disruptive, disconfirming signal conditional on the number of previously observed confirming signals. Additionally, the design makes it possible to investigate how subjects update their expectations when signals alternate (e.g. a price increase is followed by a decrease which again is followed by another increase).

Across all experiments, subjects make forecasting decisions in six consecutive periods. At the beginning of the experiment, the computer randomly determines the distribution of the risky asset (which can be good or bad) and the period in which the disconfirming signal will occur (which can be from one to six). In each of the six rounds, subjects observe either a price increase or decrease of the risky asset. After each round, we ask them to provide a probability estimate that the risky asset draws price movements from the good distribution and how confident they are about their estimate. To keep the focus on the updating task and to not test their memory performance, we display the prior outcomes in a price-line-chart next to the questions.

The experiment concluded with a brief survey about subjects' socio-economic background, self-assessed statistic skills, as well as a measure of risk preferences and financial literacy adopted from Kuhnen (2015). Subjects' belief elicitation was incentivized. Participants were paid a participation fee and a variable fee based on the accuracy of the probability estimate provided. Specifically, they received 25 Cents for each probability estimate within 10 % (+/-5%) of the objective Bayesian value. Across all studies, it took participants approximately 7 minutes to complete the experiment and participants earned $1.50 on average.

3.2 Sample of Subjects

We recruited 2061 individuals from a large crowdsourcing platform called Amazon Mechanical Turk (MTurk) to participate in four online experiments.[5] Crucial for our experimental study is that participants understood the underlying data-generating process. In particular, it is important

[5] MTurk advanced to a widely used and accepted recruiting platform for economic experiments. Not only does it offer a large and diverse subject pool compared to lab studies (which frequently rely on students), but it also provides a response quality similar to that of other subject pools (Buhrmester et al., 2011; Goodman et al., 2013).

that subjects know that a price increase is a positive signal of the risky asset's quality, while a price decrease is a negative signal. To restrict our sample to those participants who understood this learning environment, we include only participants in our analyses whose beliefs were directionally correct, that is, they were diverging at maximum by 50% from the objective Bayesian belief. As such, we give participants – as featured by our design – enough room to form subjective beliefs but can identify those participants whose beliefs evolve in the wrong direction which means that they incorrectly believe that several positive signals indicate that the asset is of bad quality and vice versa. Moreover, to assure that participants have a sufficient understanding of the learning environment, they had to correctly answer several control questions before they could continue (see Appendix A). Overall, this results in a final sample of 1482 out of 2061 participants who completed the experiment and entered the analyses.[6]

For this sample, Table 2 presents summary statistics. In our baseline experiment, participants have an average age of 34 years and 34% were female. Participants reported average statistic skills of 4.37 out of 7 and would be willing to invest 45% of a hypothetical endowment of 10,000 in a broad equity index, which serves as a proxy for their risk-aversion. Financial literacy is medium as measured by 1.31 out of three possible basic errors they make. The summary statistics of subjects are similar for our three further experiments.

[INSERT TABLE 2 ABOUT HERE]

3.3 Experimental Results

A. Main Results of the Baseline Experiment

As outlined in our framework, we begin by analyzing how individuals update their expectations in situations in which a streak of same-directed information signals is interrupted by a disconfirming signal. This requires that subjects observed *at least two* subsequent same-directed signals prior to observing the disconfirming signal. We will analyze the price paths in

[6] Our inclusion restriction which focuses on the identification of participants who mostly update in the right direction is similar to those commonly used in belief-updating experiments. The restriction is applied equally to all our treatments and therefore does not bias the findings systematically in a certain direction. If at all, the restriction makes the results rather more conservative as we can show that the effects exist even among the seemingly smarter participants. The number of participants excluded through this directionally-correct updating filter is well within the range of prior belief-updating experiments (e.g. 25% in Möbius et al., 2021; 49% in Enke & Graeber, 2019; 32% in Hartzmark et al., 2021).

which there was no prior sequence of same-directed signals or in which signals alternate later in this chapter.

As a simple first test of our main hypothesis, we compare subjects' changes in beliefs $(P_t - P_{t-1})$ after a disconfirming signal in period t to their changes in beliefs $(P_{t-1} - P_{t-2})$ after a same-directed signal that occurred in period *t-1*. Note that rational information processing would require that both changes are of identical magnitude as good and bad signals are equally diagnostic. Figure 2 displays subjects' average absolute changes in beliefs in blue bars and the average absolute changes in objective beliefs according to Bayes' rule in red bars. Results are reported individually for negative signals interrupting a sequence of prior positive signals (i.e. underlying distribution is good) and for positive signals interrupting a sequence of prior negative signals (i.e. underlying distribution is bad).

[INSERT FIGURE 2 ABOUT HERE]

From Figure 2, we observe for each distribution that the second blue bar is significantly larger than the first blue bar, consistent with a greater reaction to disconfirming signals. Subjects in the good distribution update their prior beliefs by 13.71 % on average after a disconfirming signal, whereas they update their prior beliefs by only 3.00 % on average after the same-directed signal in the period prior to the disconfirming signal. In relative terms, this means that subjects in the good distribution update their beliefs after a disconfirming signal with a magnitude that is approximately four times too much relative to their prior belief update. Compared to the Bayesian benchmark which implies equal absolute changes in beliefs for a sequence of two signals of opposite direction, we observe that subjects on average underinfer by 4.68 % after a same-directed signal ($p < 0.01$) and strongly overinfer by 6.03 % after a disconfirming signal ($p < 0.01$). As such, the observed overinference after opposite-directed information represents not only an inconsistency in belief updating relative to subjects' own prior beliefs – as *changes in beliefs* do not cancel out – but also a bias relative to absolute Bayesian beliefs – as *estimate levels* systematically deviate from objective benchmarks.[7] In the bad distribution, the findings look similar. Subjects change their prior beliefs by 14.28 % on average after a disconfirming signal, while they change their prior beliefs by 5.98 % on average after the same-directed signal

[7] For an overview of how subjects' belief estimates evolve over each round in each price paths, Figure B.1 in the Appendix can be seen.

in the period prior to the disconfirming signal. The resulting overreaction in the bad distribution is thus of similar magnitude, albeit a little bit less pronounced.

Table 3 examines this pattern in greater detail. Besides the descriptive analysis, we run the following regression, in which we control for the objective posterior probability[8]:

$$\Delta p_{i,t} = \beta_1 \Delta Objective\ Prior_{i,t} + \beta_2 Disconfirm_{i,t} + \varepsilon_{i,t},$$

where $\Delta p_{i,t}$ is the difference in subjects' probability estimates between two subsequent periods and $\Delta Objective\ Prior_{i,t}$ is the difference in the objective Bayesian probability between two subsequent periods. Finally, $Disconfirm_{i,t}$ is an indicator variable which equals one if subject i observes a disconfirming signal in period t. In the above specification we can test both for Bayesian behavior and in which way individuals depart from it. If subjects were perfect Bayesian, we would expect that $\widehat{\beta_1} = 1$, and $\widehat{\beta_2} = 0$. In other words, subjects always update their prior beliefs according to Bayes' Rule, while a disconfirming signal (which disrupts a sequence of confirming signals) would not explain any additional variation. Conversely, $\widehat{\beta_1} < (>)1$ and $\widehat{\beta_2} < (>)0$ would signal underinference (overinference), to subsequent confirming signals and to disconfirming signals, respectively.

[INSERT TABLE 3 ABOUT HERE]

The findings support our previously drawn conclusions. Even after controlling for the objective posterior, we find an economically strong and highly statistically significant overreaction after a disconfirming signal (see columns (2) and (5) in Table 3). Table 3 also implies that subjects generally underinfer which is consistent with several studies on Bayesian updating (see Benjamin, 2019). Interestingly, our results suggest that the observed underinference is mostly driven by subsequent confirming signals.[9] When differentiating between the good and the bad distribution, we find that the observed underinference is similar across distributions. If at all, it is slightly more pronounced when subjects update their beliefs from a sequence of confirming good signals than when updating their beliefs from a sequence of confirming bad signals.

[8] Since we investigate changes in subjective probability estimates, we estimate the model without constant to be consistent with the theoretical benchmark. However, results are qualitatively similar if we estimate the model on levels or with constant. For the ease of interpretation, we report the specification without constant.
[9] This finding can also be seen in Figure B.1 in the Appendix which provides an overview of how subjects' beliefs evolve in each round and treatment over the course of the experiment.

Overall, we identify a robust updating pattern in our baseline experiment: subjects strongly overreact whenever they have to incorporate disconfirming information that interrupts a prior sequence of same-directed information. In these situations, subjects violate a basic principle of rational information processing which implies that they count the difference between positive and negative signals.[10] Instead, they update their beliefs after disconfirming evidence up to four times as much as they should if they were to process the new information in a rational way.

B. Boundary Conditions

Next, we seek to pinpoint under which conditions and as such *in which situations* subjects display the observed overreaction in belief updating. To do so, we utilize the fact that our experimental design not only allows us to examine the counting-rule in situations in which a streak of same-directed evidence is interrupted, but also in situations in which information signals alternate. This is either the case when (1) a disconfirming signal after a prior sequence of same-directed signals is followed by another confirming signal (and thus matches the prior information streak) or (2) when signals alternate right at the beginning without a prior sequence of same-directed signals. In both situations, the same simple counting rule can be applied since two informationally equivalent, signals of opposite direction should cancel out irrespective of the order in which they occur. To examine the first case of alternating signals, we run the following regression model:

$$\Delta p_{i,t} = \beta_1 \Delta Objective\ Prior_{i,t} + \beta_2 Disconfirm_{i,t} + \beta_3 Correction_{i,t} + \varepsilon_{i,t},$$

where all variables are defined as in equation (4) except $Correction_{i,t}$ which is an indicator variable added to the model which equals one if subject i observes a reversion of a disconfirming signal (i.e. a correction) in period t. The results are reported in Columns (3) and (6) in Table 3. We find that subjects in the good distribution increase their probability estimate after the reversion of a previous disconfirming signal on average by 12.05 % which is almost as much as they previously decreased it (15.61% after the disconfirming signal). Similarly, in the bad distribution, subjects decrease their probability estimates on average by 11.11 % which is also almost as much as the previous increase of 14.47% after the disconfirming signal. In essence, the previously observed overreaction after a disconfirming signal is almost entirely

[10] This also implies that – in contrast to the Bayesian prediction – signals are not invariant to the order in which they occur (see Section II). We also further confirm this finding in Table B.1 in the Appendix.

corrected. While both effects are of similar magnitude and thus tend to cancel out, it remains to stress that subjects do not fully correct their prior overreaction. On average 3.56% ($p < 0.01$) of the overreaction after a disconfirming bad signal and 3.36% ($p < 0.01$) of the overreaction after a disconfirming good signal persists. Nonetheless, subjects seem to follow the simple counting principle of rational information processing much closer and revise their prior beliefs consistently correctly – even if not fully in magnitude – when signals alternate.

To examine the second case of alternating signals, we bring the results for the price paths G-1 and G-2 (B-1 and B-2) into play. These results allow us to analyze subjects' updating behavior for alternating signals in situations without prior sequence of same-directed signals. For these price paths, the single opposite-directional signal occurs either directly in the first period (G-1 and B-1) or in the second period (G-2 and B-2) resulting in sequences of either two or three alternating signals, respectively. Figure 3 displays the estimated average overreaction in expectations (measured by the difference-in-difference of probability updates as shown in Equation 3) after (i) disconfirming signals without prior information sequence (Treatments G-1/B-1 and G-2/B-2); (ii) disconfirming signals in which prior signals alternate (i.e. the correction-scenario); and (iii) disconfirming signals which follow a prior streak of same-directed information.

[INSERT FIGURE 3 ABOUT HERE]

For both alternating information cases (i.e. (i) and (ii)) we find that individuals form expectations which correspond much closer to the counting rule. First, we find once more that individuals almost fully correct a prior overreaction when they observe that the information they received from a disconfirming signal is not followed up by further disconfirming evidence (by observing another confirming signal). Second, we find that when signals alternate right from the beginning (disconfirm w/o prior sequence), subjects are almost perfect Bayesian in the bad distribution and even underreact to disconfirming evidence in the good distribution. Part of this underreaction is driven by the fact that subjects are reluctant to significantly update their beliefs downwards after short streaks of negative information, consistent with the good news-bad news effect reported by Eil and Rao (2011) as well as Möbius et al. (2021).

Taken together, we can complement our main findings as follows: We find that a sequence of same-directed signals plays a crucial role when identifying situations in which subjects overreact to new information that disconfirms their priors. In particular, our results suggest that a prior streak of at least two same-directed information signal is necessary in order to observe a pronounced overreaction after subsequent disconfirming evidence. In contrast to that, we find that subjects adhere to the counting rule and thus process information in a rational way in situations with no prior sequence of same-directed signals and in situations with exclusively alternating information signals.

C. Reducing the Diagnosticity of Information Signals

In a next step, we examine whether individuals' belief updating after a disconfirming signal depends on the diagnosticity of the signal (i.e. its informational content). Prior studies have shown that the diagnosticity of information signals strongly corresponds to the degree that individuals over- and underinfer (e.g. Green et al., 1965; Kahneman and Tversky, 1972; Donnel and DuCharme, 1975; Griffin and Tversky, 1992; Grether, 1992; or Ambuehl and Li, 2018). Note that the updating paradigm we investigate explicitly abstracts from individuals' priors being consistent with Bayes. This also implies that the diagnosticity of the signal should not affect whether individuals correctly follow the principle that two equally diagnostic signals of opposite direction should cancel out (see Equation 3).

To test whether this is indeed the case, we conduct a variation of our baseline experiment in which we change the informational content that subjects can infer from signals. This means, we change the probability of the higher outcome in the good distribution from 70 % to 60 % and of the lower outcome from 30 % to 40 %, respectively. In the bad distribution, we change the probability of the lower outcome from 70 % to 60 % and of the higher outcome from 30 % to 40 %, respectively. On the one hand, we expected to observe – as Bayes' Theorem implies – lower (higher) absolute levels of probability estimates in the good (bad) distribution given the reduced diagnosticity of signals. On the other hand, we expect to observe no impact of diagnosticity on the fundamental counting rule we are testing. Within our empirical framework, the increase (decrease) in posterior probability after a confirming signal in the good (bad) distribution should remain exactly as much as the decrease (increase) in posterior probability after a subsequent disconfirming signal, irrespective of how informative the signal is. Besides that, our second experiment serves as a robustness test of our main findings from Experiment 1.

Overall, the findings look very similar to our baseline experiment. Table 4 compares subjects' changes in beliefs after a disconfirming signal in period *t* to their changes in beliefs after a same-directed signal that occurred in period *t-1*. For direct comparability, Table 4 reports changes in beliefs for the reduced diagnosticity experiment as well as for the baseline experiment. In particular, we find that subjects in the good distribution increase their prior beliefs by 4.43 % on average after a same-directed positive signal, whereas they decrease their prior beliefs by 13.75 % on average after a disconfirming negative signal which results in an overreaction of 9.32 %. Compared to the baseline experiment, the overreaction remains highly statistically significant and of similar magnitude. In the bad distribution, the findings look similar. Subjects decrease their prior beliefs by 5.78 % on average after a same-directed negative signal, while they increase their prior beliefs by 12.88 % on average after a disconfirming positive signal. The resulting overreaction is again highly statistically significant and similar to the baseline results.

[INSERT TABLE 4 ABOUT HERE]

Moreover, the design modification allows us to analyze how less extreme priors conditional on the same streak lengths as in our baseline experiment interact with the observed overreaction. In particular, we can test whether the overreaction critically depends on subjects having extreme priors as suggested by some papers in the literature (Benjamin, 2019). However, we do not find evidence that supports this notion. Instead, even with a diagnosticity of only 60%, two subsequent confirming signals are sufficient to observe a pronounced overreaction after a disconfirming signal. In such a situation not only the experimentally observed subjective priors, but also the objective Bayesian probabilities in the reduced diagnosticity experiment are low with on average 73% and 69%, respectively.

3.4 Exploring the Mechanism

In the previous sections we documented robust and systematic patterns of overreaction whenever a sequence of identical signals is interrupted by a signal which disconfirms prior evidence. Individuals appear to overweight signals which, given the observed outcome history, appear to be more extreme even though the data generating process suggests otherwise. Conversely, we observe that individuals are conservative whenever they observe a sequence of identical signals. In this next section, we aim to provide evidence for a specific mechanism behind the effect.

A. False Belief in Trends and Shifting Regimes

The paradoxical co-existence of both over- and underreaction to new information in financial markets is frequently associated with a combination of multiple behavioral biases. Barberis, Shleifer, and Vishny (1998) for example develop a model motivated by two well-documented systematic biases: conservatism and representativeness. The interaction of both biases causes investors to form false beliefs that the price development follows a regime-shifting framework, i.e. that prices either trend (or follow a "continuation" regime) or revert (or follow a "reversal" regime). Rabin (2002) as well as Rabin and Vayanos (2010) model the interaction between the "gambler's fallacy" and the "hot hand bias", which results in individuals to overestimate the probability that a short streak reverses and that a long streak of identical signals continues.

Both a mistaken belief in mean reversion as well as a mistaken belief in a continued trend can be traced back to individuals' uncertainty about the underlying data-generated process (or the distribution from which outcomes are drawn in our experiment) and the extent to which they anticipate a signal that disconfirms prior evidence (see for example Augenblick & Rabin, 2021). To cleanly test for such a mechanism in our data, we conduct another experimental study with N = 604 participants. The experiment again builds on our baseline experiment but manipulates the information that individuals posses both about the observed outcomes but also about the underlying distribution. To do so, we change the previously framed forward-looking updating task to a backward-looking updating task. Whereas subjects in the baseline experiment are asked to make a forecasting decision without knowing the future outcome history, we now show subjects the full outcome history beforehand. Then, we ask them to provide probability estimates period by period as in the baseline experiment for exactly the same outcome history they have seen in advance. Importantly, subjects were still incentivized to provide probability forecasts which only incorporate the information subjects had in each period. In other words, the objective Bayesian probabilities are identical to the baseline experiment. By showing subjects the entire outcome history beforehand, we eliminate most of the uncertainty regarding the underlying distribution and any of the potential surprise related to the period in which the disconfirming signal occurs. As such, it would be difficult to argue that subjects wrongly infer trends or reversals as they know the outcomes that follow each signal in advance. To ensure that subjects correctly remember the outcomes history, we ask them two additional questions before the first period: (i) count the number of price increases and decreases in the outcome history and (ii) state the period in which the disconfirming signal occurs.

To investigate whether a perfect foresight of future outcomes causally reduces individuals' persistent overreaction towards signals which disconfirm prior evidence, we once again compare the difference in their belief update after a disconfirming signal to the difference in their belief update after the most recent confirming signal across experiments. Note that the intrinsic property of Bayesian updating predicts that belief changes should be of identical magnitude such that any observed difference can be attributed to the different properties of the experimental environment. Results are reported in Table 5.

[INSERT TABLE 5 ABOUT HERE]

Results in Table 5 reveal that the difference of the difference in belief updates is roughly equivalent across our baseline experiment and the perfect foresight experiment. Whereas we observe a slightly weaker overreaction towards negative signals ($p < 0.05$), the overreaction after positive signals is of similar magnitude and statistically undistinguishable across experiments ($p > 0.1$). Overall, this first test suggests that while a false belief in trends or price reversals might add to the bias, it cannot account for the over- and underreaction patterns we observe.

Next, we investigate how the overreaction in the formation of expectations evolves as a function of the streak length of prior identical signals. If the overreaction we observe is driven by a false belief in trends or mean-reversion, we should observe an interaction between the streak length of prior identical signals and the overreaction after the streak is interrupted by a signal which disconfirms prior evidence. For instance, a key implication of the Barberis, Shleifer, and Vishny (1998) model is that the probability of a streak continuing increases monotonically with streak length, which means that a decision maker exhibits a dampened overreaction after a longer streak of identical signal is interrupted. In contrast, the model by Rabin (2002) predicts a non-monotonic relationship. In the model, a decision maker initially decreases the probability of continuation for shorter streaks and then switches to increasing the probability of continuation at longer streak lengths (see Asparouhova et al., 2009). To investigate this relationship, we plot the observed overreaction (again measured by the difference-in-difference of probability updates) as a function of the streak of prior identical signals. Results are reported in Figure 4.

[INSERT FIGURE 4 ABOUT HERE]

Figure 4 shows that the overreaction we document is mostly independent of the streak length. Consistent with our prior conjecture, we once again observe that a sequence of at least two same-directed signals is necessary to observe a pronounced overreaction. Afterwards however, there is little interaction between the streak length and the absolute overreaction. Additionally, Figure 4 confirms that the overreaction is rather symmetric across distributions, which is inconsistent with behavioral models of motivated beliefs (Brunnermeier & Parker, 2005; Kunda, 1990), which predict asymmetric updating and overall optimism. Overall, both tests individually and jointly do not support the hypothesis that the systematic patterns of over- and underreaction we document are driven by a false belief in trends or mean-reversion.

B. Misallocation of Cognitive Resources

An alternative channel through which the simultaneous co-existence of both over- and underreaction operates is motivated by prior findings from cognitive psychology. Over the years, many systematic deviations from Bayesian beliefs in sequential belief updating have been document. Among the most robust findings are that individuals are conservative (e.g. Edwards, 1982), or that they pay too little attention to base-rates (e.g. Bar-Hillel, 1980). At the same time, individuals appear to overweight the extremeness of signals (e.g. Griffin and Tversky, 1992). In some instances, the biases tend to push in different directions which produces predictions that are at odds with one another. One prominent way to jointly rationalize both over- and underreaction to new information is by introducing mental updating costs (Kominers, Mu, & Peysakhovich, 2017). In this framework, individuals face a mental updating cost which they pay every time they update their beliefs. In the light of such an updating cost, individuals need to selectively allocate their cognitive resources such that they only fully update when they deem a signal "valuable" and ignore "invaluable" signals. Both Kominers, Mu, and Peysakhovich (2017) as well as Gertsman (2021) recently proposed that individuals condition the updating decision on the realization of the signal itself. As such, individuals readily update their expectations when the new information is extreme relative to the prior since it indicates that their beliefs are not in line with reality (consistent with Griffin & Tversky, 1992). Conversely, individuals tend to underweight signals which are not extreme relative to the prior since it indicates that their beliefs are still a solid representation of reality. Importantly, under the mechanism we consider, rather than affecting how information is processed (as in models of motivated beliefs or misattribution) or using a misspecified model (e.g. Benjamin et al., 2016), agents face updating costs to pay sufficient attention to the signals they observe. Prior work in cognitive psychology has shown that attention and information processing are closely related (e.g. Smith & Krajbich, 2019; Enax et al., 2016). In our setting, individuals might underweight information which is consistent with their priors since they deem their beliefs already as *good enough*. Information which disconfirms their prior is – however – more valuable to process as it indicates that they might have wrong expectations about an underlying state of the world.

To test this hypothesis, we require evidence for the following conjectures: (1) individuals take more time to incorporate disconfirming signals into their expectations; and (2) the overreaction is a positive function of the time that individuals take to process the disconfirming signal. To test for this mechanism, we conduct a variation of our baseline experiment and combine it with a change detection task from cognitive psychology ($N = 254$) to both measure their response time but also to reduce the perceived repetitiveness in observing streaks of identical outcomes. In addition to the instructions of the baseline experiment, participants were told that in each round, the announced signal would randomly be highlighted in one of three colors. Besides reporting their beliefs, participants were incentivized to correctly identify how many outcomes were displayed in one particular color. To make the change detection task sufficiently salient, the payment they could earn by identifying the correct number of colored outcomes represented a significant bonus to their overall earnings (approximately 20% of their earnings). Overall, the change detection task we use is similar to those in cognitive psychology to examine the allocation of cognitive resources (e.g. Verghese, 2001; Mrkva et al., 2019). However, we did not limit or incentivize the time they take to make a decision to keep the experiment comparable to the baseline experiment.

Consistent with our prior conjecture, we find that individuals take significantly more time to process signals which disconfirm prior evidence. Whereas subjects in our experiment require on average 9.6 seconds to update their expectations after confirming evidence, they require roughly 14.5 seconds on average after observing disconfirming evidence ($p < 0.001$, Wilcoxon signed-rank test). Figure 5 illustrates this relation in greater detail, by displaying the time individuals require to update their expectations after both confirming and disconfirming signals for each round separately. Figure 5 captures two related effects. First, we observe once more that individuals take significantly more time to process disconfirming evidence relative to confirming evidence independent of the round in which the signal occurred. Second, we find that the longer the streak of confirming signals, the faster individuals tend to process subsequent confirming signals. Conversely, the longer the streak of confirming signals before individuals observe a signal which disconfirms prior evidence, the longer they take to process the disconfirming signal. This finding lends further credence to the hypothesis that individuals face mental updating costs when processing a signal. Repeated signals which confirm prior evidence become increasingly "invaluable" for the decision maker as they consider their beliefs as *good enough*. Consequently, they spend increasingly less time to process such signals. Conversely, once a signal disconfirms a prior sequence of evidence, individuals realize that their beliefs might not be in line with reality and devote significantly more time to process it, thereby overreacting to the signal. Importantly, the longer the prior streak of consistent evidence, the more extreme the disconfirming signal occurs, and the longer individuals need to process the signal.

[INSERT FIGURE 5 ABOUT HERE]

In a next step, we investigate whether the magnitude of the observed overreaction is a function of individuals' response time. The aggregate overreaction (measured by the difference-in-difference in belief updates) in the change detection experiment was on average 4.19 %, which is lower compared to the aggregate overreaction in the baseline experiment, which was on average 8.75 % ($p < 0.01$, Wilcoxon signed-rank test). The difference in the overreaction between experiments does not come unexpected, as the change detection experiment was in part designed to raise participants' awareness and to reduce the repetitive nature of a streak of confirming evidence (by assigning identical signals different colors which participants need to correctly identify). In Table 6, we investigate whether the overreaction is a function of participants' response time when they observe the disconfirming signal. To do this, we classify participants' response time into terciles (low, medium, and high) and compute the average overreaction for each response time tercile. Consistent with the hypothesis that participants face mental costs for updating their expectations, we observe that – independent of the distribution – the overreaction is a positive function of individuals' response time when they observe the disconfirming signal. In particular, higher response times – which indicate a greater allocation of cognitive resources – lead to substantially greater overreactions when a sequence of identical signals is interrupted by a signal which disconfirms prior evidence. This finding is also consistent with the 'more is less' hypothesis of Dawes (1979), who conjectures that greater attention may lead forecasters to overweight properties of the decision problem.

[INSERT TABLE 6 ABOUT HERE]

4. From the Lab to the Field

Results from the laboratory show a substantial discrepancy between how people should incorporate information signals which disconfirm a prior streak of same-directed information and how they actually update their beliefs in these settings. Based on this insight, we turn to a real-case application in financial markets, that is, stock return reactions to earnings announcements, and test whether the sequence of earnings surprises affects asset prices. In particular, we examine how the direction of a firm's prior earnings surprises – same- or opposite-directed – relative to a firm's current surprise affects announcement and post-announcement stock returns. Quarterly earnings announcements represent the major periodic source of news released by publicly traded U.S. firms. Investors and financial analysts form expectations about the firm's earnings prior to an announcement. On the announcement day, actual earnings may deviate from those expectations leading to earnings surprises. These earnings surprises are related to stock price movements since they present new information which investors incorporate when forming future expectations of firm prospects. In the following, we examine whether this process of expectation updating is affected not only by the most recent, new information, but also by how it relates to prior information that was released by the same company. In particular, we argue – based on our experimental results – that directional changes in the sign of information signals may lead to overreactions which applied to the case of earnings expectation updating may distort market reactions to firm earnings announcements. In line with the updating behavior observed in the laboratory, we expect to observe a stronger announcement stock return reaction if the earnings surprise today is opposite-directed to the previous earnings surprises than if it is same-directed. Moreover, and given the stronger initial reaction, we assume a weaker post-announcement drift following opposite-directed earnings surprises.

4.1 Data and Methodology

We obtain data for daily stock returns, market value-weighted index returns, daily stock prices (adjusted for splits and dividends), and daily trading volumes from the CRSP-Compustat merged database from June 2009 to May 2020. Additionally, we use quarterly earnings announcement data and analyst earnings forecasts from the union of the CRSP-Compustat merged database and I/B/E/S for the same period. To clearly identify the earnings announcement dates, we compare the earnings announcements available from the CRSP-Compustat merged database with the dates available at I/B/E/S. When the announcement dates differ, we follow DellaVigna and Pollet (2009) and take the earlier date. Our primary sample of 53,168 firm-quarter observations consists of all firm quarters from June 2009 to May 2020 and 1,456 unique firms with sufficient data to estimate the required variables (i.e. current earnings surprise and future abnormal returns) as described below.

The signals we are investigating are firms' quarterly earnings surprises. Following Hirshleifer et al. (2009), we measure an earnings surprise as the difference between announced earnings as reported by I/B/E/S and the consensus earnings forecast, defined as the median of the most recent forecasts from individual analysts using the I/B/E/S detail tape. We normalize the difference by the stock price at the end of the corresponding quarter.

In line with our experimental analyses, we classify earnings surprises in either confirming or contrary signals, depending on how the sign of the most recent earnings surprise compares to the sign of the previous earnings surprises. More precisely – and in analog to our experimental setup – an earnings surprise is defined as a contrary signal if its sign is different to the sign of the previous two earnings surprises (e.g. at least two subsequent positive earnings surprises are followed by a negative earnings surprise). Similarly, we define an earnings surprise as a confirming signal if its sign is the same as the sign of the previous two earnings surprises (e.g. a at least two subsequent positive earnings surprises are followed by another positive earnings surprise).[11]

Our dependent variable is the cumulative abnormal return over two windows, the announcement window [0,1] and the post-announcement window [2,61] measured in trading days relative to the announcement date. We follow the methodology of Hirshleifer et al. (2009) and define the cumulative abnormal return as the difference between the buy-and-hold return of the announcing firm in quarter q and that of one of 25 size and book-to-market (B/M) matching portfolios in quarter q,

[11] We chose this classification to remain as close as possible to our experimental framework. However, we also consider alternative specifications in which we only condition on the last earnings surprise which either gets interrupted (contrary) or followed by a subsequent same-directed signal (confirming). Results are consistent and reported in Appendix Figure B.2 and Table B.2.

$$CAR[0,1]_{iq} = \prod_{k=t}^{t+1}(1 + r_{ik}) - \prod_{k=t}^{t+1}(1 + r_{pk})$$

$$CAR[2,61]_{iq} = \prod_{k=t+2}^{t+61}(1 + r_{ik}) - \prod_{k=t+2}^{t+61}(1 + r_{pk}),$$

where r_{ik} is the return of the firm *i* and r_{pk} is the return of the matching size-B/M portfolio on day *k*, where *t* is the announcement date of quarter *q's* earnings.[12]

4.2 Empirical Results

A. Descriptive Statistics

Before we analyze stock return reactions to same- and opposite-directed earnings surprises, we provide summary statistics of how our new measure interacts with firm characteristics. Table 7 reports descriptive statistics of the firms that announce contrary earnings surprises versus the firms that announce confirming earnings surprises. We later employ the firm characteristic variables as controls in regression analyses. We define firm characteristic variables as in Hirshleifer et al. (2009). The size and B/M ratios are calculated at the end of June of each year based on the market value of equity at the end of June and the book value of equity for the last fiscal year-end in the previous calendar year divided by the market value of equity for December of the previous calendar year. Earnings Persistence is the first-order autocorrelation coefficient of quarterly earnings per share during the past four years (split-adjusted; minimum four observations required). Earnings Volatility is the standard deviation during the preceding four years of the deviations of quarterly earnings from 1-year-ago earnings (split-adjusted; minimum four observations required). Share Turnover is defined as the average monthly share trading volume divided by the average number of shares outstanding during a 1-year period ending as the end of the corresponding fiscal quarter.

[INSERT TABLE 7 ABOUT HERE]

Overall, our sample consists of roughly 33% contrary and 67% confirming earnings surprise observations. We find that contrary earnings surprises are rather from larger firms with

[12] The results in the paper are similar if we instead use a market model to calculate buy-and-hold abnormal returns. Results for the alternative specification are available upon request.

lower B/M ratios, lower earnings persistence, less shares outstanding, and a slightly smaller number of covering analysts. There is no significant difference in earnings volatility or share turnover between contrary and confirming earnings surprises.

B. Contrary and Confirming Earnings Surprises and Stock Return Reactions

We first perform univariate analyses to examine the effect of contrary versus confirming earnings surprises on (i) announcement date returns and (ii) post-earnings announcement date returns. In each quarter, we perform a two-way independent sort of quarterly earnings surprises in that quarter into 2 x 5 = 10 groups based on the sign of the earnings surprise relative to the previous surprise (i.e. contrary versus confirming) and the magnitude of the earnings surprise (Quintile 1: most negative earnings surprises, Quintile 5: most positive earnings surprises).[13] For each group of confirming and contrary earnings surprises, we calculate the average cumulative abnormal return over the announcement and the post-announcement window for the most negative earnings surprise quintile and the most positive earnings surprise quintile and the difference in cumulative abnormal returns between the two extreme earnings surprise quintiles.

The difference in abnormal announcement day returns between earnings surprise quintile 5 and 1 measures the stock price reaction to earnings news, whereby a larger difference indicates that investors react more strongly to earnings news on the announcement day. The difference in post-announcement abnormal returns between earnings surprise quintile 5 and 1 measures underreaction to earnings news as reflected in a subsequent drift. Our experimental results predict a stronger CAR[0,1] difference (stronger announcement-day reaction) and a weaker CAR[2,61] difference (weaker post-earnings announcement drift) for contrary earnings surprises compared to confirming earnings surprises.

[INSERT FIGURE 6 ABOUT HERE]

Figure 6 illustrates the development of the average cumulative abnormal returns over 61 days after the earnings announcement for the top and the bottom earnings surprise quintile conditional on whether the earnings surprise was confirming (i.e. same-directed relative to the previous two earnings surprises of the firm) or contrary (i.e. opposite-directed relative to the

[13] We also perform sorts based on earnings surprise deciles instead of quintiles. The results are similar and reported in Figure B.3 in the Appendix.

previous two earnings surprises of the firm). In line with the literature, we observe that stock prices adjust strongly in the direction of the earnings surprise immediately after the announcement. That is, after a positive earnings surprise stock prices increase, whereas after a negative earnings surprise they decrease. However, Figure 6 shows that this initial announcement reaction depends not only on the magnitude of the announcement-day earnings surprise, but also on its sign relative to the previous earnings surprise. Consistent with our experimental findings, we find that initial announcement stock return reactions are significantly stronger for opposite-directed earnings surprises than for same-directed earnings surprises. That is, after *positive* earnings surprises, we observe on average stronger return reactions if the surprise follows prior *negative* earnings surprises, compared to when it follows prior *positive* earnings surprises. The same holds true for negative earnings surprises that interrupt a streak of prior positive earnings surprises. In that case, the initial stock return reaction is more than twice as strong as if a similar negative earnings surprise follows prior same-directed earnings surprises.

Going forward, Figure 6 demonstrates that this larger initial reaction of opposite-directed earnings surprises only partly reverts in the long-run. In particular, we observe for the top earnings surprise quintile that the cumulative abnormal return following the opposite-directed earnings surprise converges to the cumulative abnormal return following the confirming earnings surprise. Approximately 30 trading days after the announcement, they become statistically insignificant from one another. While stocks in both, the contrary and confirming top surprise quintile clearly drift upwards over the 60 days following the announcement date, the drift is stronger for those stocks that belong to the confirming earnings surprise quintile, consistent with an underreaction that is eventually corrected. For negative surprises, we find a different pattern. In particular, we cannot observe a drift in the direction of the surprise. Instead, prices of stocks with negative same-directed or negative opposite-directed earnings surprises remain on average relatively stable. Consequently, the initial difference in stock return reactions between opposite-directed and same-directed surprises in the bottom surprise quintile carries forward and persists even 60 days after the initial earnings announcement.

[INSERT TABLE 8 ABOUT HERE]

In Table 8, we provide a more detailed overview of the announcement and post-announcement abnormal returns following confirming and disconfirming earnings surprises. Consistent with our prior conjecture, we show that investors' 2-day announcement reactions to earnings surprises are more sensitive following announcements that are of opposite direction. While on average cumulative abnormal returns for the top surprise quintile are around 3.05% if they follow a previous same-directed earnings surprise, they are around 4.22% if they follow a previous opposite-directed earnings surprise, resulting in a difference in abnormal returns of 1.17% ($p < 0.01$). For the bottom surprise quintile, the difference is even larger with 2.93% ($p < 0.01$). Announcement abnormal returns for confirming earnings surprises which display the same sign as the previous two earnings surprises are around -1.39%, whereas they are around -4.86% if they are contrary to the previous earnings surprise. This indicates that the stock price reactions to earnings announcements are stronger when they are of opposite direction to the prior announcement than when they are of same direction.

For post-announcement abnormal returns, we observe a weak or even absent post-earnings announcement drift, measured as the difference between Quintile 5 and 1 if the earnings surprise is contrary to the previous quarter earnings surprise. However, for confirming earnings surprises, we find a positive post-earnings announcement drift. In this case, the difference in post-announcement abnormal returns between the top and bottom surprise quintile is 1.27% ($p < 0.05$). Comparing the post-announcement returns within quintiles between contrary and confirming surprises, we show that the effect is driven by a stronger upward drift for same-directed positive surprises than for opposite-directed positive surprises.

To back up these findings, we run multivariate regression analyses on the entire sample and a subsample containing only observations in the top and bottom earnings surprise quintile. These analyses allow us to control for other possible sources of variation in the relation between announcement date returns and earnings news, and between post-announcement returns and earnings news. In separate analyses, we regress the CAR[0,1] and the CAR[2,61] on the earnings surprise quintile rank (FE), a contrary dummy which is one if the earnings surprise is of opposite direction to the prior surprise, the interaction between the two variables, and controls, also interacted with FE. Previous studies show that investor reactions to earnings news vary with firm size, book-to-market, earnings persistence, earnings volatility, number of covering analysts, and share turnover (Della Vigna & Pollet, 2009; Hirshleifer et al., 2009). Thus, we include all of these variables in our regression model. Table 9 reports the results. Consistent with our results in Table 8, we confirm that the initial stock return reaction is significantly stronger for contrary surprises than for confirming surprises. This holds after controlling for conventional variables in the announcement drift literature. For the post-announcement return, we find no significant overall drift. However, the results suggest a tendency for a stronger drift for same-directed earnings surprises.

[INSERT TABLE 9 ABOUT HERE]

Overall, our findings are in line with the under- and overreaction patterns observed in our experimental data. Stock market reactions around the announcement date are stronger for opposite-directed than for same-directed earnings surprises. However, in the long run, it turns out that investors initially underreact after positive earnings surprises. This underreaction leads to price drifts as commonly observed in the post-earnings announcement literature. Our results suggest that drifts are particularly pronounced for confirming positive surprises, that is, if a positive surprise follows a prior positive surprise. Thus, abnormal returns of confirming and contrary surprise stocks converge over the subsequent 60 days. This convergence for the top surprise quintile is driven by the confirming surprise stocks that drift stronger in the *same* direction of surprise (i.e. correcting the initial *underreaction*). In contrast to that, after negative earnings surprises, we cannot observe a clear downward price drift. Instead, the initial difference between same- and opposite-directed earnings surprises persists even 60 days after the announcement.

5. Discussion and Conclusion

The paradoxical co-existence of both over- and underreaction to new information continues to be a major challenge for understanding individual behavior (Benjamin, 2019) as well as for aggregate market behavior and its implications for asset prices (e.g. Barberis et al., 1998; Daniel et al., 1998; Hong and Stein, 1999; Barberis et al., 2015, 2018). In this study, we investigate individuals' belief updating behavior through the lens of a simple but important property of rational information processing: two equally-diagnostic signals of opposite direction should cancel out. This allows us to judge *belief movement* solely on the intrinsic properties of Bayesian updating while remaining agnostic about the rationality of *absolute beliefs*.

Utilizing evidence from the lab and field, we show that individuals do not always follow this counting-based principle of information processing. Instead of balancing out two equally-diagnostic but opposite pieces of information, individuals strongly overweight evidence which is inconsistent with prior information. Conversely, we find that individuals generally underinfer in situations in which they cannot rely on this principle and in which they would need to know Bayes' rule (e.g. by observing consecutive pieces of same-directed information). As potential mechanism, we identify that individuals update their expectations as if they must selectively allocate their cognitive resources. Successive information signals are increasingly underweighted as individuals deem their beliefs as *good enough*. Information which disconfirms prior beliefs is – however – treated as more valuable to process as it indicates that their beliefs might not be in line with reality.

Our findings contribute to a more comprehensive understanding of the co-existence of over- and underreaction to new information both at the individual but also at the market level. Instead of relying on an interaction of multiple behavioral biases, we empirically show that both over- and underreaction can be jointly rationalized when considering that computing posterior beliefs is (mentally) costly for individuals. Our findings leave open a number of interesting questions regarding how the documented phenomenon interacts with other contextual domains. While we show that market reactions to earnings surprises display similar patterns as documented in the laboratory, earnings announcements are not the only source of new information in financial markets. An important focus for future research should be to further the understanding of how psychological biases found in the lab manifest in real-world settings.

References

Ambuehl, S., & Li, S. (2018). Belief updating and the demand for information. *Games and Economic Behavior*, 109, 21-39.

Asparouhova, E., Hertzel, M., & Lemmon, M. (2009). Inference from streaks in random outcomes: Experimental evidence on beliefs in regime shifting and the law of small numbers. *Management Science*, 55(11), 1766-1782.

Augenblick, N., & Rabin, M. (2021). Belief movement, uncertainty reduction, and rational updating. *The Quarterly Journal of Economics*, 136(2), 933-985.

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

Barberis, N., Greenwood, R., Jin, L., & Shleifer, A. (2018). Extrapolation and bubbles. *Journal of Financial Economics*, 129(2), 203-227.

Barberis, N., Shleifer, A., & Vishny, R. (1998). A model of investor sentiment. *Journal of Financial Economics*, 49(3), 307-343.

Bar-Hillel, M. (1980). The base-rate fallacy in probability judgments. *Acta Psychologica*, 44(3), 211-233.

Benjamin, D. J. (2019). Errors in probabilistic reasoning and judgment biases. *In Handbook of Behavioral Economics: Applications and Foundations 1* (Vol. 2, pp. 69-186). North-Holland.

Bernard, V. L., & Thomas, J. K. (1989). Post-earnings-announcement drift: delayed price response or risk premium?. *Journal of Accounting Research*, 27, 1-36.

Bernard, V. L., & Thomas, J. K. (1990). Evidence that stock prices do not fully reflect the implications of current earnings for future earnings. *Journal of Accounting and Economics*, 13(4), 305-340.

Brunnermeier, M. K., & Parker, J. A. (2005). Optimal expectations. *American Economic Review*, 95(4), 1092-1118.

Buhrmester, M., Kwang, T., & Gosling, S. D. (2011). Amazon's Mechanical Turk: A new source of inexpensive, yet high-quality, data?. *Perspectives on Psychological Science*, *6*(1), 3-5.

Camerer, C. F. (1987). Do biases in probability judgment matter in markets? Experimental evidence. *American Economic Review*, 77(5), 981-997.

Camerer, C. (1995). Individual decision making. *Handbook of Experimental Economics.*

Charness, G., & Dave, C. (2017). Confirmation bias with motivated beliefs. *Games and Economic Behavior*, 104, 1–23.

Daniel, K., Hirshleifer, D., & Subrahmanyam, A. (1998). Investor psychology and security market under-and overreactions. *The Journal of Finance*, 53(6), 1839-1885.

Dave, C., & Wolfe, K. (2003). On Confirmation Bias and Deviations from Bayesian Updating. Working Paper.

Dawes, R. M. (1979). The robust beauty of improper linear models in decision making. *American Psychologist*, 34(7), 571.

De Bondt, W. F., & Thaler, R. (1985). Does the stock market overreact? *The Journal of Finance*, 40(3), 793-805.

DellaVigna, S., & Pollet, J. M. (2009). Investor inattention and Friday earnings announcements. *The Journal of Finance*, 64(2), 709-749.

Donnell, M. L., & Du Charme, W. M. (1975). The effect of Bayesian feedback on learning in an odds estimation task. *Organizational Behavior and Human Performance*, 14(3), 305-313.

DuCharme, W., & Peterson, C. (1968). Intuitive inference about normally distributed populations. *Journal of Experimental Psychology*, 78 (2), 269–275.

Edwards, W. (1982). Conservatism in human information processing. In A. T. Daniel Kahneman, Paul Slovic (Ed.), *Judgment under Uncertainty: Heuristics and Biases*, pp. 359-369. Cambridge University Press.

Eil, D., & Rao, J. M. (2011). The good news-bad news effect: asymmetric processing of objective information about yourself. *American Economic Journal: Microeconomics*, 3(2), 114-38.

Enax, L., Krajbich, I., & Weber, B. (2016). Salient nutrition labels increase the integration of health attributes in food decision-making. *Judgment & Decision Making*, 11(5).

Enke, B., & Graeber, T. (2019). *Cognitive Uncertainty* (No. w26518). National Bureau of Economic Research.

Geller, E.S., & Pitz, G.F. (1968). Confidence and decision speed in the revision of opinion. *Organizational Behavior and Human Performance*, 3 (2), 190–201.

Gertsman, G. (2021). Selective Learning and Price Over- and Under-reaction. *Working Paper*.

Goodman, J. K., Cryder, C. E., & Cheema, A. (2013). Data collection in a flat world: The strengths and weaknesses of Mechanical Turk samples. *Journal of Behavioral Decision Making*, *26*(3), 213-224.

Green, P. E., Halbert, M. H., & Robinson, P. J. (1965). An experiment in probability estimation. *Journal of Marketing Research*, 2(3), 266-273.

Grether, D. M. (1980). Bayes rule as a descriptive model: The representativeness heuristic. *The Quarterly Journal of Economics*, 95(3), 537-557.

Grether, D.M. (1992). Testing Bayes rule and the representativeness heuristic: some experimental evidence. *Journal of Economic Behavior and Organization*, 17 (1), 31–57.

Griffin, D., & Tversky, A. (1992). The weighing of evidence and the determinants of confidence. *Cognitive Psychology*, 24(3), 411-435.

Hartzmark, S. M., Hirshman, S. D., & Imas, A. (2021). Ownership, learning, and beliefs. *The Quarterly Journal of Economics*, 136(3), 1665-1717.

Hartzmark, S. M., & Shue, K. (2018). A tough act to follow: Contrast effects in financial markets. *The Journal of Finance*, 73(4), 1567-1613.

Hirshleifer, D., Lim, S. S., & Teoh, S. H. (2009). Driven to distraction: Extraneous events and underreaction to earnings news. *The Journal of Finance*, 64(5), 2289-2325.

Hong, H., & Stein, J. C. (1999). A unified theory of underreaction, momentum trading, and overreaction in asset markets. *The Journal of Finance*, 54(6), 2143-2184.

Hung, M., Li, X., & Wang, S. (2015). Post-earnings-announcement drift in global markets: Evidence from an information shock. *The Review of Financial Studies*, 28(4), 1242-1283.

Kahneman, D., & Tversky, A. (1972). Subjective probability: A judgment of representativeness. *Cognitive Psychology*, 3(3), 430-454.

Kominers, S. D., Mu, X., & Peysakhovich, A. (2018). Paying (for) attention: The impact of information processing costs on bayesian inference. *Working Paper*.

Kuhnen, C. M. (2015). Asymmetric learning from financial information. *The Journal of Finance*, 70(5), 2029-2062.

Kunda, Z. (1990). The case for motivated reasoning. *Psychological Bulletin*, 108(3), 480.

Levitt, S. D., & List, J. A. (2007a). On the generalizability of lab behaviour to the field. *Canadian Journal of Economics/Revue canadienne d'économique*, 40(2), 347-370.

Levitt, S. D., & List, J. A. (2007b). What do laboratory experiments measuring social preferences reveal about the real world?. *Journal of Economic perspectives*, 21(2), 153-174.

Martineau, C. (2019). The evolution of financial market efficiency: Evidence from earnings announcements. *Available at SSRN 3111607*.

Möbius, M. M., Niederle, M., Niehaus, P., & Rosenblat, T. S. (2021). Managing Self-Confidence: Theory and Experimental Evidence. *Management Science*, forthcoming.

Mrkva, K., Westfall, J., & Van Boven, L. (2019). Attention drives emotion: Voluntary visual attention increases perceived emotional intensity. *Psychological Science,* 30(6), 942-954.

Peterson, C.R., & Miller, A.J. (1965). Sensitivity of subjective probability revision. *Journal of Experimental Psychology*, 70 (1), 117–121.

Peterson, C.R., Schneider, R.J., & Miller, A.J. (1965). Sample size and the revision of subjective probabilities. *Journal of Experimental Psychology*, 69 (5), 522–527.

Pitz, G.F. (1969). The influence of prior probabilities on information seeking and decision-making. *Organizational Behavior and Human Performance*, 4 (3), 213–226.

Pitz, G.F., Downing, L., & Reinhold, H. (1967). Sequential effects in the revision of subjective probabilities. *Canadian Journal of Psychology*, 21 (5), 381–393.

Rabin, M. (2002). Inference by believers in the law of small numbers. *The Quarterly Journal of Economics*, 117(3), 775-816.

Rabin, M., & Vayanos, D. (2010). The gambler's and hot-hand fallacies: Theory and applications. *The Review of Economic Studies*, 77(2), 730-778.

Richardson, S., Tuna, I., & Wysocki, P. (2010). Accounting anomalies and fundamental analysis: A review of recent research advances. *Journal of Accounting and Economics*, 50(2-3), 410-454.

Sasaki, S., & Kawagoe, T. (2007). Belief Updating in Individual and Social Learning: A Field Experiment on the Internet. Discussion Paper 690. The Institute of Social and Economic Research, Osaka University.

Smith, S. M., & Krajbich, I. (2018). Attention and choice across domains. *Journal of Experimental Psychology: General*, 147(12), 1810.

Verghese, P. (2001). Visual search and attention: A signal detection theory approach. *Neuron*, 31(4), 523-535.

Figure 1: Empirical Framework

The figure illustrates the empirical framework of this study with an example. It shows how the posterior probability should evolve in a setting in which an agent learns from good and bad, equally-diagnostic signals about one of two underlying distributions. The sequence of signals is good, good, and bad, leading to an overall sequence of two same-directed signals that gets interrupted by a disconfirming signal. The blue dots present the objective probabilities (i.e. the beliefs according to Bayes' Theorem) that the asset is in the good state. The dashed lines demonstrate the counting rule implied by rational information processing which is that two equally-diagnostic signals of opposite direction should cancel out.

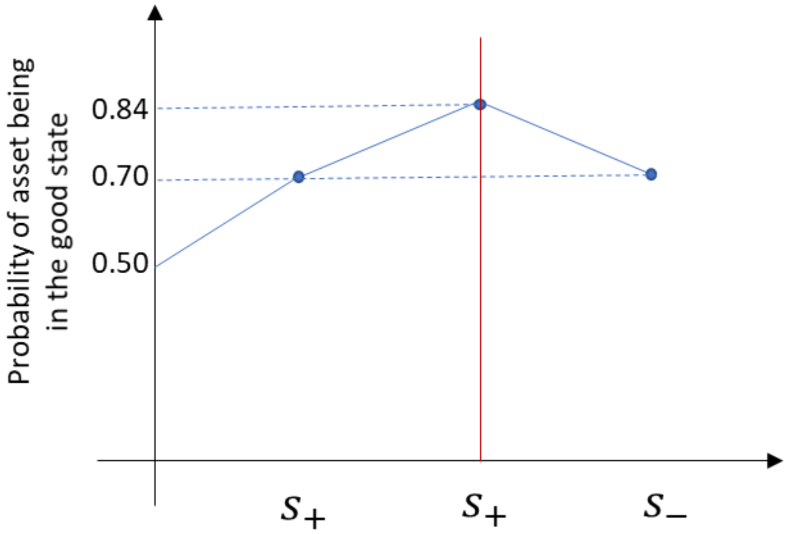

Figure 2: Subjects' Updating Behavior for Opposite- and Same-Directed Signals – Experiment 1

The figure displays in blue bars subjects' average absolute changes in probability estimates after a disconfirming signal in period t and after a same-directed signal in period t-1. The average absolute changes in objective Bayesian posterior probabilities are displayed in red bars. Results are shown individually for the good and the bad distribution. Displayed are 95% confidence intervals.

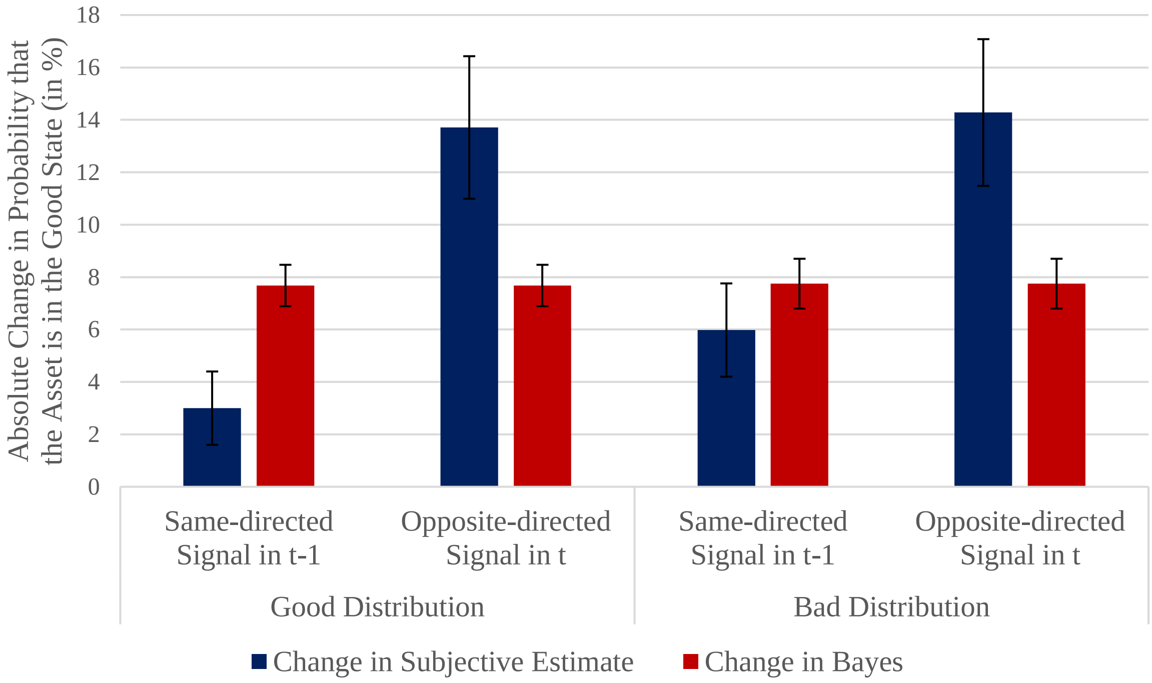

Figure 3: Subjects' Updating Behavior for Alternating Signals – Experiment 1

The figure shows participants' absolute average overreaction in belief updating after disconfirming signals which (i) have no prior sequence of same directed evidence, (ii) follow alternating information, and (iii) follow a prior sequence of same directed evidence split by the underlying distribution. Reported are both subjective beliefs (in red) and Bayesian beliefs (in green). Displayed are 95% confidence intervals.

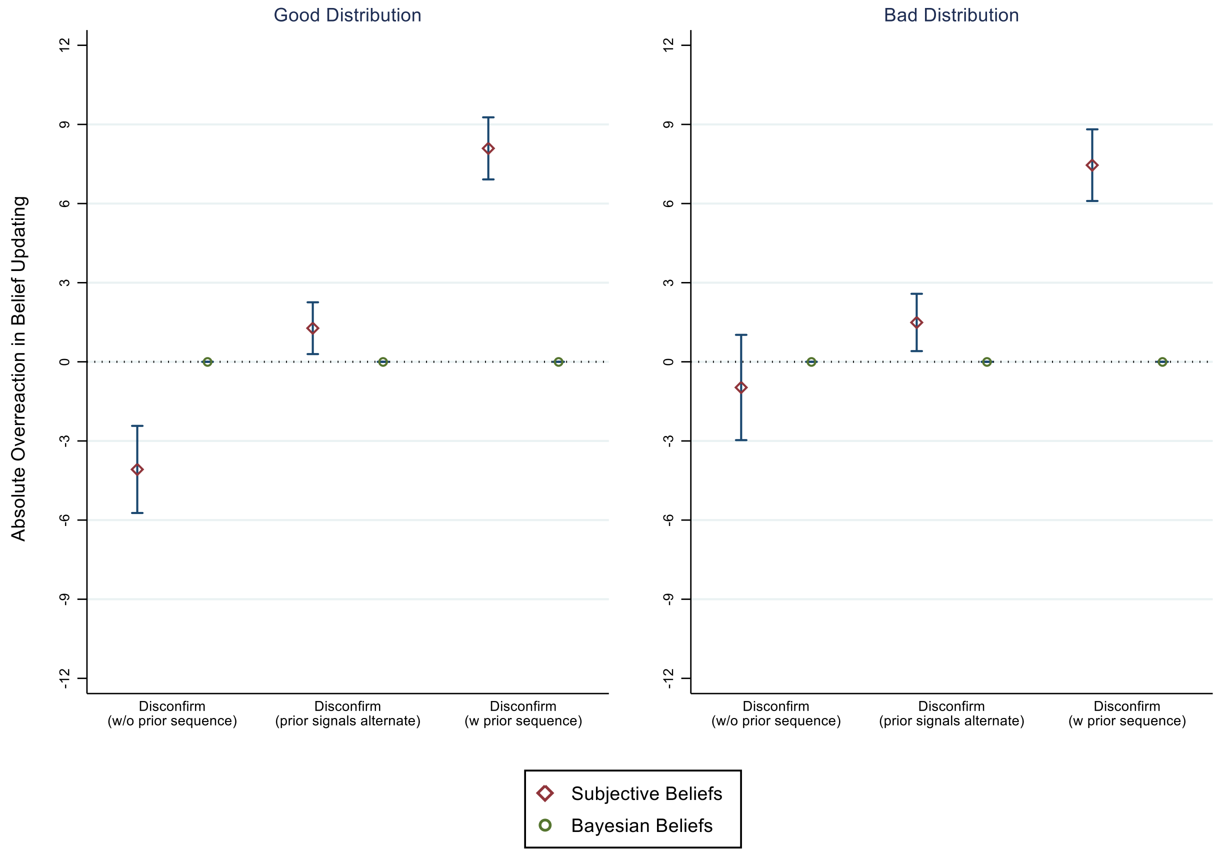

Figure 4: Overreaction and Length of Sequence of Same-Directed Signals

This figure displays participants' overreaction (in %) as a function of the streak length of prior same-directed signals. The dashed line displays results for the bad distributions (i.e. good signal interrupts a streak of bad signals) and the solid line represents results for the good distribution (i.e. bad signal interrupts a streak of good signals). Displayed are 90% confidence intervals.

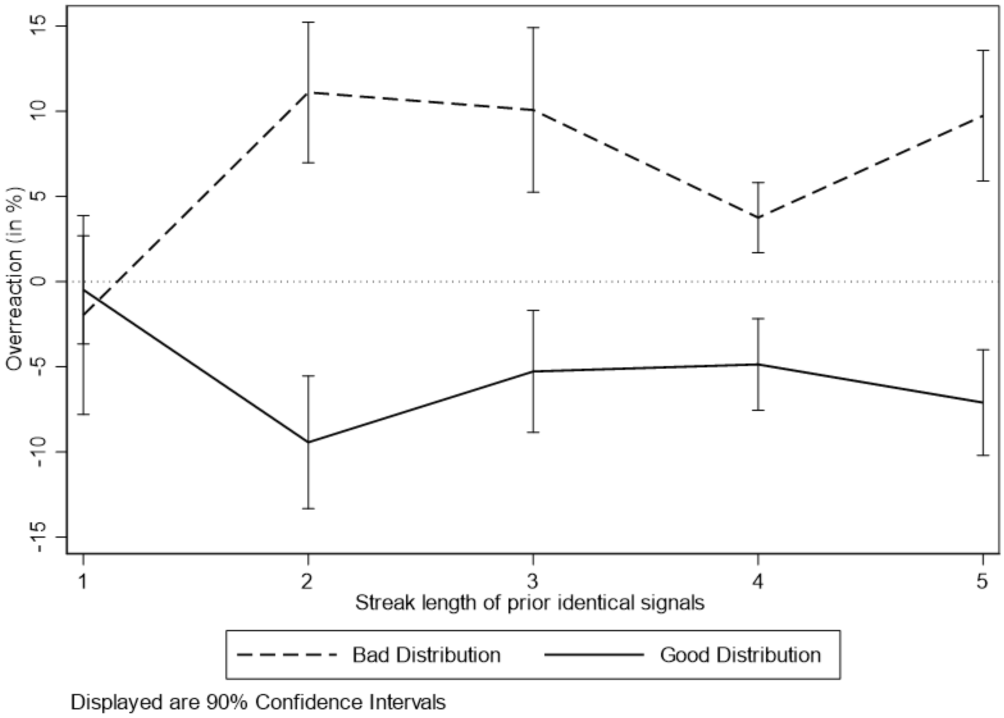

Figure 5: Subjects' Response Time for Same- and Opposite-Directed Signals

This figure displays participants' response time (in seconds) to provide their probability estimate for each round. The solid line represents average response times for disconfirming signals whereas the dashed line represents average response times for confirming signals. Displayed are 90% confidence intervals as well as a trend line for both confirming and disconfirming signals (in thin dotted lines).

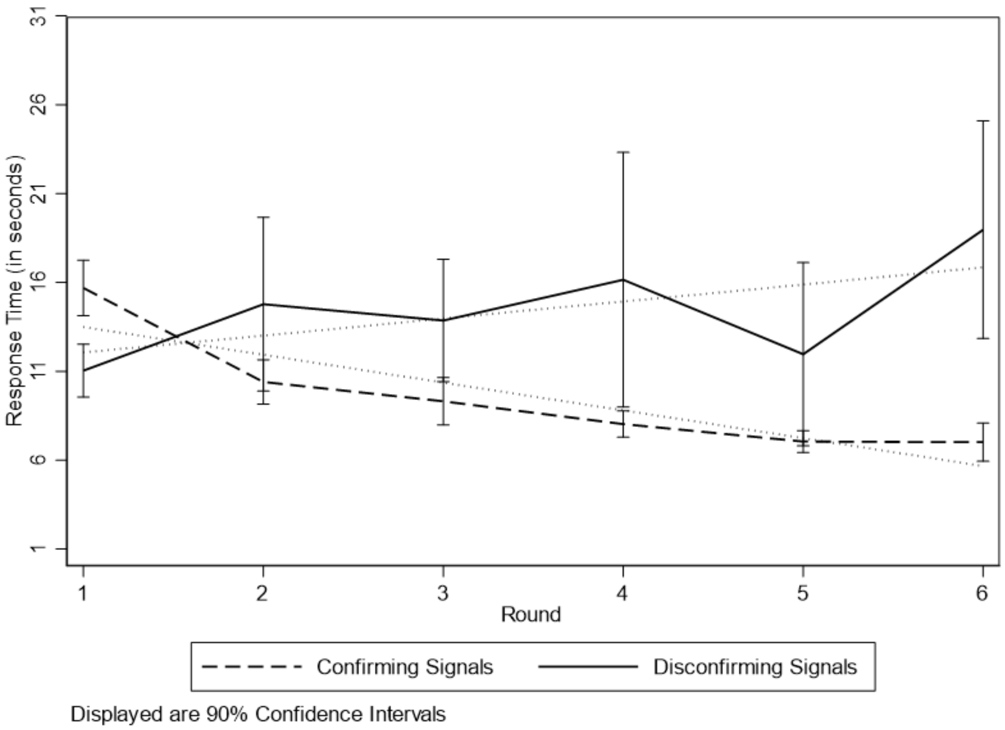

Figure 6: The Effect of Confirming and Disconfirming Earnings Surprises on Stock Returns

Using quarterly earnings announcements from June 2009 to May 2020, we calculate the average cumulative abnormal returns (CAR[0,t]) for extreme earnings surprise quintiles by whether the announcement was following a same-directed announcement surprise (confirming) or an opposite-directed announcement surprise (disconfirming). The solid lines present the development of cumulative abnormal returns for same-directed earnings surprises (blue color for the most positive surprises, red color for the most negative surprises) and the dashed lines present the development of cumulative abnormal returns for opposite-directed earnings surprises.

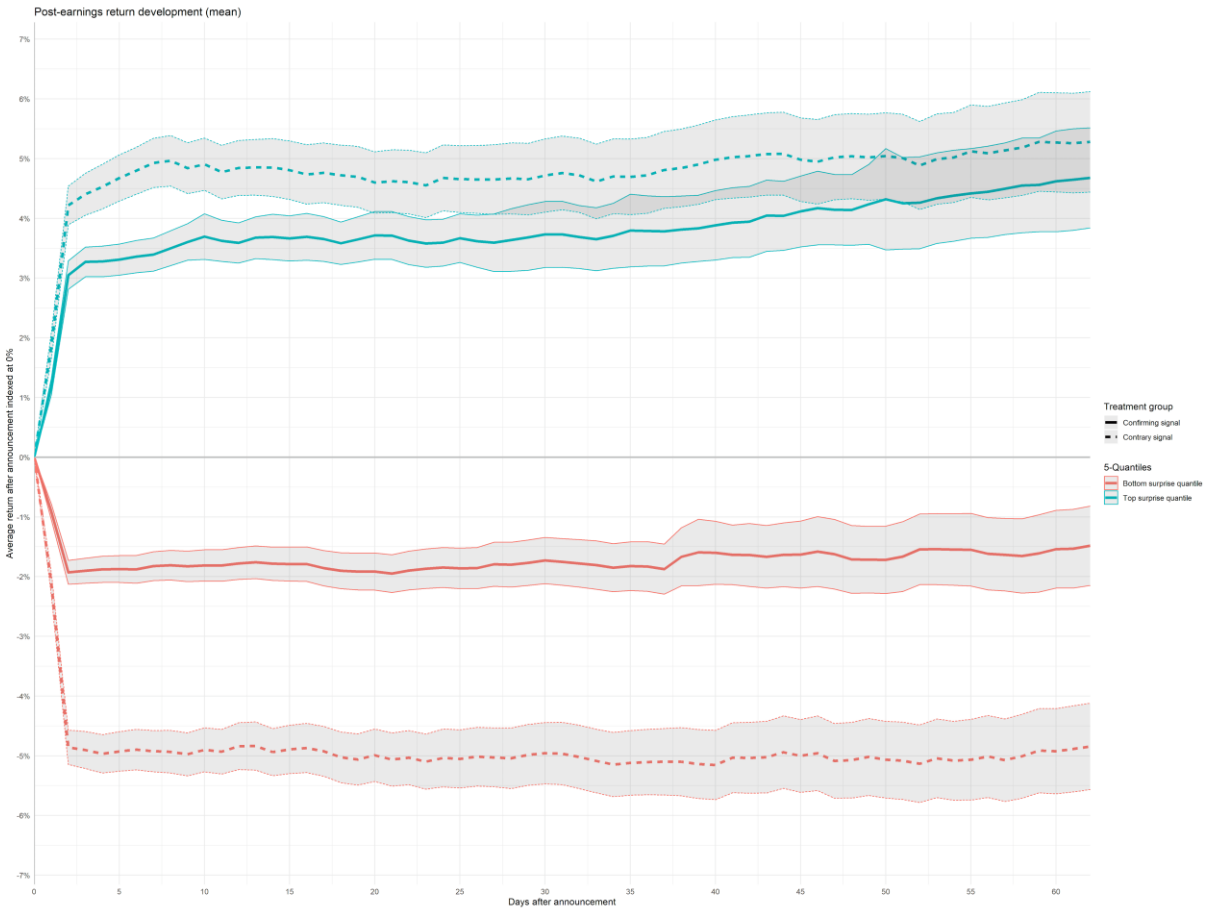

Table 1: Overview of Price Paths

This table provides an overview of the price paths in our experiments. Overall, there are twelve possibilities, six in the good distribution and six in the bad distribution, defined by the period in which the disconfirming signal occurs. The "–" sign represents a negative (bad) signal and the "+" sign a positive (good) signal.

			Good Distribution			
Path	1	2	3	4	5	6
G-1	–	+	+	+	+	+
G-2	+	–	+	+	+	+
G-3	+	+	–	+	+	+
G-4	+	+	+	–	+	+
G-5	+	+	+	+	–	+
G-6	+	+	+	+	+	–
			Bad Distribution			
Path	1	2	3	4	5	6
B-1	+	–	–	–	–	–
B-2	–	+	–	–	–	–
B-3	–	–	+	–	–	–
B-4	–	–	–	+	–	–
B-5	–	–	–	–	+	–
B-6	–	–	–	–	–	+

Table 2: Summary Statistics

This table shows summary statistics for our experimental data. Reported are the mean and the standard deviation (in parentheses) for each experiment individually. *Female* is an indicator variable that equals 1 if a participant is female. *Statistic skills* denotes participants' self-assed statistical skills on a 7-point Likert scale. *Risk preferences* are elicited by asking subjects to split an endowment between a risky and a risk-free asset (reported is the fraction invested risky). *Financial literacy* was assessed by asking subjects to identify the correct formula for calculating the expected value of the portfolio they selected. Through multiple choice answers, participants could make three basic errors (reported is the number of basic errors).

	Experiment 1	**Experiment 2**	**Experiment 3**	**Experiment 4**
	Baseline	Reduced Diagnosticity	Reduced Uncertainty	Cognitive Resources
Variable	(N = 408)	(N = 475)	(N = 417)	(N = 182)
Age	34.09	34.52	35.70	38.71
	(9.98)	(9.21)	(10.37)	(10.52)
Female	0.34	0.41	0.33	0.38
	(0.48)	(0.49)	(0.47)	(0.49)
Statistic Skills (1-7)	4.37	4.29	4.24	4.09
	(1.58)	(1.57)	(1.64)	(1.51)
Risk Preferences	45.14%	45.14%	49.17%	49.47%
	(29.19)	(28.58)	(29.33%)	(29.67)
Financial Literacy	1.31	1.32	1.36	1.47
	(0.94)	(0.92)	(0.97)	(1.04)

Table 3: Updating Behavior for Disconfirming Signals and Correction – Experiment 1

This table reports the results of six OLS regressions on how subjects update their posterior beliefs after a disconfirming signal and a correction in the baseline experiment. We report the results of OLS regressions for each distribution individually (good and bad distribution). The dependent variable in the regression model, *Change in Posterior Probability Estimate*, is the change in subjective posterior beliefs that the asset is paying from the good distribution between period *t* and period *t-1*. Independent variables include the *Disconfirm* dummy, an indicator variable that equals 1 if participants observe a disconfirming signal and zero otherwise, the *Correction* dummy, an indicator variable that equals 1 if a disconfirming signal is subsequently reverted, as well as *Change in Bayes*, which is the change in the correct Bayesian probability that the asset is good between period *t* and period *t-1*. Reported are coefficients and t-statistics (in parentheses) using robust standard errors clustered at the individual level. *, **, *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Dep. Variable	*Change in Posterior Probability Estimate*					
	Good Distribution			**Bad Distribution**		
Change in Bayes	0.851***	0.754***	0.577***	0.884***	0.790***	0.677***
	(14.48)	(14.02)	(12.89)	(14.00)	(13.21)	(11.47)
Disconfirm		-7.920***	-9.279***		8.154***	9.032***
		(-6.44)	(-6.98)		(6.13)	(6.46)
Correction			7.667***			-5.762***
			(5.65)			(-4.19)
Observations	964	964	964	668	668	668
R^2	0.309	0.347	0.378	0.352	0.393	0.411

Table 4: Reduced Diagnosticity and Overreaction – Experiment 2

This table reports the strength of the overreaction after a disconfirming signal in the reduced diagnosticity and the baseline experiment. The overreaction is defined as in Section 2 as the difference of the changes in beliefs after and prior to the disconfirming signal. Results are displayed individually for the good and the bad distribution. Reported are average absolute changes in probability estimates and standard errors (in parentheses). *, **, *** indicate significance at the 10%, 5%, and 1% level, respectively, using non-parametric Wilcoxon signed-rank tests with robust standard errors (in parentheses).

Good Distribution	Baseline	Lower Diagnosticity	Difference
Change in Beliefs (ΔP_t)	13.71	13.75	0.04
	(1.38)	(1.18)	
Prior Change in Beliefs (ΔP_{t-1})	3.00	4.43	1.43
	(0.71)	(0.79)	
Difference	10.72***	9.32***	1.39

Bad Distribution	Baseline	Lower Diagnosticity	Difference
Change in Beliefs (ΔP_t)	14.28	12.88	1.40
	(1.41)	(1.19)	
Prior Change in Beliefs (ΔP_{t-1})	5.98	5.78	0.20
	(0.90)	(1.23)	
Difference	8.30***	7.10***	1.20

Table 5: Reduced Uncertainty and Overreaction – Experiment 3

This table reports the strength of the overreaction after a disconfirming signal in the full information (reduced uncertainty) and the baseline experiment. The overreaction is defined as in Section 2 as the difference of the changes in beliefs after and prior to the disconfirming signal. Results are displayed individually for the good and the bad distribution. Reported are average absolute changes in probability estimates and standard errors (in parentheses). *, **, *** indicate significance at the 10%, 5%, and 1% level, respectively, using non-parametric Wilcoxon signed-rank tests with robust standard errors (in parentheses).

Good Distribution	Baseline	Full Information	Difference
Change in Beliefs (ΔP_t)	13.71	9.43	4.28**
	(1.38)	(1.43)	
Prior Change in Beliefs (ΔP_{t-1})	3.00	2.67	0.32
	(0.71)	(0.97)	
Difference	10.72***	6.77***	3.95**

Bad Distribution	Baseline	Full Information	Difference
Change in Beliefs (ΔP_t)	14.28	12.7	1.57
	(1.41)	(1.52)	
Prior Change in Beliefs (ΔP_{t-1})	5.98	3.96	2.02
	(0.90)	(1.05)	
Difference	8.30***	8.75***	-0.45

Table 6: Subjects' Response Time and Overreaction – Experiment 4

This table displays the absolute overreaction for the change detection experiment split by participants' response time after the disconfirming signal. The response time was classified into terciles (low, medium, and high) and the overreaction was computed as the average overreaction for each response time tercile. *, **, *** indicate significance at the 10%, 5%, and 1% level, respectively, using non-parametric Wilcoxon signed-rank tests with robust standard errors (in parentheses).

Absolute Overreaction	Low Response Time	Medium Response Time	High Response Time	Difference (High – Low)
Good Distribution	-2.77 % (3.30)	4.31 % (3.15)	6.90 %** (2.79)	9.67 %**
Bad Distribution	1.25 % (1.59)	6.10 % (3.84)	8.46 %** (3.31)	7.21 %*

Table 7: Descriptive Statistics

Using quarterly earnings announcements and analysts' earnings forecasts for the period from June 2009 to May 2020, we calculate the quarterly earnings surprise for each firm in our sample. In each calendar quarter, we sort quarterly earnings surprises during that quarter into either contrary or confirming in the direction of surprise relative to the previous-quarter earnings surprise. The table shows the average Size, B/M ratio, Earnings Surprise, Earnings Persistence, Earnings Volatility, Shares Outstanding, Number of Analysts, and Share Turnover as well as the number of observations.

Firm Characteristics	Full Sample	Confirming Surprises	Disconfirming Surprises	Difference (p-value)
Size (Mil $)	13,794	15,443	10,422	-5021 (<0.01)
B/M	0.54	0.51	0.61	0.1 (<0.01)
Earnings Surprise	0	0	0	0 (0.08)
Earnings Persistence	0.25	0.32	0.12	-0.2 (<0.01)
Earnings Volatility	0.02	0.02	0.02	0 (0.81)
Shares Outstanding	243,736	252,023	226,793	-25,230 (<0.01)
Number of Analysts	10.23	10.57	9.53	-1.04 (<0.01)
Share Turnover	21.29	21.36	21.14	-0.22 (0.56)
Common Equity (Mil $)	5029.26	5193.81	4723.26	-470.54 (0.02)
N	33,315	22,373	10,942	

Table 8: Market Reactions to Earnings News: Univariate Analysis

Using quarterly earnings announcements fr om June 2009 to May 2020, we calculate the average 2-day announcement cumulative abnormal returns (CAR[0,1]) and 60-day post-announcement cumulative abnormal return (CAR[2,61]) for extreme earnings surprise quintiles by whether the announcement was following a sequence of at least two same-directed announcement surprises (confirming) or opposite-directed announcement surprises (disconfirming). Earnings surprise quintiles are formed based on independent sorts of quarterly earnings announcements by the corresponding forecast error.

	Average CAR[0,1] for Earnings Surprise Quintile 1 and 5			**Average CAR[2,61] for Earnings Surprise Quintile 1 and 5**		
	Quintile 1	Quintile 5	Difference	Quintile 1	Quintile 5	Difference
Confirming	-1.93%	3.05%	4.98%***	0.49%	1.76%	1.27%**
Disconfirming	-4.86%	4.22%	9.08%***	0.01%	0.90%	0.89%
Difference	2.93%***	-1.17%***	-4.10%***	0.48%	0.86%	0.38%

Table 9: Market Reactions to Earnings News: Multivariate Analysis

This table reports the multivariate tests of the effects of the sign of the announcement relative to the prior announcement on the relation between announcement or post-announcement returns and earnings surprises. The dependent variable is indicated under each column heading. FE is the earnings surprise quintile (FE=1: lowest, 5: highest) and Disconfirming is a dummy that equals 1 if the sign of the current earning surprise is inconsistent with the sign of the prior earning surprise, and 0 otherwise. FE5 is an indicator variable for the top earnings deciles (FE=5). Control variables include size and book-to-market deciles, Earnings Volatility, Earnings Persistence, Share Turnover, and indicator variables for year, month, day of week, and Fama-French 10 industry classification. Standard errors adjusted for heteroskedasticity and clustering by the day of announcement are in parentheses. *,**,***indicate significance at the 10%, 5%, and 1% level, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dependent Variable	CAR[0,1]	CAR[0,1]	CAR[0,1]	CAR[2,61]	CAR[2,61]	CAR[2,61]	CAR[0,1]	CAR[2,61]
FE	0.01367***	0.01376***	0.01414***	0.00122	0.00109	0.00220		
	(0.00033)	(0.00033)	(0.00060)	(0.00098)	(0.00098)	(0.00171)		
Contrary	-0.03323***	-0.03280***	-0.03197***	0.00873*	0.000743	0.000742		
	(0.00173)	(0.00172)	(0.00171)	(0.00492)	(0.00492)	(0.00493)		
FE x Contrary	0.00865***	0.00847***	0.00836***	-0.00046	-0.00053	-0.00058		
	(0.00054)	(0.00054)	(0.00053)	(0.00151)	(0.00151)	(0.00151)		
FE5							0.05799***	0.00526
							(0.00275)	(0.00852)
Contrary							-0.01651***	0.00615
							(0.00171)	(0.00502)
FE5 x Contrary							0.02829***	-0.00143
							(0.00248)	(0.00703)
Controls		X	X		X	X	X	X
Controls interacted with FE			X			X	X	X
Constant	0.96420***	0.96776***	0.96623***	0.98816***	0.98562***	0.98205***	0.97737***	0.99391***
	(0.00104)	(0.00122)	(0.00191)	(0.00320)	(0.00399)	(0.00624)	(0.00179)	(0.00632)
Observations	51536	51536	51536	51536	51536	51536	20613	20613
R^2	10.10%	10.45%	11.98%	0.03%	0.44%	0.46%	16.80%	0.30%

Appendix

A. Experimental Instructions and Screenshots

Instructions Bayesian Updating (Exemplary for Experiment 1)

In this part we would like to test your forecasting abilities.

You will make forecasting decisions in one block consisting of 6 rounds.

Suppose you find yourself in an environment, in which a risky asset with an initial value of 50 can either increase by 5 or decrease by 5. The probability of either outcome (5 or -5) depends on the state in which the asset is (**good** state or **bad** state). If the risky asset is in the **good** state, then the probability that the risky asset increases in value by 5 is 70% and the probability that it decreases in value by 5 is 30%. If the risky asset is in the **bad** state, then the probability that the risky asset increases in value by 5 is 30% and the probability that it decreases in value by 5 is 70%.

The computer determines the state at the beginning of the block (consisting of 6 rounds). Within a block, the state does not change and remains fixed.

At the beginning of the block, you do not know which state the risky asset is in. The risky asset may be in the good state or in the bad state with equal probability.

At the beginning of each round, you will observe the payoff of the risky asset (5 or -5). After that, we will ask you to provide a probability estimate that the risky asset is in the good state and ask you how sure you are about your probability estimate. While answering these questions, you can observe the price development in a chart next to the question.

There is always an objective correct probability that the risky asset is in the good state. This probability depends on the history of payoffs of the risky asset already. As you observe the payoffs of the risky asset, you will update your beliefs whether or not the risky asset is in the good state.

Objective Bayesian Posterior Probabilities

This table provides all possible values for the objectively correct probability that the asset is in the good state for every possible combination of trials and outcomes. The initial prior for good and bad distribution is set to 50%. The objective Bayesian posterior probability that the asset is in the good state, after observing t high outcomes in n trials so far is given by: $\frac{1}{1+\frac{1-p}{p}\cdot(\frac{q}{1-q})^{n-2t}}$, where *p* is the initial prior before any outcome is observed that the stock is in the good state (50% here), and *q* is the probability that the value increase of the asset is the higher one (70% in Experiment 1 & 3, and 60% in Experiment 2).

		Experiment 1 & 3 (q = 70%)	**Experiment 2 (q = 60%)**
n (number of trials so far)	**t (number of high outcomes so far)**	**Probability [stock is good t high outcomes in n trials]**	**Probability [stock is good t high outcomes in n trials]**
0	0	50.00%	50.00%
1	0	30.00%	40.00%
1	1	70.00%	60.00%
2	0	15.52%	30.77%
2	1	50.00%	50.00%
2	2	84.48%	69.23%
3	0	7.30%	22.86%
3	1	30.00%	40.00%
3	2	70.00%	60.00%
3	3	92.70%	77.14%
4	0	3.26%	16.49%
4	1	15.52%	30.77%
4	2	50.00%	50.00%
4	3	84.48%	69.23%
4	4	96.74%	83.51%
5	0	1.43%	11.64%
5	1	7.30%	22.86%
5	2	30.00%	40.00%
5	3	70.00%	60.00%
5	4	92.70%	77.14%
5	5	98.57%	88.36%
6	0	0.62%	8.7%
6	1	3.26%	16.49%
6	2	15.52%	30.77%
6	3	50.00%	50.00%
6	4	84.48%	69.23%
6	5	96.74%	83.51%
6	6	99.38%	91.93%

Screenshots of Experiment 1

Figures B1 to B3 present the screens of the forecasting task as seen by subjects in the experiment (example round 4). One round consists of three sequential screens. First, subjects saw the payoff of the risky asset in the respective round. Second, the cumulated payoffs of the risky asset are shown in a price-line-chart and subjects are asked to provide a probability estimate that the risky asset pays from the good distribution. Finally, subjects are asked on a 9-point Likert scale how confident they are in their probability estimate.

Figure A1: Payoff screen

Figure A2: Probability estimate screen

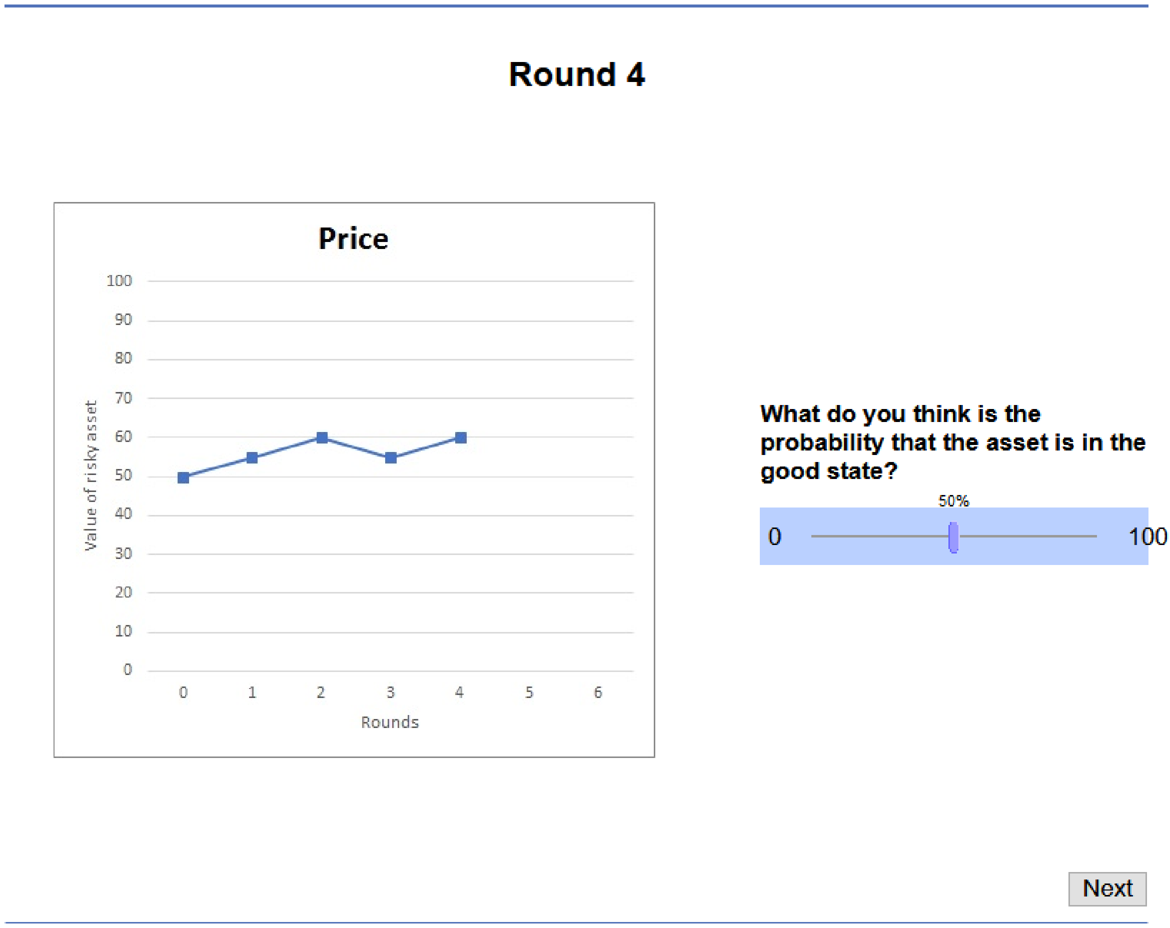

Figure A3: Confidence level screen

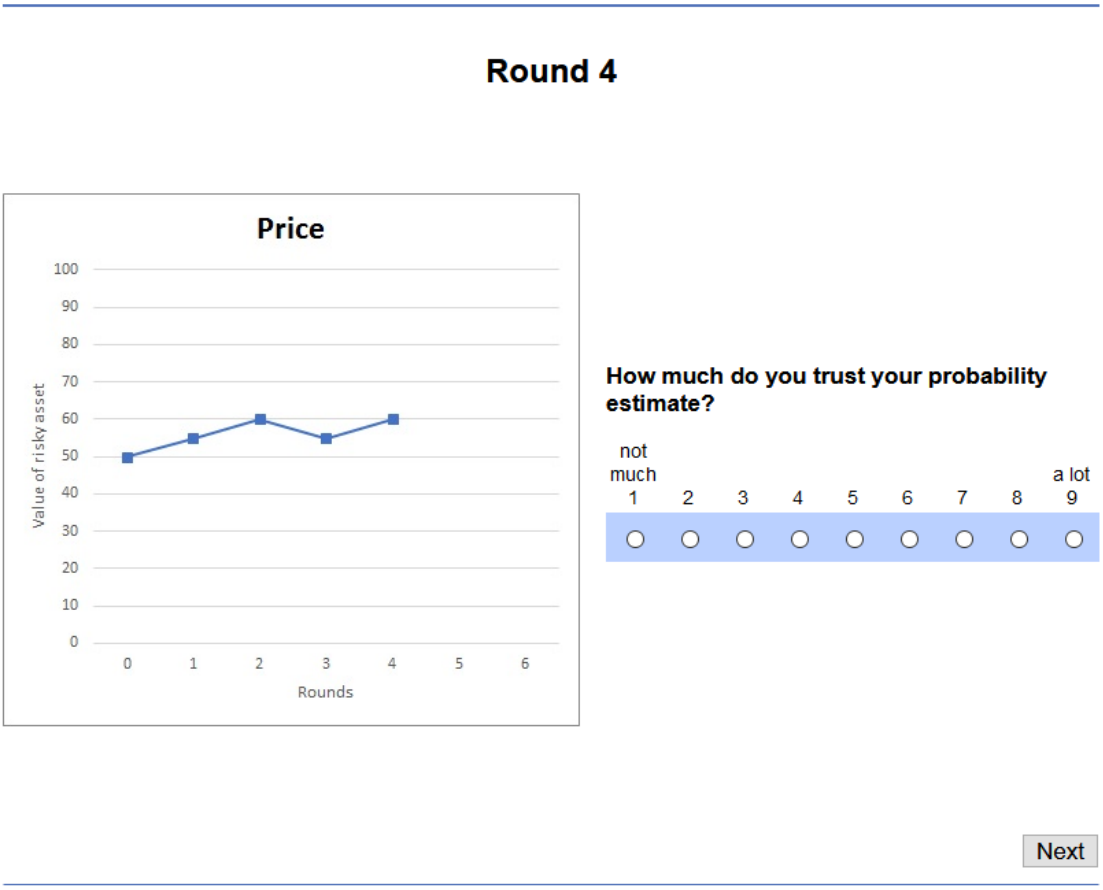

Comprehension Question for Bayesian Updating Task

Below we report the comprehension questions that participants had to answer correctly after reading the instructions to proceed to the Bayesian Updating task. Correct responses are displayed in *italic*.

1. If you see a series of +5, what is more likely?
   a. The risky asset is in the good state.
   b. The risky asset is in the bad state.

2. You observe a -5, how do you have to update your probability estimate that the asset draws from the good distribution?
   a. I reduce the probability estimate that the asset is in the good distribution.
   b. I increase the probability estimate that the asset is in the good distribution.

3. The correct probability estimate is let's say 0.70. Which probability estimate(s) would be in the range such that you earn 25 cents? [Note: You can check multiple boxes.]
   a. 0.55
   b. 0.67
   c. 0.75
   d. 0.85
   e. 0.87

4. At the beginning of the first period, the probability that the risky asset is in the good state is 50%.
   a. True
   b. False

B. Additional Results

Figure B.1: Subjects' Average Updating Behavior – Experiment 1

Panel A displays subjects' average probability estimates over six consecutive periods in the good distribution for price paths G-1 to G-6 individually. Panel B displays subjects' average probability estimates over six consecutive periods in the bad distribution for price paths B-1 to B-6 individually. The dashed line shows the objective Bayesian posterior probabilities and the solid line shows subjects' average probability estimates. Displayed are 95% confidence intervals.

Panel A: Good Distribution

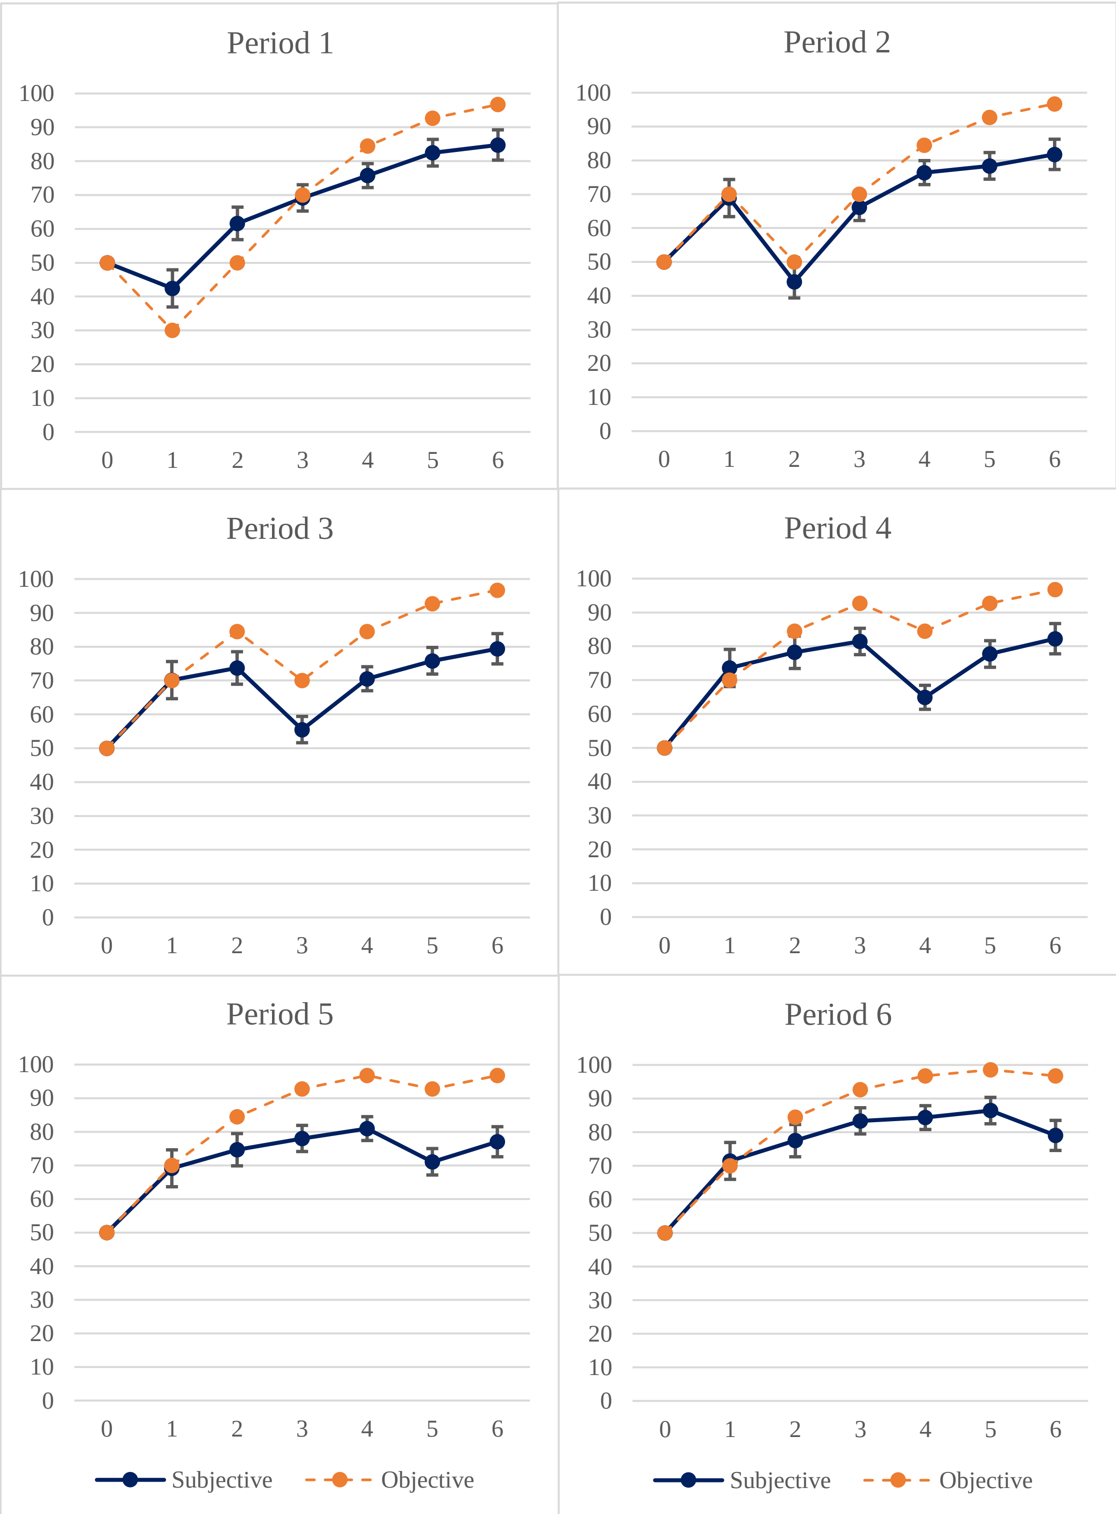

Panel B: Bad Distribution

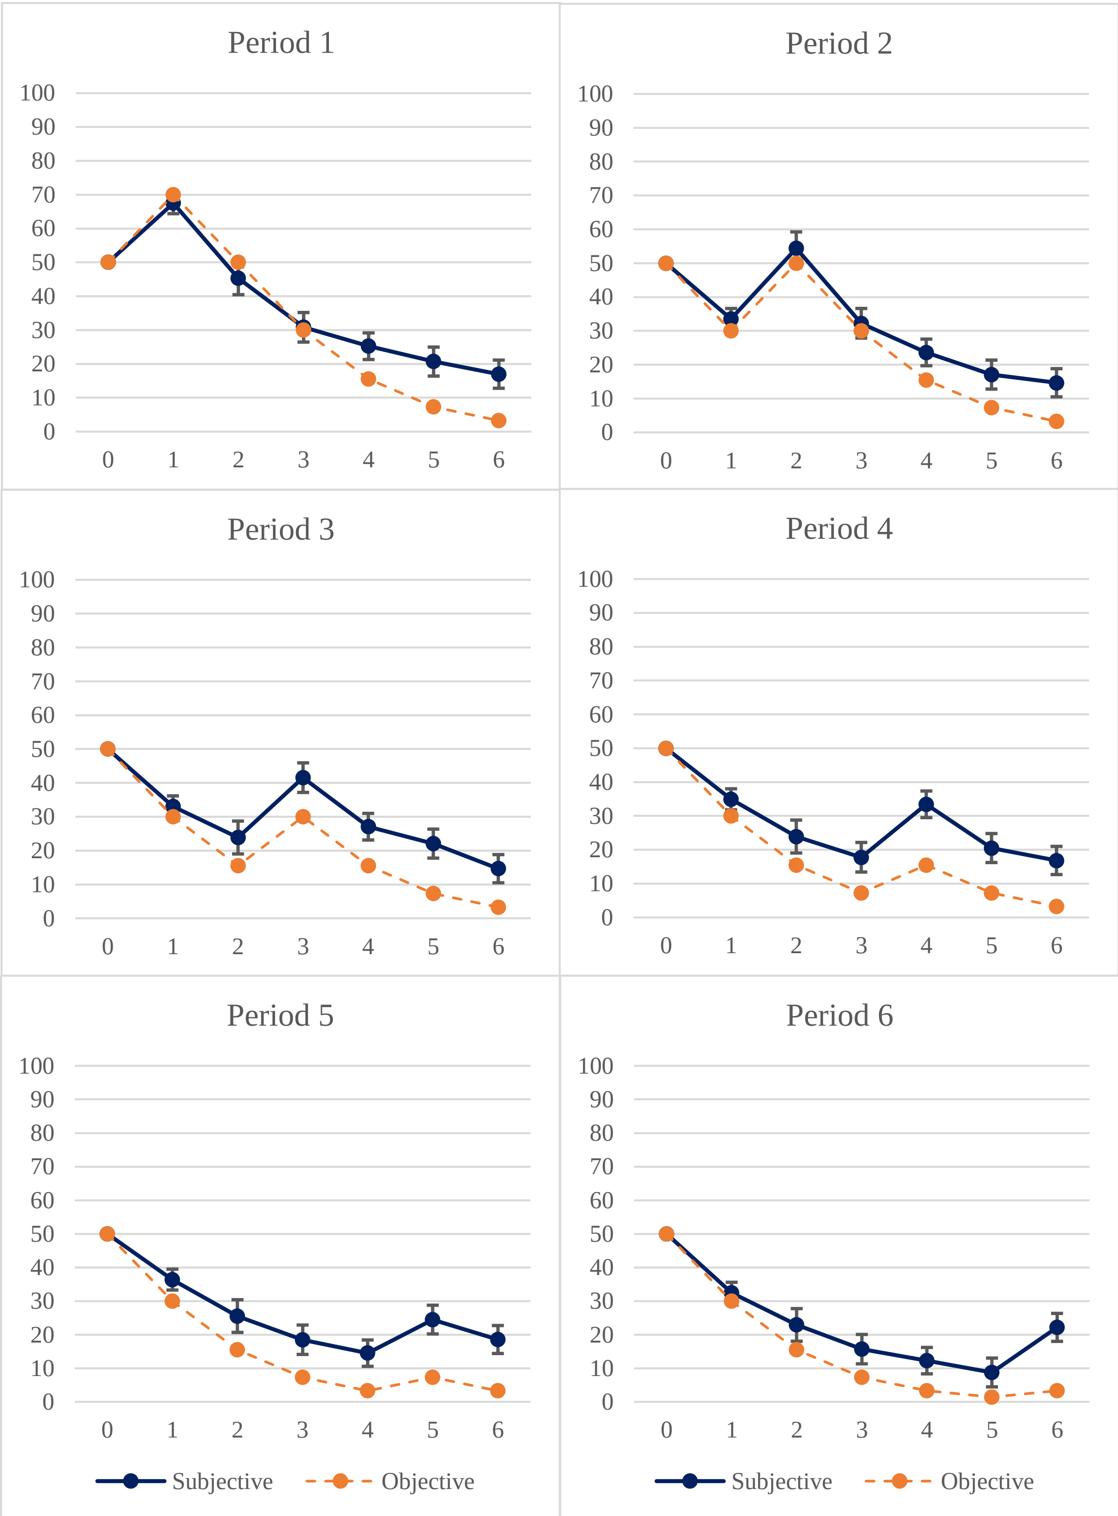

Figure B.2: The Effect of Confirming and Disconfirming Earnings Surprises on Stock Returns – Single Prior Signal

Using quarterly earnings announcements from June 2009 to May 2020, we calculate the average cumulative abnormal returns (CAR[0,t]) for extreme earnings surprise quintiles by whether the announcement was following a same-directed announcement surprise (confirming) or an opposite-directed announcement surprise (disconfirming). The solid lines present the development of cumulative abnormal returns for same-directed earnings surprises (blue color for the most positive surprises, red color for the most negative surprises) and the dashed lines present the development of cumulative abnormal returns for opposite-directed earnings surprises.

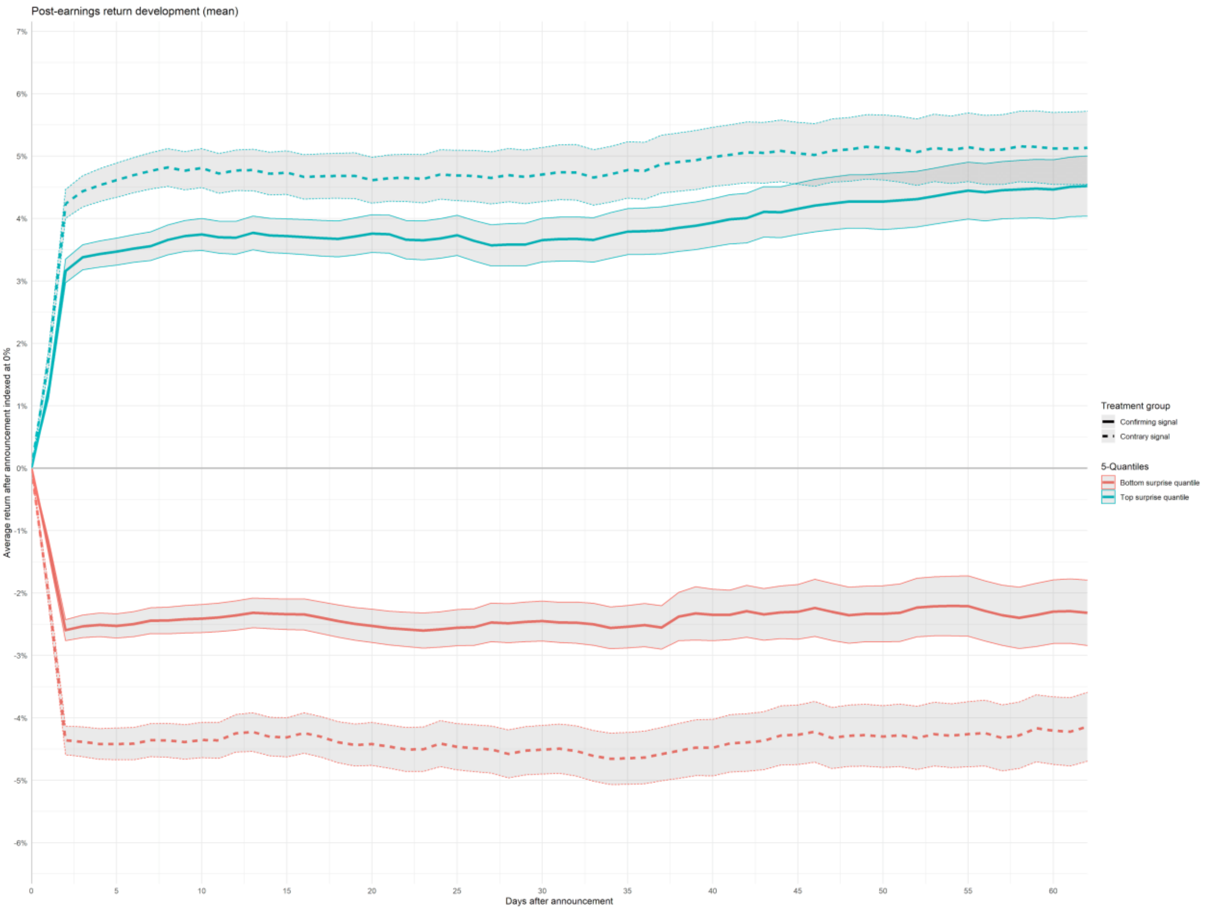

Figure B.3: The Effect of Confirming and Disconfirming Earnings Surprises on Stock Returns – Decile Sorts

Using quarterly earnings announcements from June 2009 to May 2020, we calculate the average cumulative abnormal returns (CAR[0,t]) for extreme earnings surprise deciles by whether the announcement was following a same-directed announcement surprise (confirming) or an opposite-directed announcement surprise (contrary). The solid lines present the development of cumulative abnormal returns for same-directed earnings surprises (blue color for the most positive surprises, red color for the most negative surprises) and the dashed lines present the development of cumulative abnormal returns for opposite-directed earnings surprises.

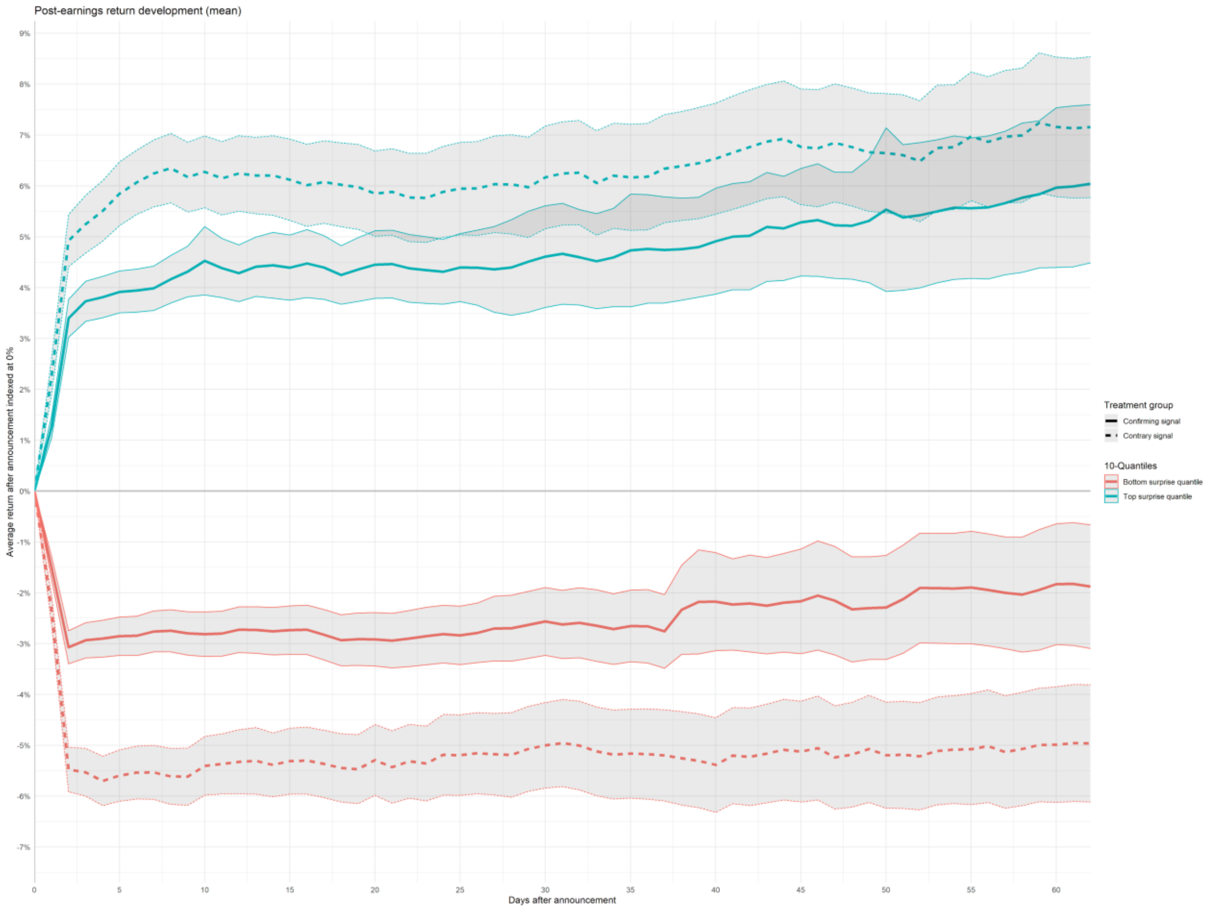

Table B.1: Signal Ordering and Updating Behavior – Experiment 1

This table reports the results of OLS regressions on how subjects updating behavior depends on the ordering of signals. We report the results for each posterior probability (from round 3 to round 6) and distribution (good and bad distribution) individually. The dependent variable in the regression model, *Probability Estimate in Period t*, is the subjective posterior belief that the asset is paying from the good distribution in period t. The independent variable includes *Round_Disconfirm* which is a variable that states the round in which the disconfirming signal is observed. Reported are coefficients and t-statistics (in parentheses) using robust standard errors. *, **, *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Dependent Variable	*Probability Estimate in Round t*							
	Good Distribution				**Bad Distribution**			
	Prob. Round 3	Prob. Round 4	Prob. Round 5	Prob. Round 6	Prob. Round 3	Prob. Round 4	Prob. Round 5	Prob. Round 6
Round_Disconfirm	-6.926***	-4.173***	-2.296***	-1.121**	5.248***	2.503**	0.973	0.902
	(-4.08)	(-4.35)	(-3.17)	(-2.19)	(2.81)	(2.20)	(1.16)	(1.45)
Constant	78.53***	83.22***	84.97***	85.60***	25.67***	22.22***	19.31***	15.14***
	(21.00)	(31.75)	(36.46)	(44.45)	(6.55)	(7.46)	(7.15)	(6.67)
Observations	131	171	205	241	94	119	148	167
R^2	0.115	0.101	0.047	0.020	0.079	0.040	0.009	0.013

Table B.2: Market Reactions to Earnings News: Univariate Analysis – Single Prior Signal

Using quarterly earnings announcements fr om June 2009 to May 2020, we calculate the average 2-day announcement cumulative abnormal returns (CAR[0,1]) and 60-day post-announcement cumulative abnormal return (CAR[2,61]) for extreme earnings surprise quintiles by whether the announcement was following a sequence of at least two same-directed announcement surprises (confirming) or opposite-directed announcement surprises (disconfirming). Earnings surprise quintiles are formed based on independent sorts of quarterly earnings announcements by the corresponding forecast error.

	Average CAR[0,1] for Earnings Surprise Quintile 1 and 5			**Average CAR[2,61] for Earnings Surprise Quintile 1 and 5**		
	Quintile 1	Quintile 5	Difference	Quintile 1	Quintile 5	Difference
Confirming	-2.59%	3.16%	5.76%***	0.31%	1.27%	0.96%**
Disconfirming	-4.36%	4.24%	8.60%***	0.27%	0.88%	0.61%
Difference	1.77%***	-1.07%***	-2.84%***	0.03%	0.39%	0.36%