

Enke, Benjamin; Schwerter, Frederik; Zimmermann, Florian

Working Paper

Associative Memory and Belief Formation

CESifo Working Paper, No. 7916

Provided in Cooperation with:

Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

Suggested Citation: Enke, Benjamin; Schwerter, Frederik; Zimmermann, Florian (2019) : Associative Memory and Belief Formation, CESifo Working Paper, No. 7916, Center for Economic Studies and ifo Institute (CESifo), Munich

This Version is available at:

<https://hdl.handle.net/10419/207307>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Associative Memory and Belief Formation

Benjamin Enke, Frederik Schwerter, Florian Zimmermann

Impressum:

CESifo Working Papers

ISSN 2364-1428 (electronic version)

Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH

The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute

Poschingerstr. 5, 81679 Munich, Germany

Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de

Editor: Clemens Fuest

www.cesifo-group.org/wp

An electronic version of the paper may be downloaded

- from the SSRN website: www.SSRN.com
- from the RePEc website: www.RePEc.org
- from the CESifo website: www.CESifo-group.org/wp

Associative Memory and Belief Formation

Abstract

This paper experimentally studies the role of associative memory for belief formation. Real-world information signals are often embedded in memorable contexts. Thus, today's news, and the contexts they are embedded in, may cue the selective retrieval of similar past news and hence contribute to the widely documented pattern of expectation overreaction. Based on a stylized version of models of associative memory in the literature, we develop a simple and tightly controlled experimental setup in which participants observe sequences of news about the stock market value of hypothetical companies. Here, identical types of news are associated with identical stories and images. In this setup, participants' expectations strongly overreact to recent news. We successfully verify the model's predictions about how the magnitude of overreaction should depend on the history of news. For example, once today's news are associated with the stories and images of previous opposite news, expectations systematically underreact. By exogenously manipulating the scope for imperfect and associative recall in our setup, we further provide direct causal evidence for the role of memory in belief formation and overreaction. Finally, we use our experimental data to estimate the model parameters that govern the strength of imperfect and associative recall over different time horizons.

JEL-Codes: D010.

Keywords: beliefs, expectations, memory, bounded rationality.

Benjamin Enke
Harvard University
Cambridge / MA / USA
enke@fas.harvard.edu

Frederik Schwerter
University of Cologne / Germany
frederik.schwerter@uni-koeln.de

Florian Zimmermann
University of Bonn & briq / Germany
florian.zimmermann@briq-institute.org

October 9, 2019

The experiments in this paper were pre-registered in the AEA RCT Registry under AEARCTR-0004247. For helpful discussions and comments we thank Ned Augenblick, Pedro Bordalo, Nicola Gennaioli, Alex Imas, Peter Schwardmann, Josh Schwartzstein, Jesse Shapiro, Adam Szeidl, Yves le Yaouanq, and in particular Andrei Shleifer. Schwerter thanks the Sloan Foundation for post-doctoral funding.

1 Introduction

This paper experimentally studies the role of imperfect and associative memory for the formation of expectations or beliefs. In the textbook model of Bayesian updating, memory imperfections play no role: agents entertain a prior belief, update this belief upon receipt of information, and yesterday’s posterior equals today’s prior. Our paper starts from the premise that people do not constantly entertain beliefs about every potentially relevant state of the world. Rather, when people are prodded to act on or update their beliefs, they need to reconstruct their prior knowledge and beliefs from memory. This observation raises the empirical question how people retrieve prior information, and which features of news make it more or less likely for memory traces to get recollected.

The second observation that motivates our paper is that real-world information signals typically do not just consist of abstract information. Rather, information is often embedded in memorable contexts, including stories, images, emotions, or sounds. Moreover, *similar news are frequently embedded in similar contexts*. For example, when individuals receive negative feedback about their performance, these negative news are often associated with scolding, public shaming, and emotions of insufficiency. Similarly, when good news prevail in the stock market, people are disproportionately exposed to bulls, upward-sloping trend lines, and good-times stories.¹ To take yet another example, when immigration opponents relay negative information about the “typical” character traits of immigrants, then this often occurs through similar stories and images involving theft and other forms of violence.

These observations motivate the question about the role of associative recall in belief formation. The associative nature of memory has recently received increased attention in the theory literature ([Mullainathan, 2002](#); [Bordalo et al., 2017a](#)). A central prediction that emerges from this body of work is that asymmetric context-cued recall could lead to overreaction: after receipt of a piece of news, people reconstruct past knowledge from memory, yet predominantly remember those past news that appeared in similar contexts as today’s news. As a consequence, expectations might *look like* they overreact to recent news. However, even though modeling errors in belief formation is often argued to require knowledge about the underlying psychological micro-foundations ([Fudenberg, 2006](#)), direct empirical evidence on a potential link between overreaction and memory is scarce.

To make progress, we present laboratory experiments that are structured around the predictions of a simple formal framework that applies the idea of associative recall to belief formation, based on the formulations in [Bordalo et al. \(2017a\)](#) and [Mullainathan \(2002\)](#). In this model, decision-makers (i) have imperfect memory; (ii) are more likely

¹See [Akerlof and Shiller \(2009\)](#) and [Shiller \(2017\)](#).

to recollect a piece of news from the past if the context in which it is experienced is similar to today’s context; and (iii) are naïve about their memory imperfections. This stylized version of existing models predicts overreaction in expectations. Importantly, this overreaction does not occur because people incorporate the last signal in an irrational manner, but rather because they behave rationally *conditional on what they recall*. The model makes predictions about how such overreaction depends on the precise signal history, the correlation structure between signals and contexts, the imperfection of memory, or the strength of associative recall. Each of our nine experimental treatments with a total of 650 lab subjects is designed around these predictions.

We propose a new experimental paradigm to investigate the role of associative memory for belief formation in an economic decision context. This paradigm builds a bridge between (i) the types of tightly-controlled, quantitative, and financially incentivized designs that dominate modern experimental economics research on bounded rationality and (ii) psychological research on cued recall. In our experiment, participants predict the stock market value of multiple hypothetical companies. We adopt this particular framing just to make the task intuitive for subjects, rather than because we think of our work as a finance application. The experiment comprises two distinct periods that we think of as “past” and “present.” Across both periods, a subject sequentially observes multiple pieces of news about a company on their decision screen, where each piece of news takes on the value +10 or −10. The value of a company is deterministic and given by 100 plus the sum of all news that were shown up to a given point in time. As in the motivating examples, the news are not communicated to subjects as mere abstract numbers but are embedded in a context, which consists of a story and an image that relate to the piece of news. For example, for one company, a positive signal would be shown with a story about the company having launched a successful advertisement campaign with a celebrity, accompanied by a picture of that celebrity.

In the baseline version of the experiment, again like in the motivating examples, identical news are embedded in identical contexts: there is a one-to-one mapping between {Company $\times$ type of news} and context. That is, for each company, all positive news are communicated using the same context, and all negative news are communicated using the same context. However, the same context is never used for different types of news or for different companies. All of this is known to subjects.

In the first period of the experiment, a subject sequentially observes a weakly positive number of pieces of news for a company and then states a first belief about the value of a company. This process is repeated for all companies. Using the data on first-period beliefs, we successfully verify that – absent memory constraints – subjects understand our new paradigm and are well-capable of aggregating the signals into a rational guess.

After the first period of the experiment, we implement a time gap during which sub-

jects work on an unrelated real effort task. In the second period, subjects observe up to one additional piece of news for a company and immediately after state their second-period belief about the value of that company. In addition, subjects explicitly indicate how many positive or negative signals they recall having seen throughout the experiment. Again, this procedure is repeated for all companies. As before, the true value of a company is given by 100 plus the sum of all signals that have accumulated throughout the entire experiment, including in the first period. The basic intuition behind this experimental setup is that observing a particular piece of news and context today might make it more likely for subjects to (asymmetrically) remember similar first-period news.

In this setup, our interest lies in evaluating the extent to which second-period beliefs overreact with respect to the second-period signal. Because of the simple deterministic structure of the experiment, the prediction of a rational (Bayesian) model is that the OLS coefficient in a regression of second-period beliefs on second-period signals equals one. Likewise, a model with imperfect but no associative memory predicts a regression coefficient of one. In contrast, our formal framework predicts that (i) the OLS coefficient is larger than one, meaning that second-period beliefs overreact; (ii) overreaction increases in the number of first-period signals that take on the same realization as the second-period signal (because more first-period news can be cued); (iii) overreaction disappears if memory is exogenously manipulated to be perfect; (iv) overreaction disappears if associative recall is exogenously shut down; and (v) expectations under- rather than overreact if recent news appear in a context that was previously associated with the opposite type of signal. We provide causal tests of each of these predictions. All of our experiments were pre-registered, including a pre-analysis plan.

We test predictions (i) and (ii) using the baseline treatment variation *Main* discussed above. We find that, in the experimental data, second-period beliefs strongly overreact: the aggregate OLS regression coefficient of the second-period signal is 1.10, substantially larger than its rational or imperfect-but-no-associative-recall benchmark of one. Moreover, as predicted, across tasks the magnitude of overreaction is strongly increasing in the number of first-period signals that take on the same value as the second-period signal. For instance, when subjects do not observe any first-period signals that match the second-period signal, their expectations do not overreact at all. Yet overreaction monotonically increases in the number of first-period signals that match the second-period signal (prediction (ii)). We verify that subjects' direct recall data further support our findings on participant expectations. Thus, taken together, these patterns point to the importance of associative recall in our setup and cannot be explained by other accounts such as recency effects.

To provide causal evidence for the role of imperfect and associative memory, we turn to testing predictions (iii) and (iv). To this effect, we exogenously manipulate the

strength (or relevance) of both imperfect and associative memory. To show that imperfect memory is necessary in order for overreaction to arise in our setup, we introduce treatment *Reminder*. This treatment follows exactly the same structure as condition *Main*, except that before subjects observe the second-period signal for a given company, they are reminded of their own first-period belief. Viewed through the lens of our formal framework, this treatment manipulation eliminates the imperfection of memory, so that asymmetric recall and hence overreaction can no longer take place. We find that, in treatment *Reminder*, subjects' beliefs do indeed not overreact.

Having documented the role of imperfect memory for overreaction, we next study the role of associative recall. In our model, associativeness operates via the similarity of contexts. Thus, in order to show that it is indeed associative memory that generates overreaction in our experiments, treatment *NoCue* follows the same structure as condition *Main*, except that each piece of news is communicated with a different context. That is, subjects never observe the same story or image twice, even if they receive the same signal for a given company twice. Thus, viewed through the lens of our formal framework, this treatment manipulation eliminates (or at least substantially reduces) the extent to which context-based associativeness can affect recall. The results show that overreaction disappears entirely in *NoCue*, and the treatment difference in overreaction between *Main* and *NoCue* is quantitatively large and statistically significant. Again, all of these results hold not only when we consider participants' expectations but also when we directly look at their reported recall of first-period signals.

In all experiments reported above, types of news and contexts (stories and images) were connected through a one-to-one mapping. That is, all positive signals for a given company appeared with the same context, and all negative signals appeared with the same (yet different) context. In treatment *Underreaction*, we modify this correlation structure between signals and contexts to test prediction (v) above. Specifically, in the second period of the experiment, positive signals are communicated with the context that was associated with negative signals for that same company in the first period. Likewise, negative signals for a company appear in contexts that were previously associated with positive signals for that same company. In this treatment, our formal framework predicts that expectations should underreact to the last signal because it selectively cues the retrieval of signals that took on a different value than the current signal. To causally isolate the role of memory, we again implement a control condition *Underreaction reminder*, in which subjects were reminded of their first-period belief before receiving a second-period signal.

We find that expectations in *Underreaction* systematically underreact, also relative to treatment *Underreaction reminder*. Moreover, based on our simple model, we again derive and empirically verify predictions about how the magnitude of underreaction should

depend on the signal history. These results highlight that associative memory generates predictable patterns of over- and underreaction, purely depending on the precise ways in which contexts are linked to pieces of news.

We conclude our experimental analysis with two sets of treatments aimed at documenting robustness. First, while all experiments summarized up to this point implemented a time lag between the first and second period of 15 minutes, we verify that we find very similar patterns when we extend the time lag to three days. Second, in a separate set of treatments, we show that our experimental results on overreaction and the corresponding role of associative memory are not restricted to participants' beliefs. Rather, we find almost identical results when we instead elicit subjects' willingness-to-pay for the companies in our experiment. Thus, the specific ways in which associative memory distorts beliefs in our setup also translates into corresponding choices.

All of our main results are derived from theoretically-motivated and pre-registered, yet ultimately reduced-form regression specifications. In complementary analyses in the final part of the paper, we directly estimate our stylized model of memory and belief formation, in particular the parameters that govern the imperfection of memory and the strength of associative recall. The results of these estimations suggest that associative recall plays a quantitatively large role in generating observed beliefs. For example, our parameter estimates suggest that the probability of accurately recalling a piece of news is 50% (30 percentage points) higher if it got cued by the second-period signal. We further document that our simple model matches the observed data well: the model parameters that are estimated from participants' beliefs data accurately predict subjects' explicitly stated recall of news.

In summary, the central contribution of our paper is a theoretically-structured experimental analysis of the role of associative recall for belief formation. This paper hence fits into an emerging literature that has argued for the importance of associative memory for economics. [Mullainathan \(2002\)](#) and [Bordalo et al. \(2017a\)](#) present models of how cued recall shapes economic decision-making across a broad set of domains, including how it generates overreaction in belief formation. Related work has investigated the implications of memory in applied settings such as updating biases ([Gennaioli and Shleifer, 2010](#); [Wilson, 2014](#)), financial markets ([Bodoh-Creed, 2013](#); [Bordalo et al., 2017b, 2018](#); [Wachter and Kahana, 2019](#)), self-esteem ([Koszegi et al., 2019](#)), or personal experiences ([Herz and Taubinsky, 2017](#); [Malmendier and Nagel, 2015](#)).

As much of the simple formalism that structures our experiments directly draws from this literature, we view our experiments as providing some of the first direct evidence from tightly structured and quantitative economic decision making tasks in relation to this emerging body of theoretical work. Psychological experiments on associative recall exhibit a different structure than the experiments that are presented here (see [Kahana,](#)

2012, for an overview). These experiments typically consist of explicit cued recall problems (such as with words), rather than of the types of quantitative reasoning tasks that characterize modern experimental economics research. [Bordalo et al. \(2019\)](#) present an experiment on selective recall of abstract images that shows a link between associative memory and the representativeness heuristic.² Finally, because we experimentally identify one potential mechanism behind overreaction, our work also relates to a literature that documents overreaction or excess movement in survey expectations about macroeconomic variables or geopolitical events ([Augenblick and Rabin, 2018](#); [Augenblick and Lazarus, 2018](#); [Bordalo et al., 2018, 2017b](#)).

The remainder of the paper proceeds as follows. Section 2 offers a stylized formal framework that motivates the experimental design and structures the empirical analysis. Section 3 describes the experimental design, implementation, and pre-registration. Sections 4 and 5 present the evidence on overreaction and the roles of imperfect and associative memory therein. Section 6 considers the case of underreaction, while Section 7 presents treatments aimed at documenting robustness. Section 8 explicitly estimates the model and Section 9 concludes.

2 Theoretical Framework

2.1 Setup

This section presents a stylized model to guide the design of the experiments and to structure the empirical analysis. The mechanics of the model directly build on some of the formulations in [Mullainathan \(2002\)](#) and [Bordalo et al. \(2017a\)](#). The framework rests on three key assumptions: (i) people do not permanently entertain a prior belief but may instead forget it over time, so that once they are prompted to update their beliefs they first need to reconstruct past knowledge from memory; (ii) this recollection process is subject to associative recall, meaning that news are more likely to get remembered if they were observed in a context that is similar to the context in which today’s signal is observed; and (iii) people are (at least partially) naïve about their biased memory technology. For simplicity, we abstract away from additional behavioral assumptions that the literature on associative memory has incorporated, such as salience or rehearsal.

Consider a decision-maker (DM) who forms beliefs about the state of a time-varying

²More indirectly, our paper also relates to recent experimental work on the role of motivated (self-serving) memory, as in [Zimmermann \(forthcoming\)](#) and [Carlson et al. \(2018\)](#). More broadly, our paper also fits into the recent experimental literature on bounded rationality, in particular work that has focused on the micro-foundations behind behavioral anomalies ([Enke and Zimmermann, 2019](#); [Enke, 2017](#); [Enke and Graeber, 2019](#); [Esponda and Vespa, 2016](#); [Martínez-Marquina et al., 2017](#); [Dertwinkel-Kalt et al., 2017](#); [Frydman and Jin, 2018](#)).

stochastic variable θ_t with initial value v . We consider two periods that we will think of as “past” and “present.” In any given period, the value of θ is given by its initial value plus the sum of all news n_j that have accumulated up to this point, where $n_j \in \{-q, q\}$. News are equally likely and i.i.d. We will use the terms “news” and “signal” interchangeably.

A piece of news n_j is associated with a memorable context $c_j \in \{L, H\}$. In the “past”, k news arrive, so that $\theta_1 = v + \sum_{x=1}^k n_x$. In $t = 1$, there is a one-to-one mapping between type of news (positive or negative) and context (high or low): $n_x = n_y \Leftrightarrow c_x = c_y$.

In the “present” ($t = 2$), the DM observes one final piece of news n_{k+1} . Thus:

$$\theta_2 = v + \sum_{x=1}^k n_x + n_{k+1} \quad (1)$$

Just as in $t = 1$, the piece of news is associated with a context. We will consider two regimes. In the first, second-period news and contexts are associated in the same way as in the first period. In the second regime, the DM receives second-period news in a context opposite to what he was exposed to in the first period, meaning that he observes positive news in a “negative” context and vice versa. As a shorthand for this “correlation” between news and context, we define

$$\rho \equiv \begin{cases} 1 & \text{if } P(c_2 = H | n_{k+1} = q) = P(c_2 = L | n_{k+1} = -q) = 1 \\ -1 & \text{if } P(c_2 = L | n_{k+1} = q) = P(c_2 = H | n_{k+1} = -q) = 1 \end{cases}$$

2.2 Memory and Beliefs

In period 1, the DM observes a collection of news which deterministically pin down the true state. Thus, when prompted for their belief, the DM will state $b_1 = v + \sum_{x=1}^k n_x$.

Our object of interest is the extent to which the DM’s belief about θ responds to the latest piece of news n_{k+1} . A rational (or Bayesian, though there is no uncertainty here) DM would observe the collection of n_x and then correctly predict $\theta_2 = v + \sum_{x=1}^k n_x + n_{k+1}$.

Suppose instead that the DM potentially forgets some of the news between $t = 1$ and $t = 2$. That is, when our DM wakes up in $t = 2$, he has potentially forgotten some of the $n_1 \dots, n_k$. Thus, his belief (after observing n_{k+1}) is given by

$$b_1 = v + \sum_{x=1}^k m_x n_x + n_{k+1} \quad (2)$$

where $m_x \in \{0, 1\}$ denotes whether the DM remembers piece of news n_x .

Whether or not the DM remembers a piece of news is determined by both (i) imperfect and (ii) associative memory. First, by imperfect recall we mean that, irrespective of the piece of news, there is some probability that the DM will forget. The reduced-form

assumption of imperfect recall is a shorthand for different mechanisms that have been highlighted in the psychological literature, in particular that of interference (Kahana, 2012). By the logic of interference, it is harder for people to recall a specific item if they have been exposed to many similar items in the past. In our experiments, we generate interference (and hence imperfect memory) by implementing the same type of judgment task multiple times with different sets of signal realizations and contexts.

Second, by associative recall we mean that the probability of recalling a piece of news from the past is higher if it is cued by today's signal. That is, a signal is more likely to get remembered if it occurred with the same context as today's signal.

We assume that the DM forms beliefs exclusively from what he recalls and is not aware of his biased memory technology. This implies naïveté about memory imperfections as in Mullainathan (2002). In principle, naïveté could come in two facets: (i) the DM fails to realize that he sometimes forgets, i.e., that there are signals he does not recall; (ii) the DM realizes that he sometimes forgets, but he does not take into account that his recall is associative and hence asymmetric. In Appendix A.1, we formalize these types of naïveté and show that our predictions are robust to assuming partial naïveté.

We formalize the probability of remembering signal n_x as follows. First, the baseline probability of recall is $r < 1$. In addition, there is an increase in the probability of recalling $(1 - r)a$, $a < 1$, if the context c_{k+1} that is associated with n_{k+1} is the same as the context that is associated with news n_x . This formulation implies that associative memory matters more for DM with highly imperfect memory (low r). Formally:

$$m_x = \begin{cases} 1 & \text{with probability } r + (1 - r)a \mathbb{1}_{c_x = c_{k+1}} \\ 0 & \text{else} \end{cases} \quad (3)$$

Denote by $z \geq 0$ the number of news in $t = 1$ that were observed in the same context as n_{k+1} . The expected belief in period $t = 2$ is then given by:

$$\begin{aligned} E[b_2 | n_x, n_{k+1}] &= v + n_{k+1} + \sum_{x=1}^k m_x n_x \\ &= v + n_{k+1} + \sum_{x=1}^k r n_x + \sum_{x=1}^z (1 - r)a n_x \end{aligned} \quad (4)$$

$$= v + [1 + \rho z(1 - r)a] n_{k+1} + r \sum_{x=1}^k n_x \quad (5)$$

$$= v + n_{k+1} + [r + \rho(1 - r)a] \sum_{x=1}^z n_x + r \sum_{x=z+1}^k n_x \quad (6)$$

Equation (5) is the core expression that we subject to systematic experimental tests.

It implies the rational (Bayesian) prediction that second-period beliefs will respond with a coefficient of one to variation in the second-period signal. On the other hand, viewed through the lens of imperfect and associative memory, equation (5) suggests that expectations will overreact. This is the central prediction of our framework. At the same time, the equation also clarifies that overreaction does not occur because people incorporate the last signal in some suboptimal way, but only because people selectively retrieve first-period signals. Moreover, equation (6) clarifies that this overreaction can equivalently be understood as increased sensitivity of beliefs to past news that were communicated in the same context as today’s news (the third term), relative to news that were communicated in a different context (the fourth term in eq. (6)).

Equation (5) suggests the following abstract hypotheses, which we concretize for our experimental implementation in Section 3:

Hypotheses.

1. *If the correlation between news and context is positive ($\rho = 1$), expectations overreact to today’s news, on average. Put differently, expectations are more sensitive to past news that took on the same realization as today’s news.*
2. *Overreaction increases in the number of past news that were communicated in the same context as today’s news (z).*
3. *Overreaction increases in the imperfection of memory ($1 - r$).*
4. *Overreaction increases in the strength of associative recall (a).*
5. *If the correlation between news and context is negative ($\rho = -1$), expectations underreact to today’s news, on average.*
6. *This underreaction increases in the number of past news that there were communicated in the same context as today’s news.*

It is worth highlighting that these predictions rely on the presence of associative recall $a > 0$. Thus, models of recency bias (Fudenberg et al., 2014) or optimized responses to imperfect memory (Wilson, 2014) do not generate this joint set of predictions. For example, recency bias predicts overreaction, but not that overreaction depends on the history of news in nuanced ways, or that it disappears once memory imperfections are shut down.

3 Experimental Design

An environment in which the role of memory for belief formation can be studied requires (i) a dynamic setup in which subjects state beliefs twice with a delay period

in-between, such that previously formed posteriors do not mechanically translate into current priors; (ii) variation in signal histories that allow for nuanced predictions about when overreaction should be more or less pronounced; (iii) treatment variations that allow the exogenous manipulation of memory constraints, the role of associativeness, and the correlation structure between signals and contexts; and (iv) incentive-compatible belief elicitation. Our design was built to accommodate these features.

3.1 Experimental Setup

To isolate the role of memory, we implemented a simple deterministic decision environment in which, absent potential memory constraints, behaving rationally is trivial. Subjects were asked to guess the stock market value of twelve hypothetical companies at a given point in time. Continuing the notation from Section 2, the value of a company at time t , θ_t , was given by the baseline value, $v = 100$, plus the sum of all news about that company up to t :

$$\theta_t = 100 + \sum_{x=1}^k n_x. \quad (7)$$


where k is the number of signals that were shown up to t . News were equally likely to be positive, $n_x = 10$, or negative, $n_x = -10$, and were randomly and independently drawn by the computer. All of this was known to subjects.

News were not only communicated as abstract numbers, but were shown on subjects' computer screens with what we refer to as a context: an image and a story. The written instructions clarified that these images and stories were supplied to "explain" to subjects why a particular piece of news for a company was observed. For instance, all stories that accompanied positive news gave some rationale for why the value of the company had gone up, such as a successful marketing campaign or a recent technological innovation. The content of the story and the picture were tailored to match each other. The signal, picture and story were displayed on subjects' computer screens for 15 seconds. The time was calibrated such that subjects would have sufficient time to process the news, as well as to fully grasp the content of the picture and the story. Appendix E contains examples of these images and stories (see Figures 12 and 13).

The experiment consisted of two periods, as summarized in Figure 1. In both periods, participants estimate the stock market value of hypothetical companies. It was made salient to subjects that information from the first period of the experiment would also be relevant for their estimates in the second period.

In period 1, subjects sequentially observed news and corresponding contexts for a particular company on their computer screens. Then, they were asked to estimate the

Figure 1: Experimental Timeline



company’s current stock market value. This procedure was repeated for all twelve companies. Across companies, the number of signals varied between zero and three.³ Beliefs in period 1 allow us to verify whether subjects understand the basic information structure, had sufficient time to process the information, and are in principle able to form correct posteriors in our decision environment. As we will see below, first-period beliefs are indeed always very close to rational beliefs, which lends credence to our assumption that (absent memory constraints) subjects understand our design and are well-capable of behaving optimally.

After period 1, we implemented a time gap in which subjects went through a 15 minute real effort task. The real-effort task required subjects to type multiple combinations of letters and numbers into the keyboard. Subjects had 15 minutes to type in as many combinations as they could. For each correctly solved task, subjects received 5 cents. The purpose of the real effort task was to trigger memory constraints, such that previously formed beliefs (period 1) do not necessarily translate into second-period priors. While the memory literature contains many demonstrations that 15 minutes are sufficient to activate long-term memory and the corresponding memory constraints, in Section 7.1 we show the robustness of our findings when we increase the time lag to three days rather than 15 minutes. It is further worth mentioning that our procedure of working with twelve similar companies induces what psychologists call “interference”, which makes it harder for subjects to memorize and correctly attribute separate signals or beliefs. This effect likely contributes to our reduced-form assumption of imperfect memory, but ultimately our paper does not need to take a stand on where exactly imperfect memory comes from.

In the second period, for each company, subjects were shown up to one additional piece of news. Specifically, for ten companies, subjects received an additional piece of news, while for two companies, there were no additional news. We included two companies with no additional news because these allow us to directly assess whether subjects perfectly remember their first-period belief in the second period.

Immediately after observing the additional piece of news for a company, subjects

³All subjects saw three companies with three pieces of news, three with two pieces of news, three with one pieces of news and three with zero pieces of news.

were asked to state a second-period belief about the value of that company. Second-period beliefs constitute our main outcome of interest. In addition, on a subsequent decision screen, subjects were asked to recall the total number of positive and negative news that were shown to them in the course of the entire experiment for that company. These recall measures were not financially incentivized. Again, this procedure was repeated for all twelve companies. The experimental instructions and comprehension questions emphasized that first-period signals would be relevant for the second period.

The experiment was independently randomized across subjects across the following layers: (i) the order of companies in the first period; (ii) the order of companies in the second period; (iii) whether or not a company received a piece of news in the second period; and (iv) the actual signal realizations.

Beliefs were incentivized using a binarized scoring rule, which is incentive-compatible regardless of subjects' risk attitudes (Hossain and Okui, 2013). Under this scoring rule, subjects could potentially earn a prize of 10 euros. The probability of receiving the prize was given by 100 minus the squared distance between a subject's belief and the true value of the asset. In order to avoid hedging motives, at the end of the experiment one of the 24 beliefs was randomly selected for payment. Since second-period beliefs are our main outcome measure, we incentivized them more heavily, in expectation: with 90% probability a second-period belief was randomly selected for payment, and with 10% probability a first-period belief. To avoid extreme outliers due to typing mistakes, the computer program restricted beliefs to be in $[50, 150]$.

3.2 Treatment Variations and Sources of Exogenous Variation

We conducted five treatments, referred to as *Main*, *Reminder*, *No Cue*, *underreaction* and *Underreaction reminder*. In combination, these treatments allow for causal tests of all of the abstract predictions laid out in Section 2. That is, the treatments were designed to identify (i) potential overreaction in expectations; (ii) the ways in which the quantitative magnitude of such overreaction depends on the precise signal history; and (iii) the causal roles of imperfect and associative memory for overreaction; and (iv) the role of the correlation between context and news.

Treatment *Main*. In treatment *Main*, there is a one-to-one mapping between type of news (positive or negative) for a given company and the context with which the signal is communicated. That is, every positive news for company A is communicated with the same context (image and story). Likewise, every negative news for company A is communicated with the same context (albeit a different one than the positive news). The same logic holds for all other companies. Thus, it can never happen that a context

is communicated with news for different companies, or with both positive and negative news. A context deterministically identifies a piece of news.

Thus, treatment *Main* resembles our opening examples and implements a situation in which we hypothesize to observe overreaction. In addition, note that because the number and realizations of the signals vary across companies and subjects, the twelve tasks exhibit substantial variation in signal histories. We leverage this source of exogenous variation to test the within-treatment predictions derived in Section 2 about how the presence or quantitative magnitude of overreaction depends on the number of first-period signals that occurred in the same context (have the same realization) as the second-period signal. 80 subjects participated in this treatment.

Treatment *Reminder*. In treatment *Reminder*, we seek to remove subjects' memory constraints, holding everything else constant. The setup is exactly the same as in *Main*, except that at the beginning of the second period (i.e., before a subject observes the last signal for a company), subjects were reminded of their own first-period belief for that company. Thus, in contrast to treatment *Main*, we assist subjects in the recall of their first-period belief, so that they presumably no longer need to reconstruct their prior knowledge from memory. Conceptually, we think of this treatment as exogenously setting the parameter $r = 1$ in the framework of Section 2 (meaning perfect memory). Comparing treatments *Main* and *Reminder* allows us to cleanly identify the role of memory imperfection r for overreaction. 50 subjects participated in this treatment.

Treatment *NoCue*. Treatment *NoCue* was designed to isolate the role of associative recall. In terms of implementation, the setup in this treatment was exactly the same as in *Main*, except that each piece of news is communicated with a different context. That is, a given context (image and story) never appears twice, even if the company and type of news is identical. Thus, it is no longer the case that every positive news for a given company is communicated with the same context, and every negative news for a given company is communicated with the same context. As a consequence, stories and images can no longer trigger associative recall. At the same time, all other features of the environment remain unchanged. Conceptually, we think of this treatment as setting the associative recall parameter $a = 0$ (meaning no associative recall). Comparing treatments *Main* and *NoCue* therefore allows us to cleanly identify the role of associative recall a . 80 subjects participated in this treatment.

Treatments *Underreaction* and *Underreaction reminder*. All treatments described above relied on a design in which the observation of a positive piece of news in the second period cues the asymmetric recollection of positive first-period news (and analogously

for negative news), which corresponds to $\rho = 1$ in our formal framework. Treatments *Underreaction* and *Underreaction reminder* conceptually correspond to setting $\rho = -1$. In both treatments, the first period proceeded exactly as in treatment *Main*. In the second period, however, the remaining news were communicated on subjects' decision screens along with the *opposite* story and image, relative to the first period. That is, a positive piece of news for company A was communicated along with the story and image that were associated with negative news for company A in the first period of the experiment. Analogously, a negative piece of news for company A was communicated along with the story and the image that were associated with positive news for company A in the first period of the experiment. 80 subjects participated in this treatment.

The instructions in *Underreaction* emphasized that period 2 news were communicated along with the opposite story and image, and targeted control questions verified subjects' understanding of this. In addition, to allow for causal inference, we conducted condition *Underreaction reminder*. This treatment was identical to *Underreaction*, except that subjects were reminded of their own first-period belief right before they received the second-period signal for a company. To preview results, subjects in *Underreaction reminder* correctly incorporate period 2 news, suggesting that subjects had no problem understanding the *Underreaction* structure. In all other respects, these two treatments followed the same procedure as treatments *Main* and *Reminder*. 50 subjects took part in *Underreaction reminder*.

3.3 Predictions

Equation (5) in the conceptual framework directly suggests the following estimating equation for subject i 's second-period belief in task j :

$$b_2^{i,j} = \beta_1 n_{k+1}^j + \beta_2 \sum_{x=1}^k n_x^j + \epsilon \quad (8)$$

That is, we regress a subject's second-period belief on the value of the last signal as well as the first-period stock value (or the corresponding first-period belief). Appendix A.2 formally derives the properties of the OLS estimator $\hat{\beta}_1$ and shows that $E[\hat{\beta}_1] \approx 1 + \rho(1-r)a\bar{z}$, where $\bar{z}$ is the average number of first-period signals that were observed in the same context as the second-period signal. By applying the abstract predictions derived in Section 2 to this experimental design and estimating equation, we are hence ready to state the following predictions:

Predictions.

1. In treatment *Main*, we observe overreaction: $\hat{\beta}_1 > 1$.

Table 1: Mapping from model predictions to experimental predictions

Abstract model prediction	Treatments	Experimental prediction
1. Overreaction if news and context co-occur	<i>Main</i>	$\hat{\beta}_1^{Main} > 1$
2. Overreaction increases in # identical past contexts	<i>Main</i>	$\hat{\beta}_1^{Main}$ increases in z
3. Overreaction increases in imperfection of memory	<i>Main</i> & <i>Reminder</i>	$\hat{\beta}_1^{Main} > \hat{\beta}_1^{Reminder}$
4. Overreaction increases in strength of associative recall	<i>Main</i> & <i>NoCue</i>	$\hat{\beta}_1^{Main} > \hat{\beta}_1^{NoCue}$
5. Underreaction if news and context negatively corr.	<i>Underreaction</i>	$\hat{\beta}_1^{Under} < 1$
6. Underreaction increases in imperfection of memory	<i>Under.</i> & <i>Under. rem.</i>	$\hat{\beta}_1^{Under} < \hat{\beta}_1^{Under. rem.}$

2. *In treatment Main, overreaction increases in the number of first-period signals that were observed in the same context as the second-period signal.*
3. *Overreaction is stronger in treatment Main than in Reminder.*
4. *Overreaction is stronger in treatment Main than in No Cue.*
5. *In treatment Underreaction, we observe underreaction: $\hat{\beta}_1 < 1$. Moreover, underreaction increases in the number of first-period signals that were observed in the same context as the second-period signal.*
6. *Underreaction is stronger in treatment Underreaction than in Underreaction reminder.*

For clarity, Table 1 explicitly spells out which abstract model prediction from Section 2 maps into which specific experimental prediction, and which experimental treatments we use to test a given prediction.

3.4 Procedures and Logistics

Upon arrival in the lab, subjects received written instructions about the experiment.⁴ Subjects were given unlimited time to read the instructions and could ask questions at any point in time. After all subjects had indicated that they had finished the instructions, they completed a total of seven computerized control questions to verify adequate comprehension. Whenever a subject did not solve a control question correctly, a computer screen pointed out the mistake and explained the correct solution. As we pre-registered (see below), we exclude subjects from the analysis that answered more than one control question incorrectly (7% of potential participants).

⁴Appendix D contains the full set of instructions, translated into English. The Appendix also contains a link where readers can click through the entire experiment as subjects saw it (treatment *Main*, translated into English).

Treatments *Main*, *Reminder*, and *NoCue* were conducted in the BonnEconLab of the University of Bonn. Since we had exhausted the subject pool of the BonnEconLab, treatments *Underreaction* and *Underreaction reminder* were conducted in the University of Cologne’s Laboratory for Experimental Economics. Assignment to the relevant treatments was randomized within experimental sessions: *Baseline*, *Reminder*, and *NoCue* were all implemented in the same sessions, as were *Underreaction* and *Underreaction reminder*. We only compare treatments that were randomized within session. The experiments were computerized using Qualtrics. Experimental sessions lasted up to 90 minutes, and subjects earned an average of 15 euros.

3.5 Pre-Registration

All experiments in this paper were pre-registered in the AEA RCT registry, including a pre-analysis plan. The different pre-registration files include (i) the design of all treatments reported in this paper; (ii) the heterogeneity analysis discussed in Section 4.3; (iii) the regression equation (8) through which we analyze all data; (iv) all predictions outlined in Section 3.3; (v) the sample size in each treatment; (vi) that subjects would be dropped from the sample (and replaced) if they answer more than one control question incorrectly; and (vii) the labs in which we ran the experiments.

We proceeded in two steps. We first pre-registered treatments *Main*, *Reminder*, *NoCue*, as well as treatments *Time lag* and *Time lag reminder* (to be discussed in Section 7.1). Based on results from these treatment conditions, we pre-registered treatments *Underreaction* and *Underreaction reminder* (Section 6) and treatments *WTP* and *WTP Reminder* (Section 7.2). All pre-registration documents are available at <https://www.socialscienceregistry.org/trials/4247>.

4 Baseline Results on Overreaction

4.1 Preliminaries

Before we present the results, we conduct two checks on our experimental data. First, we verify people’s understanding of the experimental setup by investigating the accuracy of people’s beliefs at the end of the first period, before memory constraints become relevant. Table 7 in Appendix C shows that beliefs almost perfectly correspond to the true value of a company, in each of the three treatments (*Main*, *Reminder* and *No Cue*): in a regression of subjects’ beliefs on actual company values, the OLS coefficient is always almost exactly one and hence rational. The average percentage deviation between first-period beliefs and the truth is only 0.4%, while the median deviation is zero. This provides reassuring

evidence that subjects appear to understand the decision task well.

Second, to provide credence to our assumption that subjects can no longer perfectly remember their first-period belief once the second period starts, we consider the relationship between subjects' second-period and first-period beliefs in those tasks in which there was no second-period signal. Table 8 in Appendix C reports the results. In a regression of second-period on first-period beliefs, the OLS coefficient is only 0.56 and hence far away from the perfect memory benchmark of one. This suggests that memory is indeed imperfect in our setup, hence opening up a potential role for associative recall. Moreover, this regression coefficient is very similar in treatments *Main* and *No Cue*.

4.2 Treatment *Main*: Overreaction in Expectations

Throughout the empirical analysis, we present OLS regressions to test the hypotheses outlined in Section 3.3. Table 2 presents the results for treatment *Main*. In columns (1)–(3), we present three regression specifications. First, a regression in which we regress second-period beliefs on the value of the last signal (+10 or –10), controlling for the first-period belief. Second, an analogous regression in which we control for the objective first-period stock value as opposed to the first-period belief. Third, a comprehensive specification in which we control for experimental session fixed effects, first-period signal history fixed effects, company fixed effects, experimental order fixed effects, and subject fixed effects. In this third specification, controlling for first-period beliefs or stock values is redundant as these are implicitly accounted for by the first-period signal history fixed effects. In each regression specification, an observation corresponds to a subject-task, for a total of ten tasks per subject.⁵ Throughout, we cluster the standard errors at the subject level.

For all coefficients that are reported in Table 2, the rational prediction is that they equal one. The simple framework outlined in Section 2 instead predicts that the coefficient of first-period beliefs or stock values is less than one (due to imperfect memory) and that the coefficient of the last signal is greater than one (due to imperfect and associative memory). This is indeed what we find, see columns (1)–(3). In terms of magnitude, the OLS coefficient suggests that expectations substantially overreact, by 11 percent relative to the rational prediction of one. The last row of Table 2 reports the p-value for the null hypothesis that the coefficient of the second-period signal equals one. We reject this rational null hypothesis at all conventional levels of significance.

As highlighted by equation (6) in the formal framework, our hypothesis is that such overreaction occurs because the first-period signals get recollected more successfully if

⁵Naturally, and as specified in the pre-analysis plan, we restrict attention to those tasks in which a subject indeed received a signal in the second period.

Table 2: Treatment *Main*

	Dependent variable: 2nd period belief						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
2nd period signal	1.10*** (0.02)	1.11*** (0.02)	1.11*** (0.03)	0.87*** (0.04)	0.87*** (0.04)	0.85*** (0.04)	0.87*** (0.04)
Belief in 1st period	0.75*** (0.03)					0.59*** (0.05)	
Stock price in 1st period		0.74*** (0.03)					
Value of cued 1st period signals				0.92*** (0.03)	0.90*** (0.03)		
Value of non-cued 1st period signals				0.59*** (0.05)	0.59*** (0.05)		
2nd period signal × # 1st period signals in same context						0.34*** (0.05)	0.31*** (0.05)
Session FE	No	No	Yes	No	Yes	No	Yes
1st period signal history FE	No	No	Yes	No	No	No	Yes
Company FE	No	No	Yes	No	Yes	No	Yes
Order FE	No	No	Yes	No	Yes	No	Yes
Subject FE	No	No	Yes	No	Yes	No	Yes
Observations	800	800	800	800	800	800	800
Adjusted R^2	0.80	0.80	0.80	0.81	0.81	0.81	0.81
p-value for $H_0: \beta$ (2nd period signal) = 1:	<0.01	<0.01	<0.01	n/a	n/a	n/a	n/a

Notes. OLS estimates, robust standard errors (in parentheses) are clustered at the subject level. The sample includes all observations from treatment *Main* where subjects observed a second-period signal. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

they are cued, that is if they take on the same value as the second-period signal. To investigate this more explicitly, columns (4) and (5) of Table 2 include as separate regressors the overall value of those first-period signals that do (or do not) equal the second-period signal. The results show that expectations are much more responsive to the value of the cued first-period signals. Here, the difference in regression coefficients is statistically significant at all conventional levels.

Figure 2 visualizes the results. For each set of possible signal frequencies in the first period of the experiment, we regress second-period beliefs on the value of the second-period signal, and then plot the OLS coefficient and corresponding standard error. The figure shows that this coefficient is almost always larger than one. At the same time, there appears to be significant variation in the quantitative magnitude of this effect. Visual inspection suggests that the coefficient is increasing in the number of first-period signals. This is intuitive because if there are no past signals that can be cued, then trivially associative memory and asymmetric recall cannot generate overreaction. Thus, it is reassuring that expectations do not at all overreact in the case of zero positive and




Figure 2: OLS coefficient in a regression of second-period beliefs on the last signal, separately for each set of signal frequencies in the first part. The regressions control for a subject's first-period belief. Standard error bars are computed based on clustering at the subject level.

zero negative first-period signals (at the very left of Figure 2). In fact, if anything, this coefficient is smaller than one, consistent with a large literature on belief updating that shows that in lab environments where the role of (associative) memory is shut down, people's belief updating typically exhibits conservatism ([Benjamin, 2018](#)).

This discussion directly relates back to Prediction 2 in Section 3.3, which posits that the magnitude of the regression coefficient should be increasing in the number of first-period signals that take on the same value as the second-period signal. This is a direct test of the role of associative memory because with either perfect memory ($r = 1$ in the model) or imperfect but no associative memory ($a = 0$), this prediction would not hold, compare equation (5). Moreover, while a particular form of recency bias could in principle generate the type of overreaction we observe in Table 2, this is not the case for the prediction that overreaction depends on the signal history in specific ways.

Columns (6)–(7) provide a formal statistical test of Prediction 2. Here, we interact the value of the second-period signal with the number of first-period signals that were communicated in the same context (z in the model). The regression table suppresses the corresponding raw term for brevity. The results show that the interaction term of interest is consistently positive and statistically highly significant. The magnitude suggests that each additional first-period signal increases the responsiveness to the second-period signal by about 30%, on average. Moreover, once the interaction term is accounted for, the regression coefficient of the second-period signal (which now econometrically corre-




Figure 3: Recall of first-period signals in Treatment *Main*, depending on whether the second-period signal was identical to or different from the first-period signals. We construct the recall variables as follows. In the case of recall of signals that are different from the second-period signal, we use the reported recall quantity. In the case of recall of signals that are identical to the second-period signal, we use the reported recall minus one. That is, in constructing the figure we make the arguably very plausible assumption that subjects always remember the value of the second-period signal that they just saw a few seconds ago.

sponds to the case of no cued first-period signals) is less than one. This result is analogous to the discussion of Figure 2: when no signals are cued, expectations do not overreact but – if anything – even underreact (Benjamin, 2018).

The evidence in support of Prediction 2 suggests that associative memory does not only help people in recalling that they saw such a signal in the first period, but also *how often* such a signal occurred. To corroborate this result, we next turn to subjects' direct recall data. Figure 3 shows average levels of reported recall of first-period signals in condition *Main*, as a function of whether these first-period signals were identical to or different from the second-period signal. That is, the figure shows how many signals subjects report to have recalled, as a function of whether those signals were cued or not. The figure shows that the recall of cued signals is very accurate, on average. In contrast, the recall of non-cued signals is more compressed: subjects overestimate the frequency of signals that did not appear at all and they substantially underestimate (under-recall) the frequency of signals that appeared often. This is indicative that associative memory helps not only with remembering *whether* a certain type of signal has appeared before, but also *how often* it appeared.




Figure 4: Subject-level distribution of regression coefficients of last signal in treatment *Main* (N=80). To estimate these coefficients, we run regressions akin to column (1) in Table 2 except that in each regression the sample is restricted to only one subject. Moreover, to adequately identify overreaction in the presence of cued recall, the sample is restricted to tasks in which a subject observed at least one first-period signal. A rational subject would have a coefficient of one.

4.3 Heterogeneity Analysis

Next, we examine across-subject heterogeneity in overreaction. To estimate the presence of such heterogeneity, we require a measure of a subject’s overall overreaction “type.” To this effect, we run our standard regression of second-period beliefs on the second-period signal, but now separately for each subject. That is, we estimate a subject-specific responsiveness parameter that should be equal to one if a subject is either rational or has imperfect-but-no-associative-memory.

Figure 4 presents the distribution of types. While the beliefs of a notable fraction of subjects do not reflect associative recall (35% have a regression coefficient of at most one), the majority of participants exhibit overreaction to varying degrees.

A natural question is what explains this heterogeneity. To investigate this, we turn to three heterogeneity analyses, all of which were specified in our pre-analysis plan: (i) performance on a Raven matrices IQ test that was administered at the end of the experiment; (ii) a measure of the strength of memory that is estimated from the experimental recall data as a proxy for r : for each subject, we regress the reported recall of non-cued signals on the actual number of corresponding signals and save this regression coefficient as a measure of the strength of (non-cued) memory; and (iii) response times.

Tables 10 and 11 in Appendix C reports the results. We find that subjects with higher

Raven scores and better non-cued recall exhibit less overreaction. The relationship between overreaction and response times is negative, but not statistically significant.

5 Exogenous Memory Manipulations

5.1 The Role of Imperfect Memory

To provide causal evidence for the role of imperfect memory in expectation overreaction, we manipulate the extent to which memory can actually play in the role in the experiment. Conceptually, treatment *Reminder* is designed to set $r = 1$. To this effect, we reminded participants of their first-period belief immediately before they received the second-period signal.

Given the explicit focus on memory in this section, the analysis considers both (i) the financially incentivized second-period beliefs and (ii) the unincentivized recall data. As specified in the pre-analysis plan, we analyze the recall data by computing the difference between recall of positive and recall of negative news and multiplying this difference by 10 so that the variable has the same scale as the beliefs data. This summary statistic of a subject's recall is highly correlated with actual second-period beliefs ($\rho = 0.95$), suggesting that the recall data are meaningful.⁶

Table 3 presents the results of the treatment comparison between *Main* and *Reminder*. As specified in the pre-registration, we again analyze our data by means of OLS regressions in which we relate subjects' second-period beliefs (or recall) to the value of the second-period signal, except that now we also interact the second-period signal with a treatment dummy. The table also includes a treatment dummy, which is suppressed for brevity. Our prediction, spelled out in Sections 2 and 3.3, is that the value of the second-period signal should matter more in treatment *Main* than in *Reminder*.

The results provide supporting evidence for this prediction. The interaction term is quantitatively large and statistically significant at all conventional levels. This is true when we consider the beliefs data as well as when we directly look at reported recall, compare columns (4)–(6). In *Main*, subjects respond 12–14% more to the value of the last signal than subjects in *Reminder*. Again, this pattern is a specific prediction of our framework, but not of an account of recency effects.

Moreover, as we can see from the regression coefficient of the second-period signal (which captures the coefficient in treatment *Reminder* only), in treatment *Reminder* we find no overreaction: the OLS coefficient is 0.94–0.99, if anything even less than one. These results provide direct causal evidence that imperfect memory is necessary in or-

⁶Table 9 in Appendix C replicates the results of Tables 2 for this summary statistic of recall, as specified in the pre-analysis plan.

Table 3: Treatments *Main* vs. *Reminder*

	<i>Dependent variable:</i>					
	2nd period belief			Δ Recall [Pos. – Neg.]		
	(1)	(2)	(3)	(4)	(5)	(6)
2nd period signal	0.99*** (0.01)	0.98*** (0.01)	0.98*** (0.01)	0.95*** (0.02)	0.95*** (0.02)	0.94*** (0.02)
2nd period signal $\times$ 1 if <i>Main</i> , 0 if <i>Reminder</i>	0.12*** (0.03)	0.13*** (0.03)	0.14*** (0.03)	0.11** (0.04)	0.12*** (0.04)	0.13*** (0.04)
Belief in 1st period	0.84*** (0.02)			0.83*** (0.02)		
Stock price in 1st period		0.84*** (0.02)			0.83*** (0.02)	
Treatment FE	Yes	Yes	Yes	Yes	Yes	Yes
Session FE	No	No	Yes	No	No	Yes
1st period signal history FE	No	No	Yes	No	No	Yes
Company FE	No	No	Yes	No	No	Yes
Order FE	No	No	Yes	No	No	Yes
Subject FE	No	No	Yes	No	No	Yes
Observations	1300	1300	1300	1300	1300	1300
Adjusted R^2	0.86	0.86	0.86	0.82	0.83	0.83

Notes. OLS estimates, robust standard errors (in parentheses) are clustered at the subject level. The sample includes all observations from treatments *Main* and *Reminder* where subjects observed a second-period signal. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

der for overreaction to arise in our setup, as predicted by our key equation (5). Figure 5 visualizes this result and the difference between *Main* and *Reminder*. Again, we see that there is no overreaction in treatment *Reminder*, which is reassuring evidence that our experimental setup is not misconstrued by subjects: in the absence of memory constraints, the second-period signal is incorporated in a rational fashion.

5.2 The Role of Associative Memory

To provide direct causal evidence for the role of associative recall, we proceed by directly manipulating the strength of associativeness of memory. According to equation (5), if there is no associative recall ($a = 0$) there should be no overreaction.

As a direct test of this hypothesis, we compare treatments *Main* and *NoCue*. Table 4 presents the results and follows a similar logic as Table 3. Again, as specified in the pre-analysis plan, we link participants' second-period beliefs to the second-period signal, interacted with a treatment dummy. Again the table also includes a treatment dummy, which is suppressed for brevity.

The regression results document that, as predicted, the interaction term is positive




Figure 5: OLS coefficient in a regression of second-period beliefs on the last signal, separately for each set of signal frequencies in the first part. The regressions control for a subject’s first-period belief. Standard error bars are computed based on clustering at the subject level.

and statistically significant. On average, subjects in *Main* react 21–30% more to the value of the second-period signal than subjects in *NoCue*. This is true for both the beliefs and the direct recall data. These results provide direct causal evidence for the role of associative recall in generating overreaction. Figure 6 visualizes the difference in regression coefficients between these two treatments.

Notice that Figure 6 also reveals some degree of conservatism in updating in treatment *NoCue*. This is analogous to findings in treatment *Main* that when no signals are cued, expectations do not overreact but, if anything, even underreact. It is also in line with a myriad of findings from updating experiments where (associative) memory cannot play a role (Benjamin, 2018). A potential explanation for this is offered in (Enke and Graeber, 2019). They argue and provide evidence that conservatism in updating might be driven by people’s “internal uncertainty”, which leads them to overweight their prior and underweight new information in situations of uncertainty. Notably, such an account would also explain why we do not see underreaction in treatment *Reminder*. When subjects are reminded of the past, uncertainty in our experimental environment is arguably very low, and hence people put appropriate weight on new information.

Table 4: Treatments *Main* vs. *NoCue*

	<i>Dependent variable:</i>					
	2nd period belief			Δ Recall [Pos. – Neg.]		
	(1)	(2)	(3)	(4)	(5)	(6)
2nd period signal	0.89*** (0.04)	0.88*** (0.04)	0.88*** (0.04)	0.76*** (0.05)	0.75*** (0.05)	0.77*** (0.06)
2nd period signal $\times$ 1 if <i>Main</i> , 0 if <i>NoCue</i>	0.21*** (0.04)	0.22*** (0.04)	0.22*** (0.05)	0.29*** (0.06)	0.29*** (0.06)	0.30*** (0.07)
Belief in 1st period	0.62*** (0.03)			0.60*** (0.03)		
Stock price in 1st period		0.62*** (0.03)			0.61*** (0.03)	
Treatment FE	Yes	Yes	Yes	Yes	Yes	Yes
Session FE	No	No	Yes	No	No	Yes
1st period signal history FE	No	No	Yes	No	No	Yes
Company FE	No	No	Yes	No	No	Yes
Order FE	No	No	Yes	No	No	Yes
Subject FE	No	No	Yes	No	No	Yes
Observations	1600	1600	1600	1600	1600	1600
Adjusted R^2	0.68	0.68	0.67	0.62	0.63	0.63

Notes. OLS estimates, robust standard errors (in parentheses) are clustered at the subject level. The sample includes all observations from treatments *Main* and *No Cue* where subjects observed a second-period signal. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

6 Over- vs. Underreaction

Next, we turn to investigating predictions 6 and 7 in Section 3.3, which conceptually correspond to exogenously varying the parameter ρ in the simple model. For this purpose, as discussed in Section 3, we implemented treatments *Underreaction* and *Underreaction reminder*. Here, second-period signals were communicated with those contexts that belonged to the respective opposite signal in the first period. Thus, a signal now cues the recollection of the opposite past signals.

Table 5 presents the results. Columns (1) and (2) show that, within treatment *Underreaction*, the coefficient of the second-period signal is 0.74–0.76, substantially smaller than one, indicating meaningful underreaction to the second-period signal. Columns (3) and (4) leverage exogenous variation in signal histories to document that underreaction strongly increases in the number of first-period signals that were communicated in the same context as the second-period signal, see the negative and statistically significant interaction term.⁷

⁷Table 12 in Appendix C shows that very similar results hold when we consider the recall data rather than subjects' beliefs.




Figure 6: OLS coefficient in a regression of second-period beliefs on the last signal, separately for each set of signal frequencies in the first part. The regressions control for a subject's first-period belief. Standard error bars are computed based on clustering at the subject level.

These results suggest that underreaction is generated by asymmetric recall. To further corroborate this claim, Figure 11 in Appendix B analyzes the self-reported recall patterns in treatment *Underreaction* as a function of the signal history, akin to Figure 3 in Section 4. Here, we see that, in contrast to treatment *Main*, subjects' recall is much more precise for those first-period signals that *differ* in value from the second-period signal than for those signals that take on the same realization as the second-period signal. Again, this pattern is expected because those first-period signals that take on a different value as the second-period signal get cued by the second-period context.

Finally, columns (5) and (6) compare treatments *Underreaction* and *Underreaction reminder*. As in Section 5.1, the coefficient of interest is now the interaction term between the second-period signal and a treatment dummy. The dummy is negative and statistically highly significant. The magnitude suggests that underreaction is 25–28% stronger in *Underreaction*. In contrast, as we can infer from the coefficient of the second-period signal, there is no underreaction in treatment *Underreaction reminder*, with a coefficient of 1.01, statistically indistinguishable from one. This pattern is again predicted because in our framework associative asymmetric recall cannot occur if memory is exogenously set to have no imperfections ($r = 1$).

Figure 7 visualizes the results in both treatments, separately for each signal history. The figure confirms that underreaction is present at all signal histories except for when there are no first-period signals, as predicted by our formal framework. In summary, the

Table 5: Treatments *Underreaction* and *Underreaction reminder*

	Dependent variable: 2nd period belief					
	Treatments:				+ Reminder	
	<i>Underreaction</i>					
	(1)	(2)	(3)	(4)	(5)	(6)
2nd period signal	0.76*** (0.04)	0.74*** (0.04)	0.99*** (0.04)	0.95*** (0.05)	1.01*** (0.02)	1.01*** (0.02)
Belief in 1st period	0.65*** (0.04)		0.50*** (0.05)		0.77*** (0.03)	
2nd period signal × # 1st period signals in same context			-0.31*** (0.06)	-0.28*** (0.07)		
2nd period signal × 1 if <i>Underreaction</i> , 0 if <i>Reminder</i>					-0.25*** (0.04)	-0.28*** (0.05)
Treatment FE	No	No	No	No	Yes	Yes
Session FE	No	Yes	No	Yes	No	Yes
Signal history FE	No	Yes	No	Yes	No	Yes
Company FE	No	Yes	No	Yes	No	Yes
Order FE	No	Yes	No	Yes	No	Yes
Subject FE	No	Yes	No	Yes	No	Yes
Observations	800	800	800	800	1300	1300
Adjusted R^2	0.67	0.68	0.68	0.70	0.79	0.79

Notes. OLS estimates, robust standard errors (in parentheses) are clustered at the subject level. The table suppresses the coefficients of the number of identical first-period signals (columns (3)–(4)) and of a binary indicator for whether the signals are mixed (columns (5)–(6)). The sample includes all observations from treatments *Underreaction* and *Underreaction reminder* where subjects observed a second-period signal. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

results show that associative memory can generate predictable patterns of both over- or underreaction, depending on which contexts a piece of news is associated with.

In summary, Sections 4, 5, and 6 have subjected equation (5) in Section 2 to a systematic experimental test. Based on initial evidence for overreaction in expectations (Table 2), we have exogenously manipulated the number of first-period signals that are identical to the second-period signal; we exogenously varied both the relevance of imperfect memory (Table 3) and the strength of associative recall (Table 4); and we exogenously varied the nature of the correlation between news and contexts (Table 5). Throughout, the results are in line with the predictions of equation (5).




Figure 7: OLS coefficient in a regression of second-period beliefs on the last signal, separately for each set of signal frequencies in the first part. The regressions control for a subject's first-period belief. Standard error bars are computed based on clustering at the subject level.

7 Robustness Treatments

7.1 Extended Time Lag

All treatments reported up to this point relied on a design in which the time lag between the first and second period was 15 minutes. As an extension and robustness check, we now investigate whether the effects of associative memory on expectation formation also prevail under a slightly longer time lag. We conducted treatments *Time lag* and *Time lag reminder*. These treatments followed the same procedure as treatments *Main* and *Reminder*, except that the time lag between the first and second period of the experiment was three days. When subjects signed up for the experiment, they registered for two separate sessions that were conducted at the same time of day, on a Tuesday and Friday.

On the first day, subjects completed the first period of the experiment, using the same experimental instructions and control questions as in the baseline treatments reported above (with obvious minor modifications regarding the timing of the second period). Thus, throughout it was obvious to participants that the information in the first period would be relevant for the second period three days later. After the first period, participants completed the real effort task, the Raven matrices test as well as the demographic questionnaire.⁸ On the second day, participants re-read the original instructions and

⁸Thus, subjects could not take notes right after the first period.




Figure 8: OLS coefficient in a regression of second-period beliefs on the last signal, separately for each set of signal frequencies in the first part. The regressions control for a subject's first-period belief. Standard error bars are computed based on clustering at the subject level.

completed the same set of control questions again. Then, they completed the second period of the experiment.

These two treatments were also pre-registered in the original pre-registration described in Section 3.5. 80 subjects participated in treatment *Time lag* and 50 in treatment *Time lag reminder*. The treatments were randomized within experimental sessions and implemented in the BonnEconLab of the University of Bonn. The sessions for the first period lasted up to 60 minutes and those for the second period up to 45 minutes. Average earnings were 25 euros, including a 10 euros show-up fee. Attrition was negligible: 95% of subjects returned for the second session.

Figure 8 summarizes the results, which are very similar to those in treatments *Main* and *Reminder*: we see (i) overreaction; (ii) stronger overreaction when more first-period signals get cued by the second-period signal; and (iii) stronger overreaction relative to a treatment with a reminder. Tables 13 and 14 in Appendix C present corresponding regression analyses.

7.2 Choices

All experiments reported above focus on the elicitation of participants' beliefs and corresponding recall data. In a final robustness check, we investigate whether similar results hold when we instead consider explicit economic actions. To this effect, we im-

plemented treatments *WTP* and *WTP reminder*, which were also pre-registered. These treatments follow the same structure as conditions *Main* and *Reminder*, except that we do not elicit participants' beliefs about the value of the hypothetical companies, neither in the first nor in the second period of the experiment. Neither do we elicit subjects' recall of positive and negative signals. Instead, in both the first and the second period of the experiment, subjects were endowed with 150 points each and then stated their willingness-to-pay (WTP) for the stock of a company. To this effect, we implemented a direct Becker-deGroot-Marschak elicitation mechanism, such that subjects directly entered the maximum number of points m that they would be willing to pay for an asset. We then randomly determined a price $p \sim U[50, 150]$ and subjects received the asset if $m \geq p$ and kept their endowment otherwise. In treatment *WTP reminder*, subjects were reminded of their first-period WTP right before they received the second-period signal for a company. These treatments are arguably meaningful robustness checks because the experimental instructions are stripped from almost all references to beliefs, and instead only talk about subjects' choices. 100 subjects participated in treatment *WTP* and 80 in *WTP reminder*, which were randomized within sessions and implemented in the University of Cologne's Laboratory for Experimental Economics.

We analyze these data by means of the same OLS regressions as previously, except that now the dependeng variable is a participant's second-period WTP rather than their belief about the value of a company. Figure 9 summarizes the results, which are very similar to those in treatments *Main* and *Reminder*: we see (i) overreaction; (ii) stronger overreaction when more first-period signals get cued by the second-period signal; and (iii) stronger overreaction relative to a treatment with a reminder. Table 15 in Appendix C presents corresponding regression analyses.

8 Estimating the Model

All analyses reported up to this point are motivated and structured through the formal framework laid out in Section 2. To supplement these reduced-form analyses, we now explicitly estimate this model, in particular its key memory parameters.

Specifically, we estimate the parameters $\hat{\beta}$, $\hat{r}$, and $\hat{a}$ by minimizing the sum of squared residuals for the non-linear regression equation (compare equation (4) above):

$$b = 100 + \beta n_{k+1} + r \sum_{x=1}^k n_x + (1-r)a \sum_{x=1}^z n_x + \epsilon \quad (9)$$

Table 6 summarizes the estimates across treatments, where our interest is in the parameters that govern the strength of memory $\hat{r}$ as well as of associative recall $\hat{a}$.




Figure 9: OLS coefficient in a regression of second-period willingness-to-pay on the last signal, separately for each set of signal frequencies in the first part. The regressions control for a subject's first-period willingness-to-pay. Standard error bars are computed based on clustering at the subject level.

Aggregate data. To start, we estimate this equation on the aggregate data across subjects, separately for each treatment. The quantitative estimates, reported in Table 6, resonate with the results documented above and provide interesting cross-treatment comparisons. In treatment *Main*, we estimate a substantial role for associative recall. The estimates imply that participants recall non-cued signals with probability 59% and cued ones with probability 91%. In treatment *Reminder*, we confirm that imperfect memory entirely disappears (by construction of the treatment), so that associative recall cannot be identified with reasonable precision (compare the huge standard error). Similar patterns prevail in *Underreaction reminder* and *Time lag reminder*, where the estimated associative recall parameters are very noisy, as should be the case from the perspective of equation (9). Analogously, we see that in treatment *NoCue*, associative recall collapses to zero, again by construction of the treatment.

Finally, note that estimated memory imperfection is larger in conditions *Underreaction* and *Time lag* than in *Baseline*. This is intuitive since (i) in *Underreaction* participants had to deal with the additional difficulty that the mapping between news and contexts changed between the first and second period and (ii) in *Time lag*, subjects had to memorize news for three days rather than 15–30 minutes.

Individual-level data and model fit. Next, we estimate the same model, separately for each individual. To assess the fit of the model at the individual level, we then use the

Table 6: Estimates of model parameters across treatments

Treatment	Memory imperfection ($1 - \hat{r}$)	Associative recall $\hat{a}$
<i>Main</i>	0.41*** (0.03)	0.79*** (0.08)
<i>Reminder</i>	0.01 (0.01)	-1.59 (4.75)
<i>NoCue</i>	0.51*** (0.04)	0.01 (0.12)
<i>Underreaction</i>	0.68*** (0.03)	0.35*** (0.06)
<i>Underreaction reminder</i>	0.05*** (0.01)	0.73** (0.33)
<i>Time lag</i>	0.51*** (0.03)	0.61*** (0.07)
<i>Time lag reminder</i>	0.02 (0.02)	0.31 (1.06)

Notes. Estimates of equation (9), standard errors (clustered at subject level) reported in parentheses. The model is estimated by pooling the data across subjects in a given treatment. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

individual-level estimates of $\hat{r}_i$ and $\hat{a}_i$ to predict participant i 's reported recall of those first-period news that did (q_i^c) or did not (q_i^n) get cued by the second-period signal:

$$\hat{q}_i^n = \hat{r}_i * (k - z) \quad (10)$$

$$\hat{q}_i^c = [\hat{r}_i + (1 - \hat{r}_i)\hat{a}_i] * z \quad (11)$$

where z again denotes the number of first-period signals that were communicated in the same context as the second-period signal, and k the total number of first-period signals. Note that the recall data do not enter the estimation and prediction procedure because the memory parameters are estimated only from the beliefs data.

Nevertheless, we find that, within treatment *Main*, the correlation between predicted and actual recall of those signals that got cued by the second-period signal is $\rho = 0.82$. Figure 10 visualizes this pattern. The correlation between predicted and actual recall of those signals that did not get cued is $\rho = 0.67$. We interpret these results as encouraging evidence that our simple two-parameter memory model fits the observed data well, and that subjects exhibit a substantial level of across-task consistency in how they use their own memory technology to recall news and estimate stock prices in our experiment.




Figure 10: Relationship between recall patterns as predicted by the model estimates and actual recall patterns. The figures represent binned scatter plots that average observed recall for a given level of (rounded) predicted recall. Predicted recall is computed by first estimating equation (9) and then applying equations (10) and (11).

9 Discussion

This paper has provided a theoretically-structured experimental analysis of the roles of imperfect and associative memory for belief formation. The notion of associative recall has recently received increased attention from economic theorists, yet direct experimental evidence on the importance of cued recall in structured economic decision environments is limited.

In our experiments, (i) participants' beliefs strongly overreact to the latest piece of news; (ii) the presence and magnitude of such overreaction depend on the precise signal history in predictable ways; (iii) exogenously manipulating the degree of memory imperfection provides causal evidence that without imperfect memory, overreaction does not occur; (iv) exogenously manipulating the strength of associative recall provides causal evidence that in our context associative memory is necessary in order for overreaction to arise; (v) associative memory generates predictable over- or underreaction in expectations, depending on the precise mapping that links types of news with certain contexts; (vi) the effect of associative memory on expectation overreaction holds both over a rather short (15 minutes) and a somewhat longer (three days) time horizon; (viii) we identify similar patterns for participants' beliefs and their purchasing decisions; and (ix) a direct estimation of our simple model suggests that associative memory plays a quantitatively large role in generating observed beliefs. In combination, we view our experiments as clean and theoretically-founded experimental investigation of the relevance of associative recall for belief formation. We believe that by offering a new experimental paradigm in which these types of effects can be studied, our paper opens up the possibility for further experimental research in an agenda on memory imperfections and belief formation.

References

- Akerlof, George A. and Robert J. Shiller**, *Animal spirits: How human psychology drives the economy, and why it matters for global capitalism*, Princeton University Press, 2009.
- Augenblick, Ned and Eben Lazarus**, “Restrictions on Asset-Price Movements Under Rational Expectations: Theory and Evidence,” *UC Berkeley-Haas and Harvard University Mimeo*, 2018.
- **and Matthew Rabin**, “Belief movement, uncertainty reduction, and rational updating,” *UC Berkeley-Haas and Harvard University Mimeo*, 2018.
- Benjamin, Daniel J**, “Errors in Probabilistic Reasoning and Judgmental Biases,” in “Handbook of Behavioral Economics” 2018.
- Bodoh-Creed, A**, “Mood, memory, and the evaluation of asset prices,” *Review of Finance*, 2013.
- Bordalo, Pedro, Katherine Coffman, Nicola Gennaioli, Frederik Schwerter, and Andrei Shleifer**, “Memory and representativeness,” *Working Paper*, 2019.
- , **Nicola Gennaioli, and Andrei Shleifer**, “Memory, Attention, and Choice,” *Working Paper*, 2017.
- , — , **Rafael La Porta, and Andrei Shleifer**, “Diagnostic expectations and stock returns,” Technical Report, National Bureau of Economic Research 2017.
- , — , **Yueran Ma, and Andrei Shleifer**, “Over-reaction in macroeconomic expectations,” Technical Report, National Bureau of Economic Research 2018.
- Carlson, Ryan W, Michel Marechal, Bastiaan Oud, Ernst Fehr, and Molly Crockett**, “Motivated misremembering: Selfish decisions are more generous in hindsight,” 2018.
- Dertwinkel-Kalt, Markus, Holger Gerhardt, Gerhard Riener, Frederik Schwerter, and Louis Strang**, “Concentration bias in intertemporal choice,” *Working Paper*, 2017.
- Enke, Benjamin**, “What you see is all there is,” *Working Paper*, 2017.
- **and Florian Zimmermann**, “Correlation Neglect in Belief Formation,” *Review of Economic Studies*, 2019.
- **and Thomas Graeber**, “Cognitive Uncertainty,” *Working Paper*, 2019.
- Esponda, Ignacio and Emanuel Vespa**, “Hypothetical Thinking: Revisiting Classic Anomalies in the Laboratory,” *Working Paper*, 2016.

- Frydman, Cary and Lawrence J Jin**, “Efficient coding and risky choice,” *Working Paper*, 2018.
- Fudenberg, Drew**, “Advancing Beyond “Advances in Behavioral Economics”,” *Journal of Economic Literature*, 2006, 44 (3), 694–711.
- , **David K Levine et al.**, “Learning with recency bias,” *Proceedings of the National Academy of Sciences*, 2014, 111, 10826–10829.
- Gennaioli, Nicola and Andrei Shleifer**, “What Comes to Mind,” *Quarterly Journal of Economics*, 2010, 125 (4), 1399–1433.
- Herz, Holger and Dmitry Taubinsky**, “What makes a price fair? An experimental study of transaction experience and endogenous fairness views,” *Journal of the European Economic Association*, 2017, 16 (2), 316–352.
- Hossain, Tanjim and Ryo Okui**, “The binarized scoring rule,” *Review of Economic Studies*, 2013, 80 (3), 984–1001.
- Kahana, Michael Jacob**, *Foundations of human memory*, OUP USA, 2012.
- Koszegi, Botond, George Loewenstein, and Takeshi Murooka**, “Fragile Self-Esteem,” *Working Paper*, 2019.
- Malmendier, Ulrike and Stefan Nagel**, “Learning from inflation experiences,” *The Quarterly Journal of Economics*, 2015, 131 (1), 53–87.
- Martínez-Marquina, Alejandro, Muriel Niederle, and Emanuel Vespa**, “Probabilistic States versus Multiple Certainties: The Obstacle of Uncertainty in Contingent Reasoning,” Technical Report, National Bureau of Economic Research 2017.
- Mullainathan, Sendhil**, “A Memory-Based Model of Bounded Rationality,” *Quarterly Journal of Economics*, 2002, pp. 735–774.
- Shiller, Robert J**, “Narrative economics,” *American Economic Review: Papers and Proceedings*, 2017, 107 (4), 967–1004.
- Wachter, Jessica A and Michael J Kahana**, “A retrieved-context theory of financial decisions,” *Available at SSRN*, 2019.
- Wilson, Andrea**, “Bounded memory and biases in information processing,” *Econometrica*, 2014, 82 (6), 2257–2294.
- Zimmermann, Florian**, “The dynamics of motivated beliefs,” *American Economic Review*, forthcoming.

ONLINE APPENDIX

A Additional Derivations

A.1 Partial Naïveté

The two forms of naivete would be captured in slightly different ways:

- Suppose the DM to some extent (captured by naivete parameter α , $\alpha = 0$ captures full naivete, $\alpha = 1$ captures full sophistication) fails to realize that he sometimes forgets. When he does realize, however, then he correctly (in a Bayesian sense) infers the realization of the information based on his memory parameters r and a . Specifically, $Pr(n_{nr} = n_{k+1} | norecall) = \frac{1-a}{2-a}$. The DM's forecast would thus be given by $f = v + n_{k+1} + \sum_{x=1}^k m_x n_x + \alpha \sum_{x=1}^k (1 - m_x) \frac{-a}{2-a} n_{k+1}$
- Suppose the DM instead fully realizes that he sometimes forgets, but is naive in the way he infers what he forgot. This form of naivete could be captured by DM's belief $\hat{a}$ about memory parameter a , $\hat{a} \leq a$ ($\hat{a} = 0$ captures full naivete, $\hat{a} = 1$ captures full sophistication). The DM would do inference as outlined above, but would use $\hat{a}$ instead of a to infer the value of signals he does not recall. The DM's forecast would thus be given by $f = v + n_{k+1} + \sum_{x=1}^k m_x n_x + \sum_{x=1}^k (1 - m_x) \frac{-\hat{a}}{2-\hat{a}} n_{k+1}$.

A.2 Derivation of OLS Estimator

We formally derive the relationship between equation (5) and the OLS estimator $\hat{\beta}_1$ in equation (8). Keeping the notation that b is the belief and n the news, then with N observations (subject-tasks) the OLS estimator is given by

$$E[\hat{\beta}] = E \left[\frac{\sum n_i b_i - 1/N \sum n_i \sum b_i}{\sum n_i^2 - 1/N (\sum n_i^2)} \right] \quad (12)$$

This expectation of a ratio can be approximated by the ratio of the expectations (also, the expectation of a ratio equals the ratio of probability limits). Denote $c = v + r \sum n_x$, which is not a function of the last signal. Substitute in for the forecast. Observing that

$E[n_i] = 0$, we get

$$E[\hat{\beta}] = \frac{\sum E[n_i b_i] - 1/N \sum E[n_i] \sum E[b_i]}{\sum E[n_i^2] - 1/N E(\sum n_i^2)} \quad (13)$$

$$= \frac{\sum E[n_i [n_i (1 + z_i (1 - r) \rho a) + c]]}{\sum E[n_i^2]} \quad (14)$$

$$= 1 + (1 - r) \rho a \frac{\sum E[n_i^2 z_i]}{\sum E[n_i^2]} \quad (15)$$

$$= 1 + (1 - r) \rho a \frac{\sum E[n_i^2] E[z_i]}{\sum E[n_i^2]} \quad (16)$$

$$= 1 + (1 - r) \rho a \frac{\bar{z} \sum E[n_i^2]}{\sum E[n_i^2]} \quad (17)$$

$$= 1 + (1 - r) \rho a \bar{z} \quad (18)$$

Because z_i and n_i are independent.

B Additional Figures




Figure 11: Recall of first-period signals in Treatment *Underreaction*, depending on whether the second-period signal was identical to or different from the first-period signals. We construct the recall variables as follows. In the case of recall of signals that are different from the second-period signal, we use the reported recall quantity. In the case of recall of signals that are identical to the second-period signal, we use the reported recall minus one. That is, we make the arguably very plausible assumption that subjects always remember the value of the second-period signal that they just saw a few seconds ago.

C Additional Tables

Table 7: Beliefs in the first period

	<i>Dependent variable: 1st period belief</i>					
	<i>Main</i>		<i>Reminder</i>		<i>No Cue</i>	
	(1)	(2)	(3)	(4)	(5)	(6)
Stock price in 1st period	0.98*** (0.01)	0.97*** (0.01)	1.00*** (0.00)	1.00*** (0.00)	0.99*** (0.01)	1.00*** (0.01)
Session FE	No	Yes	No	Yes	No	Yes
Company FE	No	Yes	No	Yes	No	Yes
Order FE	No	Yes	No	Yes	No	Yes
Subject FE	No	Yes	No	Yes	No	Yes
Observations	800	800	500	500	800	800
Adjusted R^2	0.96	0.97	1.00	1.00	0.99	0.99

Notes. OLS estimates, robust standard errors (in parentheses) are clustered at the subject level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 8: Beliefs in the second period in case of no signal in second period

	<i>Dependent variable: 2nd period belief</i>					
	<i>Main</i>		<i>Reminder</i>		<i>No Cue</i>	
	(1)	(2)	(3)	(4)	(5)	(6)
Belief in 1st period	0.56*** (0.09)	1.00*** (0.00)	1*** (0.00)	1.00*** (0.00)	0.51*** (0.08)	1.37* (0.81)
Session FE	No	Yes	No	Yes	No	Yes
1st period signal history FE	No	Yes	No	Yes	No	Yes
Company FE	No	Yes	No	Yes	No	Yes
Order FE	No	Yes	No	Yes	No	Yes
Subject FE	No	Yes	No	Yes	No	Yes
Observations	160	160	100	100	160	160
Adjusted R^2	0.39	1.00	1.00	1.00	0.36	0.34

Notes. OLS estimates, robust standard errors (in parentheses) are clustered at the subject level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 9: Treatment *Main*: Recall data

	<i>Dependent variable:</i> Δ Recall [Pos. – Neg.]						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
2nd period signal	1.05*** (0.04)	1.06*** (0.03)	1.07*** (0.04)	0.81*** (0.05)	0.83*** (0.05)	0.79*** (0.05)	0.83*** (0.05)
Belief in 1st period	0.74*** (0.03)					0.57*** (0.05)	
Stock price in 1st period		0.74*** (0.03)					
Value of cued 1st period signals				0.92*** (0.03)	0.90*** (0.03)		
Value of non-cued 1st period signals				0.58*** (0.05)	0.58*** (0.05)		
2nd period signal $\times$ # 1st period signals in same context						0.36*** (0.05)	0.31*** (0.05)
Session FE	No	No	Yes	No	Yes	No	Yes
1st period signal history FE	No	No	Yes	No	No	No	Yes
Company FE	No	No	Yes	No	Yes	No	Yes
Order FE	No	No	Yes	No	Yes	No	Yes
Subject FE	No	No	Yes	No	Yes	No	Yes
Observations	800	800	800	800	800	800	800
Adjusted R^2	0.76	0.77	0.77	0.78	0.78	0.78	0.78

Notes. OLS estimates, robust standard errors (in parentheses) are clustered at the subject level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 10: Treatment *Main*: Heterogeneity analysis

	<i>Dependent variable:</i> 2nd period belief					
	(1)	(2)	(3)	(4)	(5)	(6)
2nd period signal	1.26*** (0.08)	1.25*** (0.09)	1.23*** (0.04)	1.21*** (0.04)	1.14*** (0.04)	1.14*** (0.04)
2nd period signal × Raven score	-0.030** (0.01)	-0.028* (0.01)				
2nd period signal × Memory for non-cued signals			-0.23*** (0.06)	-0.20*** (0.06)		
2nd period signal × Response time					-0.37 (0.29)	-0.35 (0.24)
Belief in 1st period	0.75*** (0.03)		0.75*** (0.03)		0.75*** (0.03)	
Session FE	No	Yes	No	Yes	No	Yes
1st period signal history FE	No	Yes	No	Yes	No	Yes
Company FE	No	Yes	No	Yes	No	Yes
Order FE	No	Yes	No	Yes	No	Yes
Observations	800	800	800	800	800	800
Adjusted R^2	0.80	0.80	0.80	0.80	0.80	0.80

Notes. OLS estimates, robust standard errors (in parentheses) are clustered at the subject level. The table suppresses the coefficients of Raven score (columns (1)–(2)), memory for non-cued signals (columns (3) –(4)), and response time (columns (5)–(6)). Response times are measured in minutes. The sample includes all observations from treatment *Main* where subjects observed a second-period signal. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 11: Treatment *Main*: Heterogeneity analysis (recall data)

	<i>Dependent variable:</i> Δ Recall [Pos. – Neg.]					
	(1)	(2)	(3)	(4)	(5)	(6)
2nd period signal	1.29*** (0.13)	1.27*** (0.13)	1.21*** (0.06)	1.19*** (0.05)	1.15*** (0.05)	1.15*** (0.05)
2nd period signal $\times$ Raven score	-0.045* (0.02)	-0.042* (0.02)				
2nd period signal $\times$ Memory for non-cued signals			-0.30*** (0.08)	-0.26*** (0.08)		
2nd period signal $\times$ Response time recall					-0.45** (0.19)	-0.50*** (0.19)
Belief in 1st period	0.74*** (0.03)		0.74*** (0.03)		0.74*** (0.03)	
Session FE	No	Yes	No	Yes	No	Yes
1st period signal history FE	No	Yes	No	Yes	No	Yes
Company FE	No	Yes	No	Yes	No	Yes
Order FE	No	Yes	No	Yes	No	Yes
Observations	800	800	800	800	800	800
Adjusted R^2	0.76	0.77	0.76	0.77	0.76	0.77

Notes. OLS estimates, robust standard errors (in parentheses) are clustered at the subject level. The table suppresses the coefficients of Raven score (columns (1) and (4)), memory for non-cued signals (columns (2) and (5)), and response time (columns (3) and (6)). Response times are measured in minutes. The sample includes all observations from treatment *Main* where subjects observed a second-period signal. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 12: Treatments *Underreaction* and *Reminder underreaction*: Recall data

	<i>Dependent variable:</i> 2nd period belief					
	Treatments:					
	<i>Underreaction</i>				+ <i>Reminder</i>	
	(1)	(2)	(3)	(4)	(5)	(6)
2nd period signal	0.62*** (0.05)	0.60*** (0.05)	0.86*** (0.05)	0.81*** (0.06)	0.89*** (0.03)	0.91*** (0.04)
Belief in 1st period	0.66*** (0.04)		0.49*** (0.05)		0.76*** (0.03)	
2nd period signal × # 1st period signals in same context			-0.33*** (0.06)	-0.28*** (0.06)		
2nd period signal × 1 if <i>Underreaction</i> , 0 if <i>Reminder</i>					-0.28*** (0.06)	-0.32*** (0.06)
Treatment FE	No	No	No	No	Yes	Yes
Session FE	No	Yes	No	Yes	No	Yes
Signal history FE	No	Yes	No	Yes	No	Yes
Company FE	No	Yes	No	Yes	No	Yes
Order FE	No	Yes	No	Yes	No	Yes
Subject FE	No	Yes	No	Yes	No	Yes
Observations	800	800	800	800	1300	1300
Adjusted R^2	0.59	0.61	0.61	0.63	0.71	0.72

Notes. OLS estimates, robust standard errors (in parentheses) are clustered at the subject level. The table suppresses the coefficients of the number of identical first-period signals (columns (3)–(4)) and of a binary indicator for whether the signals are mixed (columns (5)–(6)). The sample includes all observations from treatments *Underreaction* and *Reminder underreaction* where subjects observed a second-period signal. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 13: Treatments *Time lag* and *Time lag reminder*

	<i>Dependent variable:</i> 2nd period belief					
	Treatments:					
	<i>Time lag</i>				<i>+ Reminder time lag</i>	
	(1)	(2)	(3)	(4)	(5)	(6)
2nd period signal	1.17*** (0.04)	1.16*** (0.04)	0.85*** (0.05)	0.84*** (0.06)	1.02*** (0.02)	1.03*** (0.02)
Belief in 1st period	0.52*** (0.04)		0.33*** (0.05)		0.71*** (0.03)	
2nd period signal × # 1st period signals in same context			0.44*** (0.05)	0.45*** (0.07)		
2nd period signal × 1 if <i>Time lag</i> , 0 if <i>Reminder t. l.</i>					0.18*** (0.04)	0.15*** (0.05)
Treatment FE	No	No	No	No	Yes	Yes
Session FE	No	Yes	No	Yes	No	Yes
Signal history FE	No	Yes	No	Yes	No	Yes
Company FE	No	Yes	No	Yes	No	Yes
Order FE	No	Yes	No	Yes	No	Yes
Subject FE	No	Yes	No	Yes	No	Yes
Observations	800	800	800	800	1300	1300
Adjusted R^2	0.68	0.67	0.70	0.70	0.77	0.77

Notes. OLS estimates, robust standard errors (in parentheses) are clustered at the subject level. The table suppresses the coefficients of the number of identical first-period signals (columns (3)–(4)) and of a binary indicator for whether the signals are mixed (columns (5)–(6)). The sample includes all observations from treatments *Time lag* and *Time lag reminder* where subjects observed a second-period signal. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 14: Treatments *Time lag* and *Reminder time lag*: Recall data

	<i>Dependent variable:</i> Δ Recall [Pos. – Neg.]					
	Treatments:					
	<i>Time lag</i>				<i>+ Reminder time lag</i>	
	(1)	(2)	(3)	(4)	(5)	(6)
2nd period signal	0.97*** (0.05)	0.96*** (0.06)	0.66*** (0.05)	0.66*** (0.06)	0.97*** (0.03)	0.98*** (0.03)
Belief in 1st period	0.52*** (0.04)		0.33*** (0.05)		0.70*** (0.03)	
2nd period signal $\times$ # 1st period signals in same context			0.43*** (0.06)	0.43*** (0.07)		
2nd period signal $\times$ 1 if <i>Time lag</i> , 0 if <i>Reminder t. l.</i>					0.022 (0.06)	-0.0011 (0.06)
Treatment FE	No	No	No	No	Yes	Yes
Session FE	No	Yes	No	Yes	No	Yes
Signal history FE	No	Yes	No	Yes	No	Yes
Company FE	No	Yes	No	Yes	No	Yes
Order FE	No	Yes	No	Yes	No	Yes
Subject FE	No	Yes	No	Yes	No	Yes
Observations	800	800	800	800	1300	1300
Adjusted R^2	0.57	0.57	0.60	0.60	0.70	0.70

Notes. OLS estimates, robust standard errors (in parentheses) are clustered at the subject level. The table suppresses the coefficients of the number of identical first-period signals (columns (3)–(4)) and of a binary indicator for whether the signals are mixed (columns (5)–(6)). The sample includes all observations from treatments *Time lag* and *Reminder time lag* where subjects observed a second-period signal. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 15: Treatments *WTP* and *WTP reminder*

	<i>Dependent variable:</i> 2nd period willingness-to-pay					
	Treatments:					
	<i>WTP</i>				+ <i>Reminder WTP</i>	
	(1)	(2)	(3)	(4)	(5)	(6)
2nd period signal	1.11*** (0.04)	1.13*** (0.04)	0.93*** (0.07)	0.98*** (0.07)	0.97*** (0.04)	0.96*** (0.05)
WTP in 1st period	0.51*** (0.04)		0.46*** (0.05)		0.66*** (0.04)	
2nd period signal $\times$ # 1st period signals in same context			0.25*** (0.07)	0.20*** (0.07)		
2nd period signal $\times$ 1 if <i>WTP</i> , 0 if <i>Reminder WTP</i>					0.14** (0.06)	0.17** (0.07)
Treatment FE	No	No	No	No	Yes	Yes
Session FE	No	Yes	No	Yes	No	Yes
Signal history FE	No	Yes	No	Yes	No	Yes
Company FE	No	Yes	No	Yes	No	Yes
Order FE	No	Yes	No	Yes	No	Yes
Subject FE	No	Yes	No	Yes	No	Yes
Observations	1000	1000	1000	1000	1800	1800
Adjusted R^2	0.61	0.72	0.62	0.73	0.69	0.74

Notes. OLS estimates, robust standard errors (in parentheses) are clustered at the subject level. The table suppresses the coefficients of the number of identical first-period signals (columns (3)–(4)) and of a binary indicator for whether the signals are mixed (columns (5)–(6)). The sample includes all observations from treatments *WTP* and *WTP reminder* where subjects observed a second-period signal. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

D Experimental Instructions

D.1 Treatment *Main*

Welcome to the Experiment!

We ask you to remain quiet throughout the experiment, and to refrain from talking to or disturbing other participants. Should you have any questions, please notify one of the experimenters. Please do so quietly in order to avoid disturbing other participants.

As is the case in all experiments in the BonnEconLab, you are free to leave the experiment at any time without explanation.

The main part of the experiment consists of two parts that belong together. Below you will receive the instructions for both parts. Please read the instructions carefully. At the end of the instructions, you will be asked a series of control questions in order to test your understanding of the instructions. You may only take part in the experiment if you answer these control questions correctly.

For your participation you will receive a participation fee of 5 euros. Depending on your decisions, you can earn additional money.

PART 1 OF THE EXPERIMENT

In this experiment, there are twelve hypothetical firms. We have invented twelve firms that are in no way related to real firms. These firms have the following names:

- Firm X
- Firm I
- Firm K
- Firm N
- Firm J
- Firm M
- Firm D
- Firm U
- Firm P

- Firm G
- Firm R
- Firm T

Each firm has a stock price that is determined by a simple formula: The stock price is given by the so-called base price plus the sum of all news you receive about the respective firm over the course of the entire experiment.

For example, suppose that there are two pieces of news for a firm. Then, the stock price of that firm is calculated as follows:

$$\text{Stock price} = \text{Base price} + \text{News 1} + \text{News 2}$$

This is just an example. In the actual experiment, you will not receive two pieces of news for each firm. Instead, the number of news varies from firm to firm. You will thus receive more news about some firms than about others. It is also possible that you receive no news at all for some firms. It is just important for you to understand that the stock price is calculated as the sum of the base price and all news. In this experiment, you can hence simply calculate the stock price of a firm by adding up the base price of a firm and all news about this firm. Other factors do not play a role in determining the stock price.

The Base Price

The base prices of the firms are known and identical across firms: the base price of each firm is 100.

The News

In this experiment, there are two types of news for each firm, where one type of news is positive and the other type of news is negative. Positive news have a value of +10, which means that the stock price of the respective firm increases by 10. Negative news have a value of −10, which means that the stock price of the respective firm decreases by 10. You can see that positive and negative news each have exactly the same value, except that one is positive and one negative.

Once the experiment begins, you will see the news for the different firms in sequential order. First, on a separate screen, you will be informed about which firm the upcoming news concern. In case you receive no news for that firm, you will be informed about this on your screen. In case you do receive news, these will be displayed one after another on

your screen (one piece of news per screen). How many news you receive for a particular firm is determined randomly by the computer and does not depend on the value of the news for the firm.

The computer determines randomly whether the news for a particular firm are positive or negative. You can think of this as the computer tossing a fair coin each time:

- Heads means positive news (Probability 50%)
- Tails means negative news (Probability 50%)

Importantly, it can happen that the same type of news occurs several times. In this case, you also have to incorporate the news several times.

Example 1: The base price of a firm is 100 and you receive news -10 twice for this firm, and news $+10$ once (because the three coin tosses of the computer turned out that way). Then, the correct stock price is given by $100 - 10 - 10 + 10 = 90$.

Example 2: The base price of a firm is 100 and you receive no news about this firm. Then, the correct stock price is 100.

Example 3: The base price of a firm is 100 and you receive one news $+10$ for the firm (because the coin toss of the computer happened to land that way). The correct stock price in this case is given by $100 + 10 = 110$.

Please note that the computer independently tosses a coin for each firm and each piece of news, such that each coin toss is completely independent of the others. This means that the development of the stock price of a firm is completely random and does not follow systematic trends. Just because the first piece of news was positive does not mean that the second piece will also be positive. Rather, the probability for positive news is again exactly 50%, because the coin tosses are completely independent of each other.

Please also note that this implies that for every firm the expected value of the news is exactly zero: positive and negative news have the same value and the probability for each is 50:50. Thus, in case you don't know the news for a firm, you know that the news is on average zero and thus no change in the stock price occurs.

Communication of the News

As already mentioned, in this experiment you will receive news about the stock prices of twelve firms. In case you receive news for a firm, the news will appear sequentially on separate screens. However, the news appear separately for each firm. This means that

you will first observe all news for one company, then all the news for another company, and so on. It will be important for you to distinguish which news belong to which firm.

The news will be communicated to you on your screen. Each piece of news is communicated along with two features:

1. Each type of news is accompanied by a particular “story”, that explains to you why this particular type of news occurred.
2. Each type of news is accompanied by a particular image that will be displayed on your screen. This image will roughly reflect the story.

In this experiment, there are 24 types of news in total: one positive and one negative type of news for each of the twelve firms. As explained above, each of these 24 types of news is accompanied by a specific image and a specific story:

- The positive news about firm X will only be communicated with story 1 and image 1.
- The negative news about firm X will only be communicated with story 2 and image 2.
- The positive news about firm I will only be communicated with story 3 and image 3.
- The negative news about firm I will only be communicated with story 4 and image 4.
- Etc.

Please note: As mentioned above, it can happen that you receive the same news several times. For example, it can happen that you receive the positive news +10 twice for a given firm. The two pieces of news would then be accompanied by exactly the same story and image. When you determine a company’s stock price, you would then have to take both of these positive news into account.

Importantly, please note that it can never happen that a story accompanies different types of news, or even belongs to different firms. Each story only belongs to one type of news for one particular firm. Likewise, it can never happen that an image is associated with different types of news. Each image and each story are assigned to only one type of news for one particular firm.

If you now enter the code “1108” on your screen, you will see an example of a piece of news. Please note that the accompanying story and image are only an example and do not correspond to those in the actual experiment.

Your Task: Determine the Stock Prices of the Twelve Firms

After you will have seen the news for a firm, you will be asked to provide an estimate of the stock price of that firm. In doing so, you can earn 10 Euros. The closer your estimate is to the actual stock price of the firm, the higher the probability that you actually receive the 10 Euros. This is determined using the following formula:

$$\text{Probability of winning 10 Euros (in percent)} = 100 - (\text{Estimate} - \text{True price})^2$$

This means that the difference between your estimate and the true value is squared. This number is then subtracted from the maximum possible probability of 100%. While this formula might seem complicated, the underlying principle is very simple: the smaller the difference between your estimate and the true value, the higher the likelihood that you win 10 Euros. Notice that the probability of winning only depends on the absolute difference. Thus, it doesn't matter for your payment whether you overestimate the true value by, say, 5 or underestimate it by 5.

PART 2 OF THE EXPERIMENT

As explained above, in the first part of the experiment your task is to provide an estimate of the stock price of each firm. In the second part, we will again ask you to estimate the stock price of each firm.

You will receive up to one additional piece of news for each company. For some companies, there will be no further news. Whether or not you receive an additional piece of news for a particular company is randomly determined by the computer and does not depend on the value of the previous news for this company.

Afterwards, you will again be asked provide an estimate of the stock price of that firms.

As in the first part of the experiment, the stock price is determined by the base price (100) plus the sum of all news for the firm. Please note that the stock price of a firm is determined by all news that you have seen for that company over the course of the entire experiment, i.e., all news from the first and all news from the second part.

As in the first part, the closer your estimate is to the actual stock price of the firm, the higher your probability of winning 10 Euros. This is determined by the same formula as in the first part of the experiment:

Probability of winning 10 Euros (in percent) = $100 - (\text{Estimate} - \text{True Price})^2$

PROCEDURE OF THE EXPERIMENT

1. You will first answer a set of control questions on the computer.
2. You complete the first part of the experiment:
 - We will first inform you about which of the twelve hypothetical firms is next.
 - You will sequentially receive pieces of news for this firm. In case you receive no news for a firm, you will be informed about this on your screen.
 - Afterwards, you will be asked to enter an estimate of the stock price of this firm.
 - We will repeat this procedure for each of the twelve firms.
3. You complete several other tasks.
4. You complete the second part of the experiment:
 - We will first inform you about which of the twelve hypothetical firms you are dealing with.
 - Then, you will potentially receive an additional piece of news for this firm. For some firms, you will receive no further news.
 - Afterwards, you will be asked to enter an estimate of the stock price of this firm. The actual stock price of the firm is given by the base price plus all news that you received over the course of the experiment (i.e., part 1 and part 2).
 - We will repeat this procedure for each of the twelve firms.

YOUR PAYMENT

In addition to the 5 Euro participation fee, you can earn money with your estimates as described above. In total, you will provide 24 estimates in this experiment: two for each of the 12 firms. At the end of the experiment, the computer randomly selects one of the twelve firms as well as one of your two estimates for this firm. The probability that the estimate from part 2 of the experiment gets selected is 90% and the probability that the estimate from part 1 gets selected is 10%. You will then be paid according to your earnings from your estimate for this firm. Thus, every decision is potentially relevant for your payments. However, only one decision will actually be paid out, so there is no point in strategizing by, for example, alternating between high and low answers. In order to

maximize your earnings, you should always enter the best estimate that you have in mind for the task at hand.

As soon as all participants have read the instructions, we will provide you with another code to start the control questions.

D.2 Treatment *Reminder*

Instructions for treatment Reminder were identical to treatment Main, except that we informed subjects in the instructions for part 2 that they would be reminded of their part 1 estimates in part 2. For completeness, we display the relevant parts below.

PART 2 OF THE EXPERIMENT

As explained above, in the first part of the experiment your task is to provide an estimate of the stock price of each firm. In the second part, we will again ask you to estimate the stock price of each firm.

For each company, we will first remind you of your estimate of the stock price for this company from part 1.

You will receive up to one additional piece of news for each company. For some companies, there will be no further news. Whether or not you receive an additional piece of news for a particular company is randomly determined by the computer and does not depend on the value of the previous news for this company.

Afterwards, you will again be asked provide an estimate of the stock price of that firms.

As in the first part of the experiment, the stock price is determined by the base price (100) plus the sum of all news for the firm. Please note that the stock price of a firm is determined by all news that you have seen for that company over the course of the entire experiment, i.e., all news from the first and all news from the second part.

As in the first part, the closer your estimate is to the actual stock price of the firm, the higher your probability of winning 10 Euros. This is determined by the same formula as in the first part of the experiment:

$$\text{Probability of winning 10 Euros (in percent)} = 100 - (\text{Estimate} - \text{True Price})^2$$

PROCEDURE OF THE EXPERIMENTS

1. You will first answer a set of control questions on the computer.
2. You complete the first part of the experiment:
 - We will first inform you about which of the twelve hypothetical firms is next.
 - You will sequentially receive pieces of news for this firm. In case you receive no news for a firm, you will be informed about this on your screen.
 - Afterwards, you will be asked to enter an estimate of the stock price of this firm.
 - We will repeat this procedure for each of the twelve firms.
3. You complete several other tasks.
4. You complete the second part of the experiment:
 - We will first inform you about which of the twelve hypothetical firms you are dealing with.
 - We will then remind you of your part 1 estimate of the stock price for this company.
 - Then, you will potentially receive an additional piece of news for this firm. For some firms, you will receive no further news.
 - Afterwards, you will be asked to enter an estimate of the stock price of this firm. The actual stock price of the firm is given by the base price plus all news that you received over the course of the experiment (i.e., part 1 and part 2).
 - We will repeat this procedure for each of the twelve firms.

D.3 Treatment NoCue

Instructions for treatment NoCue were again identical to treatment Main, except for the description of news and stories. For completeness, we display the relevant parts below.

Communication of the News

As already mentioned, in this experiment you will receive news about the stock prices of twelve firms. In case you receive news for a firm, the news will appear sequentially on separate screens. However, the news appear separately for each firm. This means that you will first observe all news for one company, then all the news for another company, and so on. It will be important for you to distinguish which news belong to which firm.

The news will be communicated to you on your screen. Each piece of news is communicated along with two features:

1. Each type of news is accompanied by a particular “story”, that explains to you why this particular type of news occurred.
2. Each type of news is accompanied by a particular image that will be displayed on your screen. This image will roughly reflect the story.

Every single piece of news is attached to its own image and its own story.

- The first piece of news for company X (should you receive one) will be communicated with an separate story and a separate image.
- The second piece of news for company X (should you receive one) will be communicated with an separate story and a separate image.
- The first piece of news for company I (should you receive one) will be communicated with a separate story and a separate image.
- The second piece of news for company I (should you receive one) will be communicated with a separate story and a separate image.
- Etc.

Please note: As mentioned above, it can happen that you receive the same news several times. For example, it can happen that you receive the positive news +10 twice for a given firm. The two pieces of news would then be accompanied by exactly two different stories and two different images. When you determine a company’s stock price, you would then have to take both of these positive news into account.

Importantly, please note that it can never happen that a story accompanies multiple news, or even belongs to different firms. Each story only belongs to one piece of news for one particular firm. Likewise, it can never happen that an image is associated with multiple news. Each image and each story are assigned to only one piece of news for one particular firm.

If you now enter the code “1108” on your screen, you will see an example of a piece of news. Please note that the accompanying story and image are only an example and do not correspond to those in the actual experiment.

E Experimental Materials

E.1 Screenshots of signal, story and image presentation



Company N tries to advertise its products through commercials with German celebrities, like, for instance, Boris Becker, Helene Fischer or Til Schweiger. Recently, a new advertisement campaign with a celebrity worked extremely well.

The value of the company increased by 10 points.

Figure 12: Example screenshot of how a piece of positive news for Company N is communicated to subjects. The signal is displayed in the last line of the text. A story and an image accompany the signal.



The head of sales of Company K is a choleric. Every once in a while, he engages in temper tantrums during which he yells at customers of Company K and insults them. These customers hence take their business elsewhere. Just now, another temper tantrum occurred.

The value of the company decreased by 10 points.

Figure 13: Example screenshot of how a piece of negative news for Company K is communicated to subjects. The signal is displayed in the last line of the text. A story and an image accompany the signal.