Associative Memory and Belief Formation

Benjamin Enke*
Frederik Schwerter**
Florian Zimmermann***

January 2020

*Harvard University and NBER; enke@fas.harvard.edu
** University of Cologne; frederik.schwerter@uni-koeln.de
***University of Bonn and briq; florian.zimmermann@briqinstitute.org

Associative Memory and Belief Formation[*]

Benjamin Enke   Frederik Schwerter   Florian Zimmermann

January 16, 2020

Abstract

Information is often embedded in memorable contexts, which may cue the asymmetric recall of similar past news through associative memory. We design a theory-driven experiment, in which participants observe signals about hypothetical companies. Here, identical signal realizations are communicated with identical contexts: stories and images. Because participants asymmetrically remember those past signals that get cued by the current context, beliefs systematically overreact. This overreaction depends in predictable ways on the signal history; the correlation between signals and contexts; and the scope for forgetting and associative memory. We quantify these results by structurally estimating a model of associative recall.

1 Introduction

This paper experimentally studies the role of associative memory for the formation of beliefs. In textbook models of belief formation, memory imperfections play no role: agents combine prior knowledge with current information, and yesterday's belief equals today's prior. Our paper starts from the premise that people do not constantly have access to their 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 be 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, by which we mean intrinsically uninformative environmental features that accompany information, such as stories and narratives, images, emotions, or sounds. Oftentimes, *similar news are embedded in similar contexts*. For example, when individuals receive negative feedback about their performance, these negative news are often associated with scolding and public shaming. Similarly, when good news prevail in the stock market, people are disproportionately exposed to bulls, upward-sloping trend lines, and good-times stories. 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.

The observations (i) that people may need to reconstruct prior information from memory and (ii) that similar news are often embedded in similar memorable contexts motivate the question about the role of *associative recall* for belief formation. Associative recall refers to the idea that people are more likely to recollect items that are cued by current items (here: the current context). The associative nature of memory has recently received increased attention in the theory literature (Mullainathan, 2002; Bordalo et al., 2019b). 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, beliefs might *look like* they overreact to recent news, purely as a result of how prior knowledge is reconstructed. Yet even though overreaction in beliefs about macroeconomic variables and geopolitical events is the object of an active line of research (Augenblick and Rabin, 2018; Bordalo et al., 2018) and has been argued to reflect the effects of memorable narratives (Shiller, 2019), direct empirical evidence on a potential link between overre-

action and memory is scarce.

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. (2019b) and Mullainathan (2002). In this model, decision-makers (i) have imperfect memory; (ii) are more likely 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 at least partially naïve about their memory imperfections. This stylized model predicts overreaction in beliefs. Importantly, in the model, this overreaction does not occur because people incorporate the current signal in some suboptimal way, but only because they asymmetrically retrieve past signals. The model makes predictions about how such overreaction depends on the signal history, the correlation structure between signals and contexts, the imperfection of memory, or the relevance of associative recall. Our experimental treatments are tightly 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, model-based, and financially incentivized designs that dominate modern experimental economics research on bounded rationality and (ii) psychological paradigms on cued recall problems. In our experiment, participants predict the stock market value of multiple hypothetical companies, an intuitive exercise for participants. The experiment comprises two distinct periods that we think of as "past" and "present." Across both periods, a subject sequentially observes 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 an (intrinsically uninformative) 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, as in the motivating examples, identical news are embedded in identical contexts: there is a one-to-one mapping between {Company × 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 news for a company and then states a first belief about the value of that company. This process is repeated for all companies. Using the data on first-period beliefs,

we 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, subjects work on an unrelated real effort task for 15 minutes to activate long-term memory. 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 in the second period might make it more likely for subjects to (asymmetrically) remember first-period news that were communicated in the same context.

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 model is that the OLS coefficient in a regression of second-period beliefs on second-period signals equals one. Likewise, a version of our model with imperfect but no associative memory also predicts a regression coefficient of one. In contrast, our framework predicts that, if context and news are positively correlated, (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 manipulated to be perfect; and (iv) overreaction disappears if associative recall is exogenously shut down. All of these predictions hold when context and news are always linked in the same way. In contrast, when signals suddenly appear in a context that was previously associated with the opposite type of signal, our framework predicts that beliefs under- rather than overreact (prediction (v)). Our experiments with a total of 830 lab subjects were pre-registered to test these predictions, including a pre-analysis plan.

We test prediction (i) using the baseline treatment variation *Main* discussed above. We find that second-period beliefs strongly overreact with respect to the second-period signal: the aggregate OLS regression coefficient is 1.10, substantially larger than its rational or imperfect-but-no-associative-recall benchmark of one.

In a follow-up treatment, we document that overreaction in beliefs extends to economic choices. We implement the same setup as in treatment *Main*, except that we do not elicit participants' beliefs about the value of the hypothetical companies but instead participants' willingness-to-pay. Here, reported willingness-to-pay also strongly overreacts with respect to the second-period signal, with an aggregate OLS coefficient of 1.11.

3

Having established the baseline phenomenon of overreaction, we next turn to exogenous variation in our model parameters and hence direct causal evidence for the role of memory in overreaction. As a first piece of causal evidence, we study the effect of the signal history (prediction (ii)). As predicted by our model, the magnitude of overreaction is strongly increasing in the number of first-period signals that get cued by the second-period signal. For instance, when subjects do not observe any first-period signals that match the second-period signal, their beliefs do not overreact at all. We verify that subjects' direct recall data further support our findings on participant beliefs: subjects are much better at recalling the frequency of those first-period signals that got communicated in the same context as the second-period signal. These patterns cannot be explained by recency effects and directly point to a role of associative recall for overreaction.

To provide further 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 eliminates the imperfection of memory, so that asymmetric recall and hence overreaction can no longer take place. We find that subjects' beliefs indeed do not overreact in treatment *Reminder*.

Having documented the role of imperfect memory for overreaction, we next directly manipulate the relevance of associative recall. In our model, associativeness operates via identical contexts. Thus, in order to show that it is indeed associative memory that generates overreaction in our experiments, treatment *No Cue* 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. As predicted by the model, overreaction disappears entirely in *No Cue*, and the treatment difference in overreaction between *Main* and *No Cue* is quantitatively large and statistically significant. Again, all of these results hold not only when we consider participants' beliefs 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: 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 the context that was previously associated with positive signals for that same company. In this treatment, our formal framework predicts that beliefs should underreact to the last signal because it selectively cues the retrieval of signals that took on a different value than the current signal. We find that beliefs in *Underreaction* systematically underreact, also relative to a control treatment *Underreaction reminder* that reminds participants of their first-period beliefs.

Because this paper focuses on the role of associative recall *conditional* on the existence of forgetting in the first place, our model and baseline treatments are agnostic about whether forgetting arises due to the time lag between the first and second period, due to general memory load, or due to what memory researchers call "interference:" memory imperfections that arise from the similarity of news (+10 and -10) across the different companies. Still, to shed some light on this issue, we also implement variations of our experiment in which we vary the length of the time lag and the presence of interference. While we discuss the details of these treatments in Section 6, the results suggest an important role for interference in shaping memory constraints.

All of our main results are derived from theoretically-motivated reduced-form regressions. In complementary analyses in the final part of the paper, we structurally estimate our stylized model, in particular the parameters that govern the imperfection of memory and the strength of associative recall. The results of our estimations show 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' stated recall.

In summary, the central contributions of our paper are threefold: (i) the first experimental economics investigation of the role of associative memory for belief formation and choice; (ii) in a novel model-based decision environment that builds a bridge between quantitative economic decision tasks and psychological cued recall problems; and (iii) the first structural quantification of the roles of imperfect and associative recall for belief formation. Our paper hence fits into an emerging literature that has argued for the importance of associative memory for economics. Mullainathan (2002) and Bordalo et al. (2019b) present models of how cued recall shapes economic decision-making across a broad set of domains. Models on cuing effects in consumption include Laibson (2001) and Bernheim and Rangel (2004). Related theoretical work has investigated the implications of associative recall in applied settings such as updating biases (Gennaioli

and Shleifer, 2010), financial markets (Bodoh-Creed, 2013; Bordalo et al., 2017, 2018; Wachter and Kahana, 2019), and self-esteem (Koszegi et al., 2019).[1] Thus far, this literature is theoretical in nature. 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 economic decision making tasks in relation to this emerging body of theoretical work.

Our work builds on a long psychology literature on episodic memory, which is the part of long-term memory that stores information about events and experiences. Associative recall is usually thought of as a component of episodic memory. Psychological experiments on associative recall exhibit a different structure than the experiments that are presented here (see Kahana, 2012, for an overview). These typically consist of explicit cued recall problems (such as with words), rather than model-driven quantitative economic decision tasks. Also, psychological experiments do not focus on the implications of associative recall for beliefs or willingness-to-pay, as we do here. An important concept in psychological research, which we also leverage in our experimental design, is that of item similarity (Tversky, 1977). For example, Bordalo et al. (2019a) present an experiment on selective recall of abstract images that shows a link between associative memory and the representativeness heuristic.

Finally, our work also relates to work that documents overreaction or excess movement in survey expectations (Augenblick and Rabin, 2018; Bordalo et al., 2018, 2017). We discuss the relationship between this line of work and our results in the Conclusion.

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 main results, and Section 6 evidence aimed at understanding mechanisms. Section 7 estimates the model and Section 8 concludes.

---

[1]Other research in economics on memory that does not focus on associative recall includes work on heuristics (Wilson, 2014), the role of personal experiences (Schwerter and Zimmermann, 2019; Herz and Taubinsky, 2017; Malmendier and Nagel, 2015) and motivated (self-serving) memory (Zimmermann, forthcoming; Carlson et al., 2018; Huffman et al., 2018). More broadly, our paper also relates to the recent experimental literature on bounded rationality, in particular work that has focused on the microfoundations 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; Hartzmark et al., 2019).

2 Theoretical Framework

2.1 Setup

This section presents a stylized model to guide the design of the experiments and structure the empirical analysis. The mechanics of the model directly build on some of the formulations in Mullainathan (2002) and Bordalo et al. (2019b). The framework rests on three key assumptions: (i) people may forget prior knowledge, so that they need to reconstruct it 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. In this model, decision-makers behave optimally conditional on what they recall. This means that 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 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 t, θ_t is given by its initial value plus the sum of all news n_x that have accumulated up to this point, where $n_x \in \{-q, q\}$. News are equally likely and i.i.d. We will use the terms "news" and "signal" interchangeably.

A piece of news n_x is associated with a memorable context $c_x \in \{L, H\}$. 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:

$$\theta_2 = v + \sum_{x=1}^{k} n_x + n_{k+1} \tag{1}$$

Just as in $t = 1$, the piece of news is associated with a context. We will consider two regimes, though for any given DM the prevailing regime is known. In the first regime, second-period news and contexts are associated in the same way as in the first period: positive news appear in a "high" context and negative news in a "low" context. 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 "low" context and negative news in a "high" one. As a shorthand for this "correlation" between news and context, we define

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

2.2 Memory and Beliefs

Our object of interest is the extent to which the DM's belief about θ_2 in $t = 2$ responds to the latest piece of news n_{k+1}. A rational (or Bayesian, though there is no uncertainty here) DM would 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 $t = 2$. Thus, his belief (after observing n_{k+1}) is given by

$$b_2 = v + \sum_{x=1}^{k} m_x n_x + n_{k+1}, \tag{2}$$

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 $r \in [0, 1)$ that the DM will remember. The reduced-form assumption of imperfect recall is a shorthand for different mechanisms that have been highlighted in the psychological literature. For now, we will assume that the parameter r is exogenously given, though our experimental design and results will shed some light on what induces imperfect memory in the first place.

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 past signal is more likely to get remembered if it occurred with the same context as today's signal. Formally, there is an increase in the probability of recalling $(1-r)a$, $a \in (0, 1]$, 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.

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).[2] We have:

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

Denote by $z \geq 0$ the number of news in $t = 1$ that were observed in the same context as n_{k+1} and hence got "cued." Doing straightforward algebra, the expected belief in period $t = 2$ is then given by:

---

[2]In principle, naïveté could come in two forms: (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.

$$E[b_2|n_x, n_{k+1}] = v + n_{k+1} + \sum_{x=1}^{k} m_x n_x$$

$$= v + n_{k+1} + \sum_{x=1}^{k} r n_x + \sum_{x=1}^{z} (1-r) a n_x \qquad (4)$$

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

$$= v + \underbrace{[1 + \rho z(1-r)a]}_{\text{Object of interest}} n_{k+1} + r \sum_{x=1}^{k} n_x \qquad (6)$$

Equation (6), in particular the bracketed term, is the core expression that we subject to systematic experimental tests. If the agent is rational ($r = 1$), second-period beliefs will respond with a coefficient of one to variation in the second-period signal. Similarly, if the agent exhibits imperfect ($r < 1$) but no associative ($a = 0$) recall, the bracketed term equals one. On the other hand, viewed through the lens of imperfect and associative memory, equation (6) suggests that beliefs will overreact if context and news are positively correlated ($\rho = 1$). At the same time, the equation clarifies that overreaction does not occur because the agent incorporates the last signal in some suboptimal way, but only because he asymmetrically retrieves first-period signals.[3]

In our experiments, we exogenously manipulate the components of the bracketed expression. This suggests the following abstract testable 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 or relevance of associative recall (a).*

5. *If the correlation between news and context is negative ($\rho = -1$), expectations underreact to today's news, on average.*

---

[3]Moreover, equation (5) 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. (5)).

6. *This underreaction increases in the number of past news that there were communicated in the same context as today's news.*

7. *Underreaction increases in the imperfection of memory $(1-r)$.*

It is worth highlighting that these predictions rely on the presence of associative recall $a > 0$. 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, or that it disappears once associative recall is shut down.

3   Experimental Design

Our experimental design is guided by the following design objectives: (i) a decision setup that is closely tied to the model in Section 2; (ii) a task that is very intuitive and simple, conditional on what is being recalled; (iii) to build a bridge between the tightly-controlled and quantitative designs that dominate modern experimental economics research on the one hand and psychological paradigms on cued recall problems on the other hand; (iv) exogenous variation in the key model parameters; and (v) incentive-compatible belief elicitation.

3.1   Experimental Setup

Task.   To isolate the role of memory, we implemented a simple deterministic decision environment in which, absent potential memory constraints, behaving rationally is trivial. This ensures that results are not conflated or noised up by subjects having to go through non-trivial Bayes'-rule-type calculations. The experiment consisted of two periods, as summarized in Figure 1. In both periods, participants estimate the stock market value of hypothetical companies.

First period.   Continuing the notation from Section 2, the value of company j in period $t = 1$ is given by a baseline value, $v = 100$, plus the sum of all news about that company in $t = 1$:

$$\theta_{j,1} = 100 + \sum_{x=1}^{k_j} n_x. \tag{7}$$

COMPANY$_i$ COMPANY$_{xii}$ COMPANY$_{iv}$ COMPANY$_{ix}$
news period-1 ⋯ *news period-1*   *news period-2 signal* ⋯ *news period-2 signal*
display belief   *display belief*   *display belief recall*   *display belief recall*

┤————————————┤————————————┤————————————┤————————————→ t

Period 1   Time Gap   Period 2   Question-
                                                                                naire

where $k_j \in \{0, 1, 2, 3\}$ is the number of signals in $t = 1$.[4] 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.

Subjects sequentially observed news for a particular company on their computer screens. Then, they were asked to estimate the company's current value. This procedure was repeated for all twelve companies. Thus, participants worked on the twelve companies strictly sequentially.

Beliefs in the first period allow us to verify whether subjects understood the basic information structure, had sufficient time to process the information, and were in principle able to determine correct estimates 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 understood our design and were well-capable of behaving optimally.

After the first period, we implemented a time gap in which subjects worked on an unrelated real effort task, which 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.

Second period. In the second period, for each company, subjects were shown up to one additional piece of news. The value of company j is hence given by:

$$\theta_{j,2} = \theta_{j,1} + n_{k+1} = 100 + \sum_{x=1}^{k_j} n_x + n_{k+1}. \tag{8}$$

For ten companies, subjects received an additional piece of news, while for two companies, there were no additional news. The experimental instructions and comprehension questions emphasized that first-period signals are also relevant for second-period guesses. 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

---

[4]Each subject 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.

second period.

Immediately after observing the additional piece of news for a company, subjects 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 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, so that participants worked on the 12 companies strictly sequentially within a given period.

To summarize, as depicted in Figure 1, the timeline of the experiment was as follows. First, a subject received all first-period signals for a company and immediately after stated a first-period belief. Then, the subject received all first-period signals for the next company and stated a first-period belief. This process was repeated for all twelve companies, after which a 15-minutes real effort task followed. Then, the subject received a second-period signal (if any) for a company and immediately after stated a second-period belief and indicated their recall of positive and negative signals. This procedure was then repeated for all twelve companies.

Communication of news and context. News were not only communicated as abstract numbers, but were shown on subjects' computer screens with what we refer to as a context. Neither our stylized model nor existing theoretical contributions define what exactly is part of a context. For the purposes of our experimental implementation, we use "context" as a shorthand for an image and a story that accompany a signal.

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 F contains examples of these images and stories (see Figures 15 and 16). An English version of the computer program that communicates the sequence of first-period and second-period news and contexts in treatment *Main* can be accessed at `https://unikoelnwiso.eu.qualtrics.com/jfe/form/SV_0MrVD2rNNrKeLGt`.

Randomization and incentives. 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 Discussion of Design in Relation to Psychological Concepts

Given that ours is arguably the first structured economics experiment on associative recall, we deem it helpful to explain our design choices in light of the psychology literature. As discussed in the Introduction, associative recall is thought to be part of episodic memory, which is that part of long-term memory that stores past events. Our experiment is based on the following concepts from memory research (Kahana, 2012).

First, because associative recall is believed to operate on long-term (rather than working) memory, we implemented a distraction task between the first and second period. The memory literature contains many demonstrations that sufficiently long distraction tasks activate long-term memory and corresponding memory imperfections.

Second, an important component of recent memory models in both economics and psychology is that of similarity (Tversky, 1977) and resulting interference (Kahana, 2012; Bordalo et al., 2019b). The key idea is that it is hard for people's memory to link a specific piece of past information to a particular variable if they have been exposed to similar information also for other variables. Based on this insight, we deliberately designed our experiment so as to include twelve companies with similar news.

Notice that both of these components could generate baseline forgetting. While the focus of our paper is an analysis of the role of associative memory *conditional on the existence of forgetting*, in Section 6, we shed light on the relative importance of the time lag and interference in generating forgetting in the first place.

3.3 Treatment Variations and Sources of Exogenous Variation

We conducted seven treatments, referred to as *Main, Reminder, No Cue, Underreaction, Underreaction reminder, WTP* and *WTP 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 beliefs and correspnding economic actions; (ii) the ways in which the quantitative magnitude of such overreaction causally depends on the precise signal history, in particular the parameter z; (iii) the causal roles of imperfect and associative memory for overreaction; and (iv) the role of the correlation between context and news.

Treatments *Main* and *WTP*. 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.

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 treatment *Main*.

Treatment *WTP* follows the same structure as *Main*, except that we do not elicit participants' beliefs about the value of the hypothetical companies. 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 for each company and then stated their willingness-to-pay (WTP) for a company. This treatment hence allows us to tie associative recall back to economic actions. To elicit WTP, 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. Because we anticipated that participants' WTP would be a slightly noisier measure than pure beliefs data, 100 subjects partici-

pated in treatment *WTP*.

Treatments *Reminder* and *WTP reminder*. In treatments *Reminder* and *WTP reminder*, we seek to remove subjects' memory constraints, holding everything else constant. The setup in *Reminder* was exactly the same as in *Main*, except that at the beginning of the second period (i.e., before a subject observes the second-period signal for a company), subjects were reminded of their own first-period belief for that company. Similarly, the setup in *WTP reminder* is the same as in *WTP*, except that subjects were reminded of their own first-period willingness-to-pay before they received a second-period signal. Thus, the reminder treatments assist subjects in the recall of their first-period knowledge, 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). 50 subjects participated in treatment *Reminder* and 80 subjects in treatment *WTP reminder*.

Treatment *No Cue*. Treatment *No Cue* was designed to manipulate the relevance of associative recall. The setup in this treatment was exactly the same as in *Main*, except that each piece of news was 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. Comparing treatments *Main* and *No Cue* therefore allows us to cleanly identify the role of associative recall. 80 subjects participated in this treatment.

Treatments *Underreaction* and *Underreaction reminder*. All treatments described above rely 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, 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

15

the image that were associated with positive news for company A in the first period of the experiment. The instructions in *Underreaction* emphasized that second-period news were communicated along with the opposite story and image, and control questions verified subjects' understanding of this aspect of the design. 80 subjects participated in this treatment.

To allow for causal inference, we additionally 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. This treatment hence allows us to verify whether any potential underreaction in treatment *Underreaction* indeed stems from memory imperfections, rather than from confusion because the mapping between news and contexts flips after the first period. 50 subjects took part in *Underreaction reminder*. In all other respects, these two treatments followed the same procedure as treatments *Main* and *Reminder*.

3.4   Predictions

Equation (6) in the conceptual framework directly suggests the following estimating equation for subject i's second-period belief about the value of company j:

$$b_2^{i,j} = \alpha + \beta_1 n_{k+1}^j + \beta_2 \sum_{x=1}^{k} n_x^j + \epsilon^{i,j} \tag{9}$$

That is, we regress a subject's second-period belief on the value of the second-period signal as well as the first-period stock value (or the first-period belief). In those treatments in which we elicited WTPs rather than beliefs, $b_2^{i,j}$ refers to the second-period WTP. 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 treatments Main and WTP, there is overreaction: $\hat{\beta}_1 > 1$.*

2. *In treatments Main and WTP, 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, and stronger in WTP than in WTP reminder.*

4. *Overreaction is stronger in treatment Main than in No Cue.*

Table 1: Mapping from model predictions to experimental predictions

Abstract model prediction	Treatments	Experimental prediction
1. Overreaction if news and context positively corr.	*Main, WTP*	$\hat{\beta}_1^{Main} > 1$
2. Overreaction increases in # identical past contexts	*Main, WTP*	$\hat{\beta}_1^{Main}$ increases in z
3. Overreaction increases in imperfection of memory	*Main, WTP, Reminder*	$\hat{\beta}_1^{Main} > \hat{\beta}_1^{Reminder}$
4. Overreaction increases in relevance of associative recall	*Main, No Cue*	$\hat{\beta}_1^{Main} > \hat{\beta}_1^{No\ Cue}$
5. Underreaction if news and context negatively corr.	*Underreaction*	$\hat{\beta}_1^{Under.} < 1$
6. Underreaction increases in # identical past contexts	*Underreaction*	$\hat{\beta}_1^{Underr.}$ decreases in z
7. Underreaction increases in imperfection of memory	*Under., Under. rem.*	$\hat{\beta}_1^{Under.} < \hat{\beta}_1^{Under.\ rem.}$

5. *In treatment Underreaction, we observe underreaction:* $\hat{\beta}_1 < 1$.

6. *In treatment Underreaction, underreaction increases in the number of first-period signals that were observed in the same context as the second-period signal.*

7. *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.

3.5   Procedures and Logistics

Upon arrival in the lab, subjects received written instructions about the experiment. Appendix E contains the full set of paper-based instructions, translated into English. 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).

Treatments *Main*, *Reminder*, and *No Cue* were conducted in the BonnEconLab of the University of Bonn. Since we had exhausted the subject pool of the BonnEconLab, treatments *WTP*, *WTP reminder*, *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 *No Cue* were all implemented in the same sessions, as were *Underreaction* and *Underreaction reminder*. In our statistical analyses, we only compare treatments that were

randomized within experimental sessions, in the same location. The experiments were computerized using Qualtrics and lasted up to 90 minutes.

3.6 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 (9) through which we analyze all data; (iv) all predictions outlined in Section 3.4; (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 multiple steps. We first pre-registered and implemented treatments *Main, Reminder, No Cue,* as well as treatments *Extended time lag* and *Extended time lag reminder* (to be discussed in Section 6). Based on results from these treatments, we pre-registered treatments *Underreaction, Underreaction reminder, WTP, WTP reminder,* and additional treatments discussed in Section 6. Table 7 in Appendix B provides an overview of all treatments that we conducted for this paper, including information on subjects' average earnings and pre-registration details. 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 participants' first-period beliefs. Table 8 in Appendix D 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 show 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 a subject did not receive a second-period signal. Table 9 in Appendix D reports the results. In a regression of second-period

on first-period beliefs, the OLS coefficient is only 0.56 and hence far 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.

4.2   Treatments *Main* and *WTP*: Overreaction in Beliefs and Choices

Throughout the empirical analysis, we present OLS regressions that correspond to variants of the estimating equation (9). Columns (1)–(4) of 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 second-period 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.

The framework outlined in Section 2 predicts that the coefficient of first-period beliefs or first-period stock values is less than one (due to imperfect memory) and that the coefficient of the second-period 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 beliefs substantially overreact with respect to second-period signals, by 10–11 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.[6]

As highlighted by equation (5) in the formal framework, our hypothesis is that overreaction occurs because the first-period signals get recollected more successfully if they get 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

---

[5]Naturally, 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.

[6]Table 10 in Appendix D replicates the results of Tables 2 for the direct 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. The results using this measure are very similar to those in Table 2.

19

Table 2: Treatments *Main* and *WTP*

	Treatment *Main*				Treatment *WTP*			
	Dependent variable:							
	2nd period belief				2nd period WTP			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
2nd period signal	1.10***	1.11***	1.11***	0.87***	1.11***	1.14***	1.13***	0.98***
	(0.02)	(0.02)	(0.03)	(0.04)	(0.04)	(0.04)	(0.04)	(0.07)
Belief in 1st period	0.75***							
	(0.03)							
Company value in 1st period		0.74***				0.60***		
		(0.03)				(0.04)		
Value of cued 1st period signals				0.90***				0.70***
				(0.03)				(0.05)
Value of non-cued 1st period signals				0.59***				0.49***
				(0.05)				(0.06)
WTP in 1st period					0.51***			
					(0.04)			
Session FE	No	No	Yes	Yes	No	No	Yes	Yes
1st period signal history FE	No	No	Yes	No	No	No	Yes	No
Company FE	No	No	Yes	Yes	No	No	Yes	Yes
Order FE	No	No	Yes	Yes	No	No	Yes	Yes
Subject FE	No	No	Yes	Yes	No	No	Yes	Yes
Observations	800	800	800	800	1000	1000	1000	1000
Adjusted R^2	0.80	0.80	0.80	0.81	0.61	0.51	0.72	0.73
p-value H0: β (2nd period signal)=1	<0.01	<0.01	<0.01	n/a	<0.01	<0.01	<0.01	n/a

Notes. OLS estimates, robust standard errors (in parentheses) are clustered at the subject level. The sample includes all observations from treatments *Main* (columns (1)–(4)) and *WTP* (columns (5)–(8)) where subjects observed a second-period signal. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

the overall value of those first-period signals that do (or do not) equal the second-period signal. The results show that beliefs 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.

Columns (5)–(8) of Table 2 report analogous analyses for treatment *WTP*, where the dependent variable is now a participant's reported WTP. Note that because the decision problem in our setup is deterministic in nature, the prediction for a rational decision-maker is that the OLS coefficient of the second-period signal equals one, regardless of subjects' risk attitudes. The results are very similar to those in treatment *Main*: (i) we see overreaction with an aggregate OLS coefficient of 1.11; and (ii) the coefficient of cued first-period signals is substantially and significantly larger than the coefficient of non-cued first-period signals.

Result 1. *Beliefs and choices overreact with respect to the second-period signal.*

4.3 Heterogeneity Analysis

Next, we examine across-subject heterogeneity in overreaction. To estimate the presence of such heterogeneity, we require a subject-level measure of overreaction. To this effect, we resort back to the beliefs data from treatment *Main* and run our standard regression of second-period beliefs on the second-period signal, but now separately for each subject. Figure 12 presents the distribution of subject-level regression coefficients. Here, both the rational and the imperfect-but-no-associative-recall predictions again correspond to a coefficient of one. 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.

To investigate the correlates of this heterogeneity, we turn to three pre-registered heterogeneity analyses: (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;[7] and (iii) response times. Tables 11 and 12 in Appendix D report the results. We find that subjects with higher Raven scores and better non-cued recall exhibt less overreaction. The relationship between overreaction and response times is negative, but not statistically significant.

5 Exogenous Variation in Model Parameters

5.1 Variation in the Signal History

We proceed by considering exogenous variation in the model parameters. We begin by considering Prediction 2: that overreaction increases in the number of cued first-period signals z. 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 (6). 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.

Figure 3 visualizes the results for treatment *Main*. We show the analogous figure for treatment *WTP* in Figure 8 in Appendix C. For each set of possible signal frequencies in the first period of the experiment, we regress second-period beliefs on the 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, indicating overreac-

---

[7]For each subject, we regress the reported recall of non-cued signals on the actual number of corresponding signals and use this regression coefficient as a measure of the strength of (non-cued) memory.

Figure 2: Subject-level distribution of regression coefficients of the second-period 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. To 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 exhibit a coefficient of one.

tion. At the same time, visual inspection suggests that the coefficient is increasing in the number of first-period signals.[8] This is intuitive and predicted by the model because if there are more past signals that can be cued, then associative recall generates more pronounced overreaction.

It is reassuring that beliefs do not at all overreact in the case of zero positive and zero negative first-period signals, as predicted by our model. In fact, this coefficient is significantly 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 underreaction or shading (Benjamin, 2018). As suggested in recent work by Enke and Graeber (2019), such shading at least partly reflects a response to cognitive uncertainty: people's subjective uncertainty about what the rational belief is. Here, cognitive uncertainty plausibly arises because subjects know that they might forget some first-period signals. Either way, in our experiments, associative memory is sufficiently strong to turn such underreaction into overall overreaction.

Table 3 provides a formal and pre-registered statistical test of Prediction 2 for both treatment *Main* and *WTP*. Here, we interact the value of the second-period signal with

---

[8]A casual inspection of Figure 3 may suggest that overreaction depends on whether first-period signals are predominantly positive or negative. In a replication that we report on below (treatment *Main replication*), we do not see such suggestive patterns, see Figure 10 in Appendix C.

Figure 3: OLS coefficient (+/- 1 SE) in a regression of second-period beliefs on the last signal, separately for each set of signal frequencies in the first period. The regressions control for a subject's first-period belief. Standard error bars are computed based on clustering at the subject level.

the number of first-period signals that were communicated in the same context as the second-period signal (z in the model). The results show that the interaction term is positive and statistically highly significant. The magnitude suggests that each additional first-period signal increases the responsiveness to the second-period signal by about 20–35%, on average.

To corroborate this result, we next turn to subjects' direct recall data. Figure 4 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. This is indicative that associative memory helps not only with re-membering *whether* a certain type of signal has appeared before, but also *how often* it appeared.[9]

---

[9]Subjects indicated their recall of signals immediately after they stated a point belief. A potential concern is that the recall data do not have independent informational content but are constructed by subjects through ex post reasoning to match their stated beliefs. At first sight, the data do not seem to support such a consistency-based interpretation. To see why, note that Figure 4 reveals very accurate recall of cued signals. Such accurate recall would not be predicted by a simple consistency account because there are often multiple combinations of positive and negative signals that would rationalize a given belief. For example, if a subject stated a belief of 110, then there are two combinations of recall data that rationalize such a belief: one positive / zero negative signals, and two positive / one negative signals.

Table 3: Treatments *Main* and *WTP*: The role of the signal history

	Treatment *Main*			Treatment *WTP*		
	Dependent variable:					
	2nd period belief			2nd period WTP		
	(1)	(2)	(3)	(4)	(5)	(6)
2nd period signal	0.85***	0.87***	0.87***	0.93***	0.95***	0.98***
	(0.04)	(0.04)	(0.04)	(0.07)	(0.07)	(0.07)
2nd period signal ×	0.34***	0.32***	0.31***	0.25***	0.26***	0.20***
# 1st period signals in same context	(0.05)	(0.05)	(0.05)	(0.07)	(0.08)	(0.07)
Belief in 1st period	0.59***					
	(0.05)					
Company value in 1st period		0.59***			0.47***	
		(0.05)			(0.06)	
WTP in 1st period				0.46***		
				(0.05)		
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	800	800	800	1000	1000	1000
Adjusted R^2	0.81	0.81	0.81	0.62	0.52	0.73

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. Columns (1)–(6) suppress the coefficient of the number of first-period signals that were communicated with the same context as the second-period signal. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Result 2. *Overreaction increases in the number of cued first-period signals.*

5.2 The Role of Imperfect Memory

To provide causal evidence for the role of imperfect memory in belief overreaction, we manipulate whether participants actually need to reconstruct prior knowledge from memory. 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.

Figure 5 summarizes the results by comparing the OLS coefficient of the second-period signal across treatments *Main* and *Reminder*. There is no overreaction in treatment *Reminder*. Instead, the tightly estimated regression coefficients equal almost exactly one. These results provide direct causal evidence that imperfect memory is necessary in order for overreaction to arise in our setup, as predicted by our key equation (6).

Figure 4: 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.

These results also represent reassuring evidence that our experimental setup is not mis-construed by subjects: in the absence of memory constraints, the second-period signal is incorporated in a rational fashion.

To formally compare treatments *Main* and *Reminder*, we again resort to OLS regressions. Columns (1)–(2) of Table 4 present the results. As specified in the pe-registration, we again analyze our data by means of OLS regressions in which we relate subjects' second-period beliefs to the value of the second-period signal, except that now we also interact the second-period signal with a treatment dummy. Our prediction, spelled out in Sections 2 and 3.4, 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. In *Main*, subjects respond 12–14% more to the value of the second-period signal than subjects in *Reminder*. Again, this pattern is a specific prediction of our framework, but not of an account of recency effects.[10]

Columns (3)–(4) present analogous analyses for actions (willingness-to-pay) by comparing treatments *WTP* and *WTP reminder*. Again, overreaction is substantially stronger

---

[10]Table 13 in Appendix D shows that almost identical results hold when we again consider the summary statistic of subjects' direct recall.

Figure 5: OLS coefficient (+/- 1 SE) in a regression of second-period beliefs on the last signal, separately for each set of signal frequencies in the first period. The regressions control for a subject's first-period belief. Standard error bars are computed based on clustering at the subject level.

in the presence of memory imperfections. Indeed, the coefficient of the second-period signal suggests that in treatment *WTP reminder* there is no overreaction at all.

Result 3. *Overreaction disappears once forgetting is shut down.*

5.3   The Role of Associative Memory

We proceed by experimentally manipulating the relevance of associative memory. According to equation (6), if there is no associative recall, there should be no overreaction. As a direct test of this hypothesis, we compare treatments *Main* and *No Cue*. Recall that in treatment *No Cue*, each signal realization was communicated with a different context, so that the current context cannot cue identical past contexts.

Figure 6 summarizes the results. As predicted, there is no overreaction in treatment *No Cue*. If anything, the data reveal slight underreaction. As discussed above, this result is consistent with a large set of findings from belief updating experiments in which associative recall cannot play a role by design (Benjamin, 2018; Enke and Graeber, 2019). In combination, the results from treatments *Main* and *No Cue* again suggest that associative recall is so strong that it turns slight shading into overreaction.

Columns (5)–(6) of Table 4 present a formal comparison of treatments *Main* and *No*

Table 4: Treatments *Main* vs. *Reminder* and *No Cue*

Treatments	2nd period belief *Main* vs. *Reminder*		2nd period WTP *WTP* vs. *WTP reminder*		2nd period belief *Main* vs. *No Cue*	
	(1)	(2)	(3)	(4)	(5)	(6)
2nd period signal	0.99***	0.98***	0.97***	0.96***	0.89***	0.88***
	(0.01)	(0.01)	(0.04)	(0.05)	(0.04)	(0.04)
2nd period signal × 1 if *Main*, 0 if *Reminder*	0.12***	0.14***				
	(0.03)	(0.03)				
2nd period signal × 1 if *WTP*, 0 if *Reminder WTP*			0.14**	0.17**		
			(0.06)	(0.07)		
2nd period signal × 1 if *Main*, 0 if *No Cue*					0.21***	0.22***
					(0.04)	(0.05)
Belief in 1st period	0.84***				0.62***	
	(0.02)				(0.03)	
WTP in 1st period			0.66***			
			(0.04)			
Treatment FE	Yes	Yes	Yes	Yes	Yes	Yes
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	1300	1300	1800	1800	1600	1600
Adjusted R^2	0.86	0.86	0.69	0.74	0.68	0.67

Notes. OLS estimates, robust standard errors (in parentheses) are clustered at the subject level. In columns (1)–(3), the sample includes all observations from treatments *Main* and *Reminder* where subjects observed a second-period signal. In columns (4)–(6), the sample includes all observations from treatments *Main* and *No Cue* where subjects observed a second-period signal. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Cue.[11] As specified in the pre-analysis plan, we link participants' second-period beliefs to the second-period signal, interacted with a treatment dummy. As predicted, the interaction term shows that subjects respond significantly more to the second-period signal in *Main* than in *No Cue*.

Result 4. *Overreaction disappears once associative recall is shut down.*

5.4 Over- vs. Underreaction

Next, we turn to investigating predictions 5–7 in Section 3.4, which conceptually correspond to setting the parameter $\rho = -1$ in the simple model. For this purpose, as

---

[11]We again replicate the analysis using the direct recall data in Table 13 in Appendix D.

Figure 6: OLS coefficient (+/- 1 SE) in a regression of second-period beliefs on the last signal, separately for each set of signal frequencies in the first period. The regressions control for a subject's first-period belief. Standard error bars are computed based on clustering at the subject level.

discussed in Section 3, we implemented treatments *Underreaction* and *Underreaction reminder*. Here, a second-period signal cues the recollection of the opposite past signals. To verify that potential underreaction in treatment *Underreaction* is not driven by some form of confusion that could arise from the change in the mapping from news to contexts after the first period, treatment *Underreaction reminder* serves as a control treatment.

Figure 7 summarizes the results for both treatments, separately for each signal history. The figure shows that (i) there is underreaction in *Underreaction* but not in *Underreaction reminder* and (ii) underreaction increases in the number of first-period signals, as predicted.

Table 5 presents the regression 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. Columns (3) and (4) leverage exogenous variation in signal histories to document that, as posited in Prediction 6, underreaction strongly increases in the number of first-period signals that were communicated in the same context as the second-period signal, see the statistically significant interaction term.[12]

Finally, columns (5) and (6) compare treatments *Underreaction* and *Underreaction reminder*. Again, the coefficient of interest is the interaction term between the second-period signal and a treatment dummy. The dummy is statistically highly significant and

---

[12]Table 14 in Appendix D shows that very similar results hold when we consider the recall data.

Figure 7: OLS coefficient (+/- 1 SE) in a regression of second-period beliefs on the last signal, separately for each set of signal frequencies in the first period. The regressions control for a subject's first-period belief. Standard error bars are computed based on clustering at the subject level. In treatment *Underreaction reminder*, there was no variation in the second-period signal for first-period signal history "3:0", so that this coefficient cannot be reported.

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. The evidence hence points to asymmetric recall as mechanism behind underreaction in the same way as it produced overreaction when $\rho = 1$.[13]

Result 5. *When the correlation between context and news is negative, beliefs underreact with respect to the second-period signal.*

Result 6. *Underreaction increases in the number of cued first-period signals.*

Result 7. *Underreaction disappears once memory imperfections are shut down.*

---

[13]To further corroborate the idea that underreaction is generated by asymmetric recall, Figure 11 in Appendix C analyzes the self-reported recall patterns in treatment *Underreaction* as a function of the signal history, akin to Figure 4 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* 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 from the second-period signal now get cued by the second-period context.

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

	Dependent variable: 2nd period belief					
	Treatments:				+ *Reminder*	
	Underreaction					
	(1)	(2)	(3)	(4)	(5)	(6)
2nd period signal	0.76***	0.74***	0.99***	0.95***	1.01***	1.01***
	(0.04)	(0.04)	(0.04)	(0.05)	(0.02)	(0.02)
Belief in 1st period	0.65***		0.50***		0.77***	
	(0.04)		(0.05)		(0.03)	
2nd period signal × # 1st period signals in same context			-0.31***	-0.28***		
			(0.06)	(0.07)		
2nd period signal × 1 if *Underr.*, 0 if *Reminder underr.*					-0.25***	-0.28***
					(0.04)	(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
p-value H0: β (2nd period signal)=1	<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. In columns (3)–(4), the table suppresses the coefficients of the number of first-period signals that were communicated in the same context as the second-period signal. The sample includes all observations from treatments *Underreaction* and *Underreaction reminder* where subjects observed a second-period signal. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

6 Extension: Mechanisms Behind Forgetting

At the most basic level, our paper is agnostic about what exactly generates imperfect recall per sé: we are mostly interested in identifying the role of associative memory conditional on the existence of forgetting in the first place. Nonetheless, this section provides a first pass at studying the relative importance of three potential mechanisms that could generate forgetting in our new experimental paradigm: (i) the existence of a time lag of 15 minutes between the first and second period; (ii) similarity-based interference, as it results from the presence of twelve similar companies with identical signal realizations (+10 / -10); and (iii) more general memory load, as it results from the presence of twelve simultaneous recall problems (even if these were not similar to each other). To provide some evidence on the relative importance of these potential mechanisms, we first manipulate the length of the time lag and then the presence of similarity-based interference. Of course, these mechanisms cannot be perfectly disentangled almost by

definition because (i) inducing interference takes time and, conversely, (ii) a time lag will under any realistic circumstances imply at least mild interference.

Extended Time Lag. Treatment *Extended time lag* followed the same procedure as treatment *Main*, except that the time lag between the first and second period of the experiment was three days. On the first day, subjects completed the first period of the experiment, using the same experimental instructions and control questions as in *Main*. After the first period, participants completed the real effort task, the Raven matrices test as well as the demographic questionnaire.[14] On the second day, participants re-read the original instructions and completed the same set of control questions again. Then, they completed the second period of the experiment. Due to the substantially increased time lag, we conducted treatment *Extended time lag reminder* as an additional bechmark condition, which is identical to treatment *Reminder*, except for the increased time lag.

These two treatments were also pre-registered in the original pre-registration. 80 subjects participated in treatment *Extended time lag* and 50 in treatment *Extended time lag reminder*. The treatments were randomized within experimental sessions and implemented in the BonnEconLab of the University of Bonn. Attrition was negligible: 95% of subjects returned for the second session.

Figure 9 in Appendix C and Tables 15 and 16 in Appendix D summarize 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. If anything, we find that overreaction is even stronger with a time lag of three days rather than 15 minutes: in the baseline regression, the OLS coefficient of the second-period signal increases from 1.10 in *Baseline* to 1.17 in *Extended time lag*. This is consistent with the model in Section 2 if one assumes that the strength of memory r decays over time, yet the strength of associative recall a remains constant or decays less.

No Time Lag. We now manipulate the length of the time lag in the opposite direction: in treatment *No time lag*, subjects start the second period immediately after the first period ends. Thus, the time lag is 15 minutes shorter than in treatment *Main*. Of course, for any given company, the time lag is not zero because participants also receive news for other companies. 60 subjects participated in this treatment, which was randomized within experimental sessions with a replication of treatment *Main* (also 60 subjects). Figure 10 in Appendix C presents the results and Table 18 in Appendix D shows corresponding regression analyses. We find that the OLS coefficient of the second-period

---

[14]Thus, subjects could not take notes right after the first period.

signal is 1.11 in both *Main replication* and *No time lag*, and statistically indistinguishable from each other. This suggest that the time lag of 15 minutes between the first and second period is inessential for forgetting and the experimental design more generally.

No interference. To study the importance of interference for forgetting, we introduce treatment *No interference*. The design of this treatment is guided by two design objectives: (i) reducing similarity-based interference, while holding overall memory load roughly constant and (ii) holding the time line of the experiment constant. In treatment *No interference*, there is only one company, which is randomly selected across participants from the set of twelve companies in the original experiments. In addition to this one company, subjects also completed eleven recall tasks that were designed to be similar to the main experimental task, without introducing interference via similarity. To this effect, we introduced eleven "groups" of colored shapes, where each group consisted of two shapes. For each of the groups, participants sequentially observed between zero and three shapes on their decision screen in the first period of the experiment. In the second period, they observed one more shape from each group and were then asked how many times they had seen shapes that belong to a particular group. We designed this experiment such that the timing was extremely similar to treatment *Main*: shapes were shown for the same period of time as news for companies, and we implemented the same 15 minutes time lag between the first and second period. We implemented this treatment with 60 additional subjects, randomized within session along with treatments *Main replication* and *No time lag* discussed above.

We find that, for the one company, overreaction almost entirely disappears: the OLS coefficient of the second-period signal is 1.01 and indistinguishable from the rational benchmark of one. Moreover, overreaction in *No interference* is significantly smaller than in treatment *Main replication*, see Figure 10 in Appendix C and Table 18 in Appendix D. In fact, as Figure 13 in Appendix C shows, subjects' recall of first-period signals is close to perfect once there is no interference.

It is worth highlighting that these results hold even though we attempted to keep working memory load constant relative to the baseline condition by introducing the eleven shapes recall tasks. Indeed, as we document in Table 19 in Appendix D, we *do* find strong associative recall for the shapes, presumably because there was similarity-based interference with respect to these shapes due to the presence of eleven similar recall problems.

While these results on the micro-foundations of memory imperfections are not the focus of our paper, we view them as suggestive that similarity-based interference plays an important role in generating forgetting, consistent with recent theoretical work in both economics (Bordalo et al., 2019b) and psychology (Kahana, 2012).

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

$$b^{i,j} = 100 + \gamma n_{k+1}^{i,j} + r \sum_{x=1}^{k} n_x^{i,j} + (1-r)a \sum_{x=1}^{z} n_x^{i,j} + \epsilon^{i,j}, \tag{10}$$

where γ measures an individual's intrinsic responsiveness to the second-period signal.

Aggregate data.   To start, we estimate this equation on the aggregate data across subjects, separately for each treatment. Table 6 summarizes the estimates. The quantitative estimates are in line with the results reported 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). Analogously, we see that in treatment *No Cue*, associative recall collapses to zero, again by construction of the treatment. Related to the discussion of the mechanisms behind memory imperfections in Section 6, we see that the estimated memory imperfection is larger in condition *Extended time lag* than in *Main*. In treatment *No interference*, the estimated memory imperfection drops to almost zero.

Finally, we see that subjects' intrinsic responsiveness to the second-period signal γ (conditional on memory parameters) is consistently estimated to be less than one, unless there are no memory imperfections. This again resonates with a large body of work on belief formation that is summarized in the recent meta-study by Benjamin (2018): intrinsically, subjects usually *underreact* to current news. However, in our setup, the effects of associative recall are sufficiently strong to turn such intrinsic underreaction into overall overreaction.

Individual-level data.   Next, we estimate the same model, separately for each individual. To assess the fit of the model at the individual level, we use the individual-level estimates of $\hat{r}_i$ and $\hat{a}_i$ to predict participant i's reported recall of those first-period news

Table 6: Estimates of model parameters across treatments

Treatment	Forgetting $(1 - \hat{r})$	Associative recall $\hat{a}$	Responsiveness $\hat{\gamma}$
Main	0.41***	0.79***	0.86***
	(0.05)	(0.07)	(0.04)
Reminder	0.01	-1.59	1.00***
	(0.01)	(4.83)	(0.01)
No Cue	0.51***	0.01	0.88***
	(0.05)	(0.11)	(0.05)
Extended time lag	0.68***	0.65***	0.85***
	(0.05)	(0.06)	(0.05)
Main replication	0.48***	0.53***	0.93***
	(0.06)	(0.09)	(0.03)
No time lag	0.43***	0.53***	0.94***
	(0.06)	(0.10)	(0.04)
No interference	0.04	0.64	0.99***
	(0.04)	(1.04)	(0.05)

Notes. Estimates of equation (10), standard errors (clustered at subject level) reported in parentheses. The model is estimated by pooling the data across subjects in a given treatment. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

that did (q_i^c) or did not (q_i^n) get cued by the second-period signal:

$$\hat{q}_i^n = \hat{r}_i * (k - z) \tag{11}$$

$$\hat{q}_i^c = [\hat{r}_i + (1 - \hat{r}_i)\hat{a}_i] * z \tag{12}$$

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. Thus, comparing *predicted* with *actual* recall allows for an assessment of model fit.

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 $\rho = 0.82$. The correlation between predicted and actual recall of those signals that did not get cued is $\rho = 0.67$, see Figure 14 in Appendix C. We interpret these results as encouraging evidence that our simple two-parameter memory model fits the observed data well.

8 Discussion

This paper has provided a theoretically-structured experimental analysis of the role of 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. We present the first set of theory-driven experiments that build a bridge between psychological paradigms on cued recall and structured, quantitative economic decision tasks. Our experiments and estimations suggest a predictable and quantitatively meaningful role for associative memory in belief formation.

Our experiments are potentially related to an active literature that documents over-reaction or excess movement in survey expectations about macroeconomic variables or geopolitical events (Augenblick and Rabin, 2018; Bordalo et al., 2018). The result of overreaction in field data is often considered to be a slight puzzle from the perspective of laboratory research on belief formation. This is because structured laboratory belief updating problems almost always find underreaction (Benjamin, 2018), at least partly due to participants' cognitive uncertainty (Enke and Graeber, 2019). However, in these laboratory experiments, memory imperfections are by design ruled out. We do not intend to claim that associative recall can explain the entire pattern of over- and under-reaction identified in the literature. However, it is conceivable that *part* of the reason why the laboratory and field literatures identify such different patterns is that memory constraints and memorable contexts likely play a more important role in the field, as exemplified by Shiller's (2017; 2019) discussion of the role of memorable "narratives." 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

Augenblick, Ned 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.

Bernheim, B Douglas and Antonio Rangel, "Addiction and cue-triggered decision processes," *American economic review*, 2004, *94* (5), 1558–1590.

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*, 2019.

_ , _ , **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," *The Review of Economic Studies*, 2019, *86* (1), 313–332.

_ **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, 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.

Hartzmark, Samuel M, Samuel Hirshman, and Alex Imas, "Ownership, Learning, and Beliefs," *Available at SSRN 3465246*, 2019.

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.

Huffman, David, Collin Raymond, and Julia Shvets, "Persistent Overconfidence and Biased Memory: Evidence from Managers," 2018.

Kahana, Michael Jacob, *Foundations of human memory*, OUP USA, 2012.

Koszegi, Botond, George Loewenstein, and Takeshi Murooka, "Fragile Self-Esteem," *Working Paper*, 2019.

Laibson, David, "A cue-theory of consumption," *The Quarterly Journal of Economics*, 2001, *116* (1), 81–119.

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.

Schwerter, Frederik and Florian Zimmermann, "Determinants of trust: The role of personal experiences," *Working Paper*, 2019.

Shiller, Robert J, "Narrative economics," *American Economic Review: Papers and Proceedings*, 2017, *107* (4), 967–1004.

_ , *Narrative Economics: How Stories Go Viral and Drive Major Economic Events*, Princeton University Press, 2019.

Tversky, Amos, "Features of similarity.," *Psychological review*, 1977, *84* (4), 327.

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

A.1.1 Type I naïveté

The main text assumes that decision-makers are fully naïve about their memory imper-
fections. We now verify robustness against assuming partial naïveté. Suppose the DM
to some extent (captured by naïveté parameter α such that $\alpha = 0$ for full naïveté and
$\alpha = 1$ for full sophistication) fails to realize that he sometimes forgets. When he does
realize that he forgot a past signal n_x, however, then he correctly (in a Bayesian sense)
infers the realization of the information based on memory parameters r and a.

We have that $Pr(n_x = n_{k+1}|norecall, n_{k+1}) = \frac{(1-r)-(1-r)a}{(1-r)-(1-r)a+(1-r)} = \frac{1-a}{2-a}$. Accordingly,
we have that $E(n_x|norecall, n_{k+1}) = n_{k+1}\frac{1-a}{2-a} - n_{k+1}\frac{1}{2-a} = \frac{-a}{2-a}n_{k+1}$.

The expected belief in period $t = 2$ is then given by:

$$
\begin{aligned}
E[b_2|n_x, n_{k+1}] &= v + n_{k+1} + \sum_{x=1}^{k} m_x n_x + \alpha \sum_{x=1}^{k}(1-m_x)\frac{-a}{2-a}n_{k+1} \\
&= v + n_{k+1} + \sum_{x=1}^{k} r n_x + \sum_{x=1}^{z}(1-r)a n_x \\
&\quad + \alpha \sum_{x=1}^{k}(1-r)\frac{-a}{2-a}n_{k+1} + \alpha \sum_{x=1}^{z}(1-r)a\frac{-a}{2-a}n_{k+1} \qquad (13) \\
&= v + [1 + \rho(1-r)a(z + \alpha z\frac{a}{2-a} - \alpha k\frac{1}{2-a})]n_{k+1} + r\sum_{x=1}^{k} n_x \qquad (14)
\end{aligned}
$$

Note that equation (14) mirrors equation (6) from Section 2. Equation (14) allows
us to directly analyze the implications of allowing for partial naïveté.

We first note that if α is small ($\alpha \to 0$), then equation (14) converges to equation (6).

Second, we note that "on average" (across possible signal histories), equation (14)
still predicts overreaction for all leves of $\alpha < 1$ and $\rho = 1$. To see this, note that there is
overreaction as long as $z + \alpha z\frac{a}{2-a} - \alpha k\frac{1}{2-a} > 0$. For a given k, z is randomly and symmet-
rically distributed with mean $\frac{k}{2}$. Due to the linear structure of equation (14), it therefore
suffices to note that $z + \alpha z\frac{a}{2-a} - \alpha k\frac{1}{2-a} > 0$ for $z = \frac{k}{2}$ and $\alpha < 1$.

A.1.2 Type II naïveté

According to type II naïveté, the DM fully realizes that he sometimes forgets, but is naïve in how she infers what he forgot. This form of naïveté is captured by the DM's belief $\hat{a}$ about memory parameter a, $\hat{a} \leq a$. Here, $\hat{a} = 0$ captures full naïvité, meaning that the DM is aware of imperfect memory but not of the associative nature of recall. $\hat{a} = 1$ captures full sophistication, meaning that the DM fully takes into account that he is more likely to retrieve information that is cued by the current context.

The DM's inference would then be as outlined in the previous section, except that it would use $\hat{a}$. The DM's forecast would thus be given by $f = v + n_{k+1} + \sum_{x=1}^{k} m_x n_x + \sum_{x=1}^{k} (1 - m_x) \frac{-\hat{a}}{2 - \hat{a}} n_{k+1}$.

Analogous to type I naïveté, the expected belief in period $t = 2$ is given by:

$$
\begin{aligned}
E[b_2 | n_x, n_{k+1}] &= v + n_{k+1} + \sum_{x=1}^{k} m_x n_x + \sum_{x=1}^{k} (1 - m_x) \frac{-\hat{a}}{2 - \hat{a}} n_{k+1} \\
&= v + n_{k+1} + \sum_{x=1}^{k} r n_x + \sum_{x=1}^{z} (1 - r) a n_x \\
&\quad + \sum_{x=1}^{k} (1 - r) \frac{-\hat{a}}{2 - \hat{a}} n_{k+1} + \sum_{x=1}^{z} (1 - r) a \frac{-\hat{a}}{2 - \hat{a}} n_{k+1} \qquad (15) \\
&= v + \left[1 + \rho (1 - r) \left(az + az \frac{\hat{a}}{2 - \hat{a}} - k \frac{\hat{a}}{2 - \hat{a}} \right) \right] n_{k+1} + r \sum_{x=1}^{k} n_x \qquad (16)
\end{aligned}
$$

As above, we first note that if $\hat{a}$ is small ($\hat{a} \to 0$), then equation (16) converges to equation (6).

Moreover, by an analogous argument to the previous section, there would again be overreaction as long as $\hat{a} < 1$. To see this, note that there is overreaction as long as $az + az \frac{\hat{a}}{2 - \hat{a}} - k \frac{\hat{a}}{2 - \hat{a}} > 0$. For a given k, z is randomly and symmetrically distributed with mean $\frac{k}{2}$. Due to the linear structure of equation (16), it therefore suffices to note that $az + az \frac{\hat{a}}{2 - \hat{a}} - k \frac{\hat{a}}{2 - \hat{a}} > 0$ for $z = \frac{k}{2}$ and $\hat{a} < 1$.

A.2 Derivation of OLS Estimator

We formally derive the relationship between equation (6) and the OLS estimator $\hat{\beta}_1$ in equation (9). Keeping the notation that b is the belief and n the news, then with N observations (subject-tasks) the OLS estimator is given by

$$E[\hat{\beta}] = E\left[\frac{\sum n_i b_i - 1/N \sum n_i \sum b_i}{\sum n_i^2 - 1/N(\sum n_i^2)}\right] \tag{17}$$

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

$$E[\hat{\beta}] = \frac{\sum E[n_i b_i] - 1/N \sum E[n_i] \sum E[b_i]}{\sum E[n_i^2] - 1/NE(\sum n_i^2)} \tag{18}$$

$$= \frac{\sum E[n_i[n_i(1 + z_i(1 - r)\rho a) + c]]}{\sum E[n_i^2]} \tag{19}$$

$$= 1 + (1 - r)\rho a \frac{\sum E[n_i^2 z_i]}{\sum E[n_i^2]} \tag{20}$$

$$= 1 + (1 - r)\rho a \frac{\sum E[n_i^2] E[z_i]}{\sum E[n_i^2]} \tag{21}$$

$$= 1 + (1 - r)\rho a \frac{\bar{z} \sum E[n_i^2]}{\sum E[n_i^2]} \tag{22}$$

$$= 1 + (1 - r)\rho a \bar{z} \tag{23}$$

Because z_i and n_i are independent.

B   Overview of Experimental Treatments

Table 7: Treatment overview

Treatment	# of subjects	Ave. earnings (euros)	Pre-registration document
Main	80	15.20	1
Reminder	50	17.80	1
No Cue	80	14.00	1
Extended time lag	80	24.00	1
Extended time lag reminder	50	27.50	1
Underreaction	80	14.70	2
Underreaction reminder	50	18.00	2
WTP	100	19.10	3
WTP reminder	80	18.80	3
Main replication	60	13.40	4
No time lag	60	12.40	4
No interference	60	19.20	4

Notes. Horizontal lines indicate which treatments were randomized within the same experimental sessions. Payments included a show-up fee of €15 in *Extended time lag / Extended time lag reminder* and of €5 in all other treatments.

C   Additional Figures

Figure 8: OLS coefficient (+/- 1 SE) in a regression of second-period willingness-to-pay on the last signal, separately for each set of signal frequencies in the first period. The regressions control for a subject's first-period willingness-to-pay. Standard error bars are computed based on clustering at the subject level.

OLS coefficient of 2nd period signal by signal history

Figure 9: OLS coefficient (+/- 1 SE) in a regression of second-period beliefs on the last signal, separately for each set of signal frequencies in the first period. The regressions control for a subject's first-period belief. Standard error bars are computed based on clustering at the subject level.

OLS coefficient of 2nd period signal by signal history

Figure 10: OLS coefficient (+/- 1 SE) in a regression of second-period beliefs on the last signal, separately for each set of signal frequencies in the first period. The regressions control for a subject's first-period belief. Standard error bars are computed based on clustering at the subject level. The figure excludes signal histories with less than four observations.

44

Recall of 1st period signals in treatment *Underreaction*

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.

Figure 12: Subject-level distribution of regression coefficients of the second-period signal in treatment *WTP* (N=100). To estimate these coefficients, we run regressions akin to column (5) in Table 2 except that in each regression the sample is restricted to only one subject. To 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 exhibit a coefficient of one.

Recall of 1st period signals in treatment *No interference*

Figure 13: Recall of first-period signals in Treatment *No interference*, 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.

Figure 14: Relationship between recall as predicted by the model estimates and actual recall. 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 (10) at the subject level and then applying equations (11) and (12).

D   Additional Tables

Table 8: Beliefs in the first period in treatments *Main*, *Reminder*, and *No Cue*

| | *Dependent variable:* 1st period belief | | | | | |
| | Main | | Reminder | | No Cue | |
	(1)	(2)	(3)	(4)	(5)	(6)
Company value in 1st period	0.98***	0.97***	1.00***	1.00***	0.99***	1.00***
	(0.01)	(0.01)	(0.00)	(0.00)	(0.01)	(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 9: Beliefs in the second period in case of no signal in second period

| | *Dependent variable:* 2nd period belief | | | | | |
| | Main | | Reminder | | No Cue | |
	(1)	(2)	(3)	(4)	(5)	(6)
Belief in 1st period	0.70***	0.70***	0.95***	0.94***	0.48***	0.45***
	(0.05)	(0.06)	(0.04)	(0.05)	(0.05)	(0.06)
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	0.29	0.28	0.57	0.58	0.17	0.16

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*: Recall data

| | *Dependent variable:* | | | | | | |
| | Δ Recall [Pos. – Neg.] | | | | | | |
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
2nd period signal	1.05***	1.06***	1.07***	0.81***	0.83***	0.79***	0.83***
	(0.04)	(0.03)	(0.04)	(0.05)	(0.05)	(0.05)	(0.05)
Belief in 1st period	0.74***					0.57***	
	(0.03)					(0.05)	
Company value in 1st period		0.74***					
		(0.03)					
Value of cued 1st period signals				0.92***	0.90***		
				(0.03)	(0.03)		
Value of non-cued 1st period signals				0.58***	0.58***		
				(0.05)	(0.05)		
2nd period signal × # 1st period signals in same context						0.36***	0.31***
						(0.05)	(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. The Δ recall variable is constructed as difference between reported recall of positive and negative signals. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 11: Treatment *Main*: Heterogeneity analysis

	Dependent variable: 2nd period belief					
	(1)	(2)	(3)	(4)	(5)	(6)
2nd period signal	1.26***	1.25***	1.23***	1.21***	1.14***	1.14***
	(0.08)	(0.09)	(0.04)	(0.04)	(0.04)	(0.04)
2nd period signal × Raven score	-0.030**	-0.028*				
	(0.01)	(0.01)				
2nd period signal × Memory for non-cued signals			-0.23***	-0.20***		
			(0.06)	(0.06)		
2nd period signal × Response time					-0.37	-0.35
					(0.29)	(0.24)
Belief in 1st period	0.75***		0.75***		0.75***	
	(0.03)		(0.03)		(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 12: Treatment *Main*: Heterogeneity analysis (recall data)

	(1)	(2)	(3)	(4)	(5)	(6)
	Dependent variable: Δ Recall [Pos. – Neg.]					
2nd period signal	1.29***	1.27***	1.21***	1.19***	1.15***	1.15***
	(0.13)	(0.13)	(0.06)	(0.05)	(0.05)	(0.05)
2nd period signal × Raven score	-0.045*	-0.042*				
	(0.02)	(0.02)				
2nd period signal × Memory for non-cued signals			-0.30***	-0.26***		
			(0.08)	(0.08)		
2nd period signal × Response time recall					-0.45**	-0.50***
					(0.19)	(0.19)
Belief in 1st period	0.74***		0.74***		0.74***	
	(0.03)		(0.03)		(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 Δ recall variable is constructed as difference between reported recall of positive and negative signals. 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 13: Treatments *Main* vs. *Reminder* and *No Cue*: Recall data

| | *Dependent variable:* Δ Recall [Pos. – Neg.] | | | | | |
| | *Main* vs. *Reminder* | | | *Main* vs. *No Cue* | | |
	(1)	(2)	(3)	(4)	(5)	(6)
2nd period signal	0.95***	0.95***	0.94***	0.76***	0.75***	0.77***
	(0.02)	(0.02)	(0.02)	(0.05)	(0.05)	(0.06)
2nd period signal × 1 if *Main*, 0 if *Reminder*	0.11**	0.12***	0.13***			
	(0.04)	(0.04)	(0.04)			
2nd period signal × 1 if *Main*, 0 if *No Cue*				0.29***	0.29***	0.30***
				(0.06)	(0.06)	(0.07)
Belief in 1st period	0.83***			0.60***		
	(0.02)			(0.03)		
Company value in 1st period		0.83***			0.61***	
		(0.02)			(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	1300	1300	1300	1600	1600	1600
Adjusted R^2	0.82	0.83	0.83	0.62	0.63	0.63

Notes. OLS estimates, robust standard errors (in parentheses) are clustered at the subject level. In columns (1)–(3), the sample includes all observations from treatments *Main* and *Reminder* where subjects observed a second-period signal. In columns (4)–(6), the sample includes all observations from treatments *Main* and *No Cue* where subjects observed a second-period signal. The Δ recall variable is constructed as difference between reported recall of positive and negative signals. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 14: Treatments *Underreaction* and *Underreaction reminder*: Recall data

	Dependent variable: 2nd period belief					
	Treatments:				+ *Reminder*	
	Underreaction					
	(1)	(2)	(3)	(4)	(5)	(6)
2nd period signal	0.62***	0.60***	0.86***	0.81***	0.89***	0.91***
	(0.05)	(0.05)	(0.05)	(0.06)	(0.03)	(0.04)
Belief in 1st period	0.66***		0.49***		0.76***	
	(0.04)		(0.05)		(0.03)	
2nd period signal × # 1st period signals in same context			-0.33***	-0.28***		
			(0.06)	(0.06)		
2nd period signal × 1 if *Underr.*, 0 if *Reminder underr.*					-0.28***	-0.32***
					(0.06)	(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 Δ recall variable is constructed as difference between reported recall of positive and negative signals. In columns (3)–(4), the table suppresses the coefficient of the number of first-period signals that were communicated with the same context as tge second-period signal. 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 15: Treatments *Extended time lag* and *Extended time lag reminder*

	Dependent variable: 2nd period belief					
	Treatments:					
	Time lag				+ *Reminder time lag*	
	(1)	(2)	(3)	(4)	(5)	(6)
2nd period signal	1.17***	1.16***	0.85***	0.84***	1.02***	1.03***
	(0.04)	(0.04)	(0.05)	(0.06)	(0.02)	(0.02)
Belief in 1st period	0.52***		0.33***		0.71***	
	(0.04)		(0.05)		(0.03)	
2nd period signal × # 1st period signals in same context			0.44***	0.45***		
			(0.05)	(0.07)		
2nd period signal × 1 if *Time lag*, 0 if *Reminder t. l.*					0.18***	0.15***
					(0.04)	(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. In columns (3)–(4), the table suppresses the coefficient of the number of first-period signals that were communicated with the same context as tge second-period signal. The sample includes all observations from treatments *Extended time lag* and *Extended time lag reminder* where subjects observed a second-period signal. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 16: Treatments *Extended time lag* and *Extended time lag reminder*: Recall data

	Time lag				+ Reminder time lag	
	Dependent variable: Δ Recall [Pos. – Neg.]					
	Treatments:					
	(1)	(2)	(3)	(4)	(5)	(6)
2nd period signal	0.97***	0.96***	0.66***	0.66***	0.97***	0.98***
	(0.05)	(0.06)	(0.05)	(0.06)	(0.03)	(0.03)
Belief in 1st period	0.52***		0.33***		0.70***	
	(0.04)		(0.05)		(0.03)	
2nd period signal × # 1st period signals in same context			0.43***	0.43***		
			(0.06)	(0.07)		
2nd period signal × 1 if *Time lag*, 0 if *Reminder t. l.*					0.022	-0.0011
					(0.06)	(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 Δ recall variable is constructed as difference between reported recall of positive and negative signals. In columns (3)–(4), the table suppresses the coefficient of the number of first-period signals that were communicated with the same context as tge second-period signal. The sample includes all observations from treatments *Extended time lag* and *Extended time lag reminder* where subjects observed a second-period signal. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 17: Treatments *Main replication*, *No time lag*, and *No interference*

	Dependent variable: 2nd period belief									
	Treatments:									
	Main repl.		*No time lag*		*No interference*		*Main repl. + No time lag*		*Main repl. + No interf.*	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
2nd period signal	1.11***	1.14***	1.11***	1.11***	1.01***	0.92***	1.11***	1.12***	1.11***	1.14***
	(0.03)	(0.04)	(0.03)	(0.04)	(0.04)	(0.06)	(0.03)	(0.04)	(0.03)	(0.03)
Belief in 1st period	0.65***		0.67***		0.98***		0.66***		0.68***	
	(0.05)		(0.05)		(0.03)		(0.04)		(0.05)	
2nd period signal × 1 if *No time lag*, 0 if *Main repl.*							-0.0011	-0.0058		
							(0.05)	(0.05)		
2nd period signal × 1 if *No interf.*, 0 if *Main repl.*									-0.16**	-0.33***
									(0.07)	(0.08)
Treatment FE	No	No	No	No	No	No	Yes	Yes	Yes	Yes
Session FE	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes
Signal history FE	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes
Company FE	No	Yes	No	Yes	No	No	No	Yes	No	Yes
Order FE	No	Yes	No	Yes	No	No	No	Yes	No	Yes
Subject FE	No	Yes	No	Yes	No	No	No	Yes	No	No
Observations	600	600	600	600	60	60	1200	1200	660	660
Adjusted R^2	0.76	0.76	0.74	0.76	0.95	0.95	0.75	0.76	0.77	0.77

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

Table 18: Treatment *Main replication*, *No time lag*, and *No interference*: Recall data

| | Main repl. | | No time lag | | No interference | | Main repl. + No time lag | | Main repl. + No interf. | |
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
2nd period signal	1.06***	1.07***	1.04***	1.03***	1.07***	1.03***	1.06***	1.06***	1.06***	1.08***
	(0.04)	(0.04)	(0.05)	(0.05)	(0.05)	(0.04)	(0.04)	(0.05)	(0.04)	(0.04)
Belief in 1st period	0.68***		0.69***		1.00***		0.68***		0.70***	
	(0.05)		(0.04)		(0.03)		(0.03)		(0.04)	
2nd period signal × 1 if *No time lag*, 0 if *Main repl.*							-0.020	-0.031		
							(0.06)	(0.07)		
2nd period signal × 1 if *No interf.*, 0 if *Main repl.*									-0.050	-0.17*
									(0.08)	(0.09)
Treatment FE	No	No	No	No	No	No	Yes	Yes	Yes	Yes
Session FE	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes
Signal history FE	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes
Company FE	No	Yes	No	Yes	No	No	No	Yes	No	Yes
Order FE	No	Yes	No	Yes	No	No	No	Yes	No	Yes
Subject FE	No	Yes	No	Yes	No	No	No	Yes	No	No
Observations	600	600	600	600	60	60	1200	1200	660	660
Adjusted R^2	0.73	0.73	0.71	0.71	0.94	0.97	0.72	0.72	0.74	0.74

The dependent variable across all columns is Δ Recall [Pos. – Neg.], shown under "Treatments:".

Notes. OLS estimates, robust standard errors (in parentheses) are clustered at the subject level. The Δ recall variable is constructed as difference between reported recall of positive and negative signals. The sample includes all observations from treatments *Main replication, No time lag*, and *No interference* where subjects observed a second-period signal. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 19: Treatment *No interference*: Associative memory in recall of colored shapes

	Dependent variable: Recall of shapes in group		
	(1)	(2)	(3)
# 1st period shapes in same context	0.74***	0.68***	0.70***
	(0.05)	(0.11)	(0.10)
# 1st period shapes in different context	0.60***	0.53***	0.56***
	(0.06)	(0.12)	(0.11)
Session FE	No	Yes	Yes
Shape history FE	No	Yes	Yes
Group FE	No	Yes	Yes
Order FE	No	Yes	Yes
Subject FE	No	No	Yes
Observations	540	540	540
Adjusted R^2	0.45	0.46	0.55

Notes. OLS estimates, robust standard errors (in parentheses) are clustered at the subject level. The sample includes all observations from treatment *No interference* where subjects reported the recall of the number of shapes in a given group. $^{*}\ p < 0.10$, $^{**}\ p < 0.05$, $^{***}\ p < 0.01$.

E Experimental Instructions

We provide translations of the paper-based instructions here. An English version of the computer program for treatment *Main* (where subjects observe news and enter their guesses) can be accessed at `https://unikoelnwiso.eu.qualtrics.com/jfe/form/SV_0MrVD2rNNrKeLGt`.

E.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:

Stock price = Base price + News 1 + 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:

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.

E.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.

E.3 Treatment *No Cue*

Instructions for treatment No Cue 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.

F   Experimental Materials

F.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 15: 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 16: 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.