



Associative memory, beliefs and market interactions[☆]

Benjamin Enke^{a,b}, Frederik Schwerter^{c,d}, Florian Zimmermann^{e,f,*}

^a Harvard University, United States of America

^b NBER, United States of America

^c Frankfurt School of Finance and Management, Germany

^d University of Cologne, Germany

^e University of Bonn, Germany

^f IZA, Germany

ARTICLE INFO

Dataset link: <https://data.mendeley.com/datasets/545wb2gnd7/1>

JEL classification:
D01

Keywords:
Beliefs
Expectations
Memory
Associations
Bounded rationality

ABSTRACT

Recent theories and narratives highlight the potential role of associative recall in driving overreaction in expectations and market behavior. Based on a simple model, we test this idea through a series of experiments in which news are communicated with memorable contexts. Because the experimental participants predominantly remember those past news that get cued by new information, their beliefs about fundamentals strongly overreact. In a betting market experiment, associative recall translates into overreaction in market prices, which makes realized prices too extreme. Our results highlight the importance of associative memory for beliefs and financial decisions.

1. Introduction

Expectations play a crucial role in financial decision making. In recent years, a growing literature has documented that investors' expectations (i) strongly predict investment behavior (e.g., [Beutel and Weber, 2022](#); [Giglio et al., 2021](#)); yet (ii) often appear at odds with the rational expectations assumption (e.g., [Greenwood and Shleifer, 2014](#); [Bordalo et al., 2020a](#)). These results have raised the question which psychological primitives researchers in economics and finance should incorporate into theoretical models and empirical applications to account for deviations from the rational benchmark. As part of this movement, investors' memory has attracted interest because the abundance of financial information makes it implausible that investors

accurately recall all relevant news (e.g., [Charles, 2022a,b](#); [Jiang et al., 2023](#); [Afrouzi et al., 2023](#); [Malmendier and Wachter, 2021](#)).¹

A hallmark result in memory research in psychology is that recall is *associative* in nature, meaning that people are more likely to remember past information that is similar to new information.² Formal economic theories and popular financial narratives – discussed below – argue that the associative nature of memory could generate systematic *overreaction* of expectations and market behavior ([Shiller, 2017, 2019](#); [Gennaioli and Shleifer, 2018](#)). The main idea is that financial information can have both a direct and an indirect “cueing” effect on expectations. The direct effect is standard: upon receiving relevant news, investors update their beliefs. The indirect memory-based effect emerges because

[☆] Dimitris Papanikolaou was the editor for this article. We thank him, an associate editor and an anonymous referee for helpful comments. We also thank Pedro Bordalo, Kai Barron, Katie Coffman, Ernst Fehr, Nicola Gennaioli, Thomas Graeber, Alex Imas, Botond Koszegi, Ryan Oprea, Chris Roth, Peter Schwardmann, Josh Schwartzstein, Jesse Shapiro, Andrei Shleifer, Yves le Yaouanq, and seminar audiences at Bergen, Chicago Booth, Cologne, Copenhagen, Davis, Innsbruck, Maastricht, München, Münster, Wien, Zurich, the 2019 CESifo Behavioral Economics conference in Munich, and the 2021 Beliefs, Narratives, and Memory workshop in Riederau for valuable comments. Schwerter thanks the Sloan Foundation for post-doctoral funding. This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (ERC Starting Grant: 948424 - MEMEB). Funding by the Deutsche Forschungsgemeinschaft (DFG), Germany through CRC TR 224 (Project A01) and under Germany's Excellence Strategy - EXC 2126/1 - 390838866 is gratefully acknowledged. All experiments in this paper were pre-registered in the AEA RCT Registry as AEARCTR-0009215.

* Corresponding author.

E-mail addresses: enke@fas.harvard.edu (B. Enke), f.schwerter@fs.de (F. Schwerter), florian.zimmermann@iza.org (F. Zimmermann).

¹ Stefan Nagel, for instance, surmises that “[m]ore research, both empirically and theoretically, is needed to better understand investors' formation of memory”, see [Brunnermeier et al. \(2021\)](#).

² See [Baddeley et al. \(2020\)](#), [Schacter \(2008\)](#), [Kahana \(2012\)](#).

financial information is often embedded in contextual features such as stories, narratives and images. Given the associative nature of memory, these strong contextual elements can induce investors to reconstruct past information in a biased way, by asymmetrically retrieving those past news that are similar to today's news. For example, when investors are exposed to "good-times narratives" or images of bulls and upward-sloping trend lines, this may lead them to asymmetrically remember positive events from the past. As a consequence of this asymmetric recall, information and its associated contextual features can have an additional, indirect effect that could lead expectations to *look like* they overreact to news.

Despite the increasing theoretical interest in memory and its intuitive importance for expectation formation, clean evidence on its role for financial decisions is limited. A major hurdle in studying the role of associative recall is that it is difficult to causally identify the indirect "cueing" effects of narratives or images. The reason is that contextual features like stories are generally correlated with objective information, which makes it difficult to infer whether an expectations revision is actually driven by the indirect (cueing) effect of a narrative or by the standard direct effect of information. In laboratory experiments, on the other hand, objective information and contextual features that trigger recall can be decoupled and exogenously varied.

We, hence, implement controlled lab experiments to study the role of associative recall for expectation formation and market behavior. Our experiments are tightly organized around the predictions of a simple formal framework that applies the idea of associative recall to belief formation, following models such as Mullainathan (2002), Wachter and Kahana (2019) and Bordalo et al. (2020b). The model transparently spells out how the direct and indirect effects of information can interact to produce overreaction. In the context of a simple experimental asset pricing experiment, we obtain two key results. First, we show that associative recall causes predictable overreaction to financial news. Direct recall measures confirm that this result is indeed driven by associative memory. Second, using a betting market experiment, we document that associations-driven overreaction in beliefs translates into overreaction in market prices, even when people have an opportunity to fully or partially select out of the market.

Experimental design. In our experiment, participants form beliefs about whether each of multiple hypothetical companies is "good" or "bad". The experiment comprises two periods that we think of as "past" and "present". Across both periods, a subject sequentially observes noisy binary signals that are informative about a company's true quality. These news are communicated in a context, which consists of a story (narrative) and an image that relate to the 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. A key feature of our experimental implementation – that is difficult to achieve with observational data – is that the contextual cues are transparently uninformative (conditional on the signal realization). This allows us to identify the indirect memory-driven "cueing" effect of information, above and beyond the objective informational content of the signals.

Subjects' financial incentives are such that their second-period beliefs about a company's quality should incorporate both first- and second-period signals. Our main object of interest is whether subjects overreact to the second-period signal due to the logic of associative recall.

We deploy two types of random variation to causally identify the role of associative recall. First, we manipulate the scope for associative memory in a within-subjects treatment variation. Each subject forms beliefs about each of 14 companies, seven of which belong to treatment *Cue* and seven of which belong to treatment *NoCue*. For companies in treatment *Cue*, identical news are embedded in identical contexts. In other words, for each company in treatment *Cue*,

all positive news are communicated with the same context, and all negative news are communicated with the same context. Thus, in this treatment, the second-period signal could cue the asymmetric retrieval of identical first-period signals. While it is rarely the case that real market participants experience multiple signals in exactly the same context, this simple setup is reflective of many applications in which similar signals are consistently associated with similar contexts. For example, whenever good news prevail in the stock market, people are disproportionately exposed to bulls, upward-sloping trend lines, and good-times stories.

For companies in treatment *NoCue*, on the other hand, each signal is communicated with a different context. Thus, subjects never observe the same story or image twice. As a result, the scope for associative recall is exogenously reduced.

A second dimension of random variation is that, within treatment *Cue*, the number of first-period signals that equal the second-period signal (the number of signals that "get cued") differs randomly across subjects and companies. This is relevant because our stylized model predicts that overreaction of second-period beliefs should systematically depend on the number of cued first-period signals, even though – conditional on first-period beliefs – the signal history is irrelevant from a normative perspective.

Results. We report two main results. First, second-period beliefs overreact substantially to the second-period signal in *Cue*, an effect that is not present in *NoCue*. This identifies a causal effect of the presence of (normatively irrelevant) contextual associations on overreaction. Second, as predicted by the model, the magnitude of overreaction strongly increases in the number of first-period signals that get cued by the second-period signal. Thus, associations cause overreaction, and this overreaction is history-dependent in specific ways predicted by models of associative recall.

Our interpretation of this overreaction is that subjects asymmetrically remember those first-period signals that equal the second-period one. We confirm this by implementing a second experiment in which we directly elicit people's recall in an incentivized fashion.

Market experiments. Almost all of the narratives on the role of associative recall for financial decisions that motivate our paper concern *market behavior*, rather than purely individual decision making. This raises the question whether associations-driven overreaction in beliefs indeed potentially also impacts market behavior. One widely-discussed reason why this need not be the case is potential self-selection: people's expectations may overreact, but perhaps those people who succumb to associative recall and overreaction would never actually bet on these beliefs in a market context because they are loosely aware of their fallibility. Yet, the psychological literature provides little empirical guidance on whether we should expect people with stronger associative memory to act less aggressively on their beliefs.

To study this, we embed our individual belief elicitation paradigm into a parimutuel betting market experiment that is frequently used by experimental researchers due to its resemblance to real financial markets. In this market, groups of three subjects each receive public signals about a company's value and then place bets on whether the company is good or bad. The parimutuel price mechanism redistributes the money bet among the market participants according to whose bet was right and the amount of money bet. A crucial feature of this market is that subjects can self-select in or out in a continuous fashion, by choosing the total amount of money they would like to bet.

Despite this potential for self-selection, we find that market prices react about twice as strongly to information when associative recall is facilitated compared to when associations are experimentally removed. As a result, just like memorable contexts induce individual beliefs to be too extreme, they also lead market prices to be too high (too low) following a positive (negative) signal. Indeed, we find that the magnitude of associations-driven overreaction in market prices is very similar to the magnitude of overreaction in beliefs.

External relevance. Our experiments provide evidence that the associative nature of recall is a sufficiently hard-wired psychological process that it even affects people's financial information processing in unfamiliar online environments like ours. Given the pervasiveness of images, narratives and stories in financial information, we conjecture that the mechanism we identify here is also at play outside of the lab. In some situations, the contextual features in which financial news are embedded might serve as cues that trigger recall, for instance upward-sloping trend lines or pictures of bulls. In other situations, financial news may come in the form of narratives (e.g., “crypto is the future”, “climate risks are underpriced”), which then trigger the recall of similar narratives from the past. Relatedly, Charles (2022b) proposes earnings announcements as memory-relevant cues, whereas Jiang et al. (2023) emphasize the role of recent returns. More generally, various scholars have recently argued that the broad idea of associations-driven overreaction may be a driver of aggregate financial events. For example, in influential writings on the role of narratives, Shiller and co-authors appeal to the role of associative recall for expectation formation and overreaction in financial markets by observing that “[o]ne new narrative may remind of another that has been lying fairly dormant...there is cue-dependent forgetting” (Shiller, 2017, p. 975, also see Shiller, 2019, Goetzmann et al., 2022). Similarly, Gennaioli and Shleifer (2018) treatment asserts that associative memory may underlie overreaction to news in the context of the 2007–2008 financial crisis. More generally, overreaction has been argued to be a key pattern in the processing of financial and macroeconomic news (Bordalo et al., 2020a; Beutel and Weber, 2022), yet the sources that potentially underlie such overreaction are not well-understood.

Contribution and related literature. Overall, our contribution to the literature is (i) to provide evidence that associative recall shapes the formation of expectations; (ii) that this produces systematic overreaction; and (iii) that this affects experimental market prices, despite strong scope for self-selection. These results tie into a growing theory literature that has argued for the importance of associative memory for economics and finance (Mullainathan, 2002; Bordalo et al., 2020b, 2023; Wachter and Kahana, 2019; Bodoh-Creed, 2020). On the experimental side, there is a large psychology literature but scant evidence from economic or financial decision problems.³ Psychological experiments have at least two features that make them less-than-directly applicable to the questions that we are interested in here. First, these experiments are generally pure recall tasks that do not involve information-processing, belief updating or overreaction. Second, psychologists have not paid attention to how associative recall affects market behavior. In economics, a small number of lab experimental papers provide evidence on the role of memory imperfections for belief formation in economic or financial decision tasks (Afrouzi et al., 2023; Bordalo et al., 2023; Graeber et al., 2023). Our contribution to this line of work is to document how associative recall produces overreaction in expectations, and that these effects persist in markets with scope for self-selection.⁴

Some recent observational and survey studies also provide field evidence suggesting that associative memory is a driver of investment behavior (Charles, 2022a,b; Jiang et al., 2023). These studies provide complementary evidence to our causally identified experiments – by their very nature, they offer more ecological validity but do not afford

the possibility to directly and exogenously manipulate the presence of associations and, hence, offer less control in identifying memory effects.

Our paper also links to the voluminous literature on experience effects (e.g., Malmendier and Nagel, 2015; D'Acunzio et al., 2021). A plausible interpretation of this literature is that past experiences matter mostly when they get cued by current events, which is the mechanism that we emphasize here. In fact, Malmendier (2021) argues that “experiences in one setting (say, the stock market) affect beliefs and future risk-taking specifically in that setting (stock investment), but not necessarily in related settings, such as other asset markets (e.g. the bond market), even if the realizations of the underlying stochastic processes are correlated”, which highlights the potential role of associative memory for experience effects.

Finally, our experiments are related to the active literature that documents overreaction in survey expectations about financial and macroeconomic variables (e.g., Greenwood and Shleifer, 2014; Bordalo et al., 2020a; Barrero, 2022; Beutel and Weber, 2022). While overreaction is an oft-emphasized phenomenon in the finance literature, much other work has provided evidence for underreaction to news (e.g. Abarbanell and Bernard, 1992; Bouchaud et al., 2019; Bordalo et al., 2020a), including in laboratory experiments and surveys (Benjamin, 2018).

Recent work suggests that underreaction largely arises when the information people receive is relatively precise (Augenblick et al., 2023; Ba et al., 2023), in no small part because people who exhibit high cognitive uncertainty tend to form beliefs that exhibit regression towards a prior (Enke and Graeber, 2023). We suspect that the relative strength of under- and overreaction in any given context will depend on the relative importance of associative recall (more overreaction) and complexity-driven underreaction. For instance, in those laboratory experiments that document underreaction, memory imperfections are by design ruled out, hence leaving no role for memory-induced overreaction. On the other hand, we work with relatively imprecise and simple signals, leaving little room for complexity-based underreaction.

The remainder of the paper proceeds as follows. Section 2 offers a stylized formal framework that motivates the experimental design and structures the analysis. Section 3 describes the experimental design and pre-registration. Sections 4 and 5 present the results of the individual belief elicitation and market experiments. Section 6 concludes.

2. Theoretical framework

2.1. Setup

This section presents a stylized model to guide the design of the experiments and to structure the empirical analysis. The mechanics of the model build on some of the formulations in Mullainathan (2002) and Bordalo et al. (2020b, 2023). The framework rests on two key assumptions: (i) people may forget prior knowledge, so that they need to reconstruct it from memory; and (ii) this recollection process is subject to associative recall.

Consider a decision-maker (DM) who forms beliefs about the value of an asset θ with possible states denoted by $\theta = G(\text{ood})$ and $\theta = B(\text{ad})$. The prior probability is $P(G) = P(B) = 0.5$. The DM receives a series of i.i.d. binary signals s_x that take on the realizations $p(\text{ositive})$ and $n(\text{egative})$. The signal diagnosticity is given by $P(p|G) = P(n|B) = q > 0.5$. In what follows, we will use the terms “news” and “signal” interchangeably. With a slight abuse of notation, we will write $s_x = 1$ for positive and $s_x = -1$ for negative signals.

There are two periods. In the first (“past”), the DM receives potentially multiple signals, s_1, \dots, s_k . Denote by N_p and N_n the number of positive and negative first-period signals. In the second period (“present”), the DM receives one additional signal, s_{k+1} . We call a first-period signal s_x *congruent* with the second-period signal if $s_x = s_{k+1}$. It is helpful to introduce a shorthand for the number of first-period signals that are congruent with the second-period signal: $z := \sum_{x=1}^k \mathbb{1}_{s_x=s_{k+1}}$.

³ In a canonical psychological task on associative recall, subjects are asked to memorize words, and are subsequently more likely to remember a word if it was shown in conjunction with another word that is currently being displayed (see Schacter, 2008; Kahana, 2012, for overviews). As we review in Section 3.5 and Online Appendix A, our experimental design applies a variant of canonical psychological paradigms to financial decision tasks.

⁴ Other recent lab experimental work in behavioral finance on belief updating includes Hartzmark et al. (2021), Charles et al. (2024), Augenblick et al. (2023).

We assume that each signal s_x is experienced in a context, c_x . By context we mean all environmental features that co-occur with a signal, except the signal realization itself. Loosely speaking, contexts are characterized by two aspects. First, conditional on the signal, they are uncorrelated with the state, meaning that they are intrinsically uninformative. Second, contexts are memorable in the sense that observing them today reminds people of occurrences of the same environmental features in the past.

We introduce two different counterfactual conditions (that will correspond to experimental treatment conditions), across which the mapping between signals and contexts differs. First, in a *Cue* condition ($\mathbb{T} = 1$), there is a one-to-one mapping between type of news (positive or negative) and context: $c_x = c_y \Leftrightarrow s_x = s_y$. Thus, all positive news appear in the same context and all negative news in the same (yet different) context. Second, in a *NoCue* condition ($\mathbb{T} = 0$), the same context never appears twice, regardless of the signal realizations: $c_x \neq c_y \forall x, y$. We say that a first-period signal gets “cued” by a second-period signal when both the signals and the contexts are identical.

2.2. Memory and beliefs

First-period beliefs. Denote by $b_t(G|S_t)$ the DM’s posterior belief in period t that the state is good, following signal history S_t . By standard arguments, the first-period Bayesian posterior belief odds can be expressed as a function of the likelihood ratio and the prior odds:

$$\frac{b_1(G|S_1)}{1 - b_1(G|S_1)} = \left(\frac{q}{1 - q} \right)^{\sum_{x=1}^k s_x} \frac{p(G)}{p(B)} \quad (1)$$

where the likelihood ratio consists of the diagnosticity odds to the power of the number of positive minus negative signals. The prior odds drop out because we assumed $P(G) = P(B) = 0.5$. A popular transformation of this expression in the literature is the so-called (Grether, 1980) decomposition. Taking logs and re-arranging, we get a linear expression for the DM’s first-period *normalized log posterior odds* (*lpo*):

$$lpo_1 := \frac{\ln \left(\frac{b_1(G|S_1)}{1 - b_1(G|S_1)} \right)}{\ln \left(\frac{q}{1 - q} \right)} = \sum_{x=1}^k s_x = N_p - N_n \quad (2)$$

The normalized log posterior odds vary one-for-one with changes in the net number of positive signals. This property of Bayesian beliefs is well-understood.

Second-period beliefs: perfect memory benchmark. By a simple extension of the above, the normalized second-period Bayesian log posterior odds can be expressed as

$$lpo_2^{bayses} := \frac{\ln \left(\frac{b_2(G|S_2)}{1 - b_2(G|S_2)} \right)}{\ln \left(\frac{q}{1 - q} \right)} = s_{k+1} + (N_p - N_n) \quad (3)$$

This expression is analytically very convenient because (i) it can be estimated using simple OLS regressions and (ii) the perfect-memory benchmark coefficient of the second-period signal is simple and given by one.

Second-period beliefs: the case of associative recall. Now consider a DM who potentially forgets some or all signals going from $t = 1$ to $t = 2$. Whether or not the DM remembers a signal is determined by two factors. First, irrespective of the piece of news, there is some probability $r \in [0, 1]$ that the DM will remember. Second, reflecting the logic of associative recall, the probability of recalling a past signal is higher if its context is identical to today’s context. We assume that recall \hat{s}_x of s_x is given by

$$\hat{s}_x = \begin{cases} s_x & \text{with probability } r + (1 - r)a \mathbb{1}_{c_x=c_{k+1}} \\ 0 & \text{else} \end{cases} \quad (4)$$

Thus, the probability of remembering a first-period signal s_x is r whenever the context of the first-period signal does not match the context

of the second-period news, $c_x \neq c_{k+1}$. If, on the other hand, the context of the first-period signal equals the context of the second-period signal, the probability of recall receives an “associations boost” parameterized by $a \in (0, 1]$.⁵

Following Mullainathan (2002) and Bordalo et al. (2023), we assume that the DM applies Bayes’ rule to the signals she retrieves from memory. Thus, we derive the DM’s posterior odds following equation (1), except that we replace the actual signals, $\sum s_x$, with the recalled ones, $\sum \hat{s}_x$.

Using the Grether decomposition again and doing a bit of algebra delivers:

$$\begin{aligned} lpo_2^{assoc} &= s_{k+1} + \sum_{x=1}^k \hat{s}_x \\ &= s_{k+1} + \sum_{x=1}^k E[\hat{s}_x | s_x, s_{k+1}] + \underbrace{\sum_{x=1}^k (\hat{s}_x - E[\hat{s}_x | s_x, s_{k+1}])}_{:=\epsilon} \\ &= [1 + \underbrace{(1 - r)az\mathbb{T}}_{\text{Overreaction}}] s_{k+1} + r \underbrace{(N_p - N_n)}_{\text{1st-period lpo}} + \epsilon \end{aligned} \quad (5)$$

where the mean-zero noise term ϵ reflects that the memory technology in (4) is random. In this Grether decomposition, the second-period log posterior odds are expressed as a function of the second-period signal and the first-period log posterior odds. Here, beliefs *look like* they overreact to the second-period signal because the overall coefficient is potentially strictly larger than one. Intuitively, the second-period signal has both a direct effect on beliefs and an indirect effect through the asymmetric recall that it generates. Indeed, when the stable contexts are removed ($\mathbb{T} = 0$), Eq. (5) does not predict overreaction. Also observe that (5) clarifies that associative recall only distorts beliefs if $z > 0$. This is intuitive: if no first-period signal equals the second-period signal, then nothing gets cued and no asymmetric retrieval takes place.

Our experiments focus on testing the distinctive (comparative statistics) predictions that arise from Eq. (5) relative to Eq. (3). In particular, our experiments will exogenously manipulate the experimental analogues of the parameters \mathbb{T} and z .

Model predictions 1.

1. If a strictly positive number of first-period signals are congruent with the second-period signal ($z > 0$), overreaction of second-period beliefs to the second-period signal is larger in the presence of associations: $\frac{\partial lpo_2}{\partial s_{k+1}} |_{\mathbb{T}=1, z>0} > \frac{\partial lpo_2}{\partial s_{k+1}} |_{\mathbb{T}=0, z>0}$.
2. If no first-period signals are congruent with the second-period signal ($z = 0$), there is no differential overreaction across treatments: $\frac{\partial lpo_2}{\partial s_{k+1}} |_{\mathbb{T}=1, z=0} = \frac{\partial lpo_2}{\partial s_{k+1}} |_{\mathbb{T}=0, z=0}$.
3. In the presence of associations, overreaction increases in the number of congruent first-period signals, z , even holding fixed first-period beliefs: $\frac{\partial^2 lpo_2}{\partial s_{k+1} \partial z} |_{\mathbb{T}=1, lpo_1} > 0$.

In a nutshell, these predictions can be summarized with two themes. First, if at least one first-period signal “gets cued” by the second-period context, associative recall produces systematic overreaction of beliefs to the second-period signal, which makes beliefs too extreme, on average. Second, associative recall implies a distinctive form of history-dependence: even holding fixed first-period beliefs, the signal history matters for second-period beliefs.

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.

⁵ In a more general model, associativeness is formalized via a continuous similarity function (Bordalo et al., 2020b). Our formulation corresponds to a simplification in which similarity is either 0 or 1.




Fig. 1. Experimental Timeline. Notes. The order in which companies appear is randomly drawn, separately for each period.

3. Experimental design

Our experiment builds a bridge between the tightly-controlled and quantitative designs that dominate modern experimental economics and finance research on the one hand and psychological paradigms on cued recall problems on the other hand. We particularly focused on the following design objectives: (i) a decision setup that is closely tied to the model in Section 2; (ii) tight control over signal structure and associations; (iii) a financial decision environment that is intuitive for participants; (iv) exogenous variation in the key model parameters; and (v) incentive-compatible belief elicitation.

3.1. Experimental setup

Task overview. We implement a stylized financial decision-making task in which subjects need to estimate the probability that each of 14 hypothetical companies is of a good rather than a bad type. Our design follows binary-state balls-and-urns experiments, except that we add a memory component. Each company is denoted by a capital letter. The experiment consists of two periods (Fig. 1 provides a timeline of the experiment). In both periods, subjects receive noisy information about each of the companies, where first-period signals are also relevant for second-period beliefs. The information is embedded in memorable contexts that potentially facilitate associative recall. These contexts are intrinsically distinct from the signals itself, which allows us to identify the causal effect of associations above and beyond the normatively relevant informational content of the signals. This latter aspect is a clear advantage of a lab experiment relative to field contexts, where contexts and signals are often inseparably intertwined.

Signal structure. The objective type of each company, G (ood) or B (ad), is independently drawn according to $P(G) = P(B) = 0.5$. Subjects receive potentially multiple binary signals that can be positive, p , or negative, n . The signal diagnosticity is given by $P(p|G) = P(n|B) = 0.65$. In the experiment, good companies are represented by a box that comprises 65 positive and 35 negative news, while bad companies are represented by a box that comprises 35 positive and 65 negative news (see Online Appendix D for a picture). The computer draws at random from these boxes.

First period. In the first period, subjects complete the updating task for each of the 14 companies sequentially and in random order. For a given company f , subjects first observe k_f i.i.d signals on separate screens, with $k_f \in \{0, \dots, 4\}$.⁶ On a final screen directly thereafter, subjects state their first-period posterior belief about whether the company is good (0%–100%). The same procedure is repeated for all companies.

After the first period, we implement a time gap in which subjects work on an unrelated real effort task, which requires subjects to type multiple combinations of letters and numbers into the keyboard. Subjects have 8 min to type in as many combinations as they can. For each correctly solved task, they receive 5 cents.

Second period. In the second period, subjects are again tasked with stating probabilistic beliefs about whether each of the 14 companies is of a good or a bad type. The true state for each company is the same as in the first period, such that all first-period signals are still relevant in the second period.

For each company, subjects receive one additional signal and immediately after state their second-period posterior belief. This procedure is repeated for each company, in random order. The experimental instructions and comprehension checks emphasize that first- and second-period signals are equally relevant for second-period beliefs.

To summarize, as depicted in Fig. 1, the timeline of the experiment is as follows. Initially, subjects receive instructions and complete comprehension checks. This material covers both periods. In the first period of the updating task, a subject first receives all first-period signals for a company and immediately after states a first-period belief. Then, the subject receives all first-period signals for the next company and states a first-period belief. This process is repeated for all 14 companies, after which an 8-minutes real effort task follows. Then, the subject receives a second-period signal for a company and immediately after states a second-period belief. This procedure is then again repeated for all 14 companies.

Communication of news and contexts. Signals are communicated on subjects' computer screens, one per screen. The signal itself is communicated as "The news for company [X] is positive [negative]". In addition, this signal is embedded in an intrinsically uninformative context. In our experiment, we implement these contexts as events that explain the occurrence of positive or negative news. We chose this implementation of contexts because it is arguably intuitive for participants: all that happens is that they do not just receive an abstract piece of information about whether the company is good or bad, but that the computer also explains to them why the news are positive or negative. Our treatment variation (to be explained below) manipulates how events are linked to signals, as captured by the parameter \mathbb{T} in the model.

All events are represented by a story and an image. For example, a positive signal may be shown along with a story about a successful hire and a picture of the new employee. Another example is a positive signal that is communicated to subjects with a short story about a successful advertising campaign with a celebrity, along with a picture of the celebrity. All stories were constructed to be of similar length and structure. See Online Appendix Figures 7 and 8 for examples.

The written instructions clarify to subjects that the images and stories have no purpose other than to provide a rationale for the positive or negative news. Conditional on the signal ("positive news" or "negative news"), they are uninformative about the true state of a company. The signal, image and story are displayed on subjects' computer screens for 15 s. The time was calibrated such that subjects had sufficient time to process the news, as well as to fully grasp the content of the image and the story.

3.2. Sources of exogenous variation

Cue and NoCue companies. To exogenously manipulate the presence of associations (\mathbb{T} in the model), our design employs a within-subjects treatment variation. For each subject, seven companies are assigned to be in the *Cue* condition, while the remaining seven companies are assigned to the *NoCue* condition. To counterbalance potential differences in news events across companies, the treatment assignment

⁶ For each subject, $k = 0, 1, 3$ for two companies each and $k = 2, 4$ for four companies each.

of companies was randomized across subjects, such that any given company was a *Cue* company for some and a *NoCue* company for other subjects.

The only difference between *Cue* and *NoCue* companies is the mapping between events (contexts) and signals. In *Cue*, every positive signal for a given company is communicated with the same story and image, and every negative signal for a given company is communicated with the same (but different) story and image. For example, if a subject received three positive and one negative signals for a company, then all three positive signals would be communicated with the same story and image, and the negative signal with a different story and image. Thus, for these companies, the second-period signal potentially triggers associative recall of congruent first-period signals through the identical contexts.

For *NoCue* companies, on the other hand, each piece of news is communicated with a unique context. Any given image and story never appear twice, even if the company and type of news are identical. Continuing the above example, if a subject received three positive and one negative signals for a company, then each signal would be communicated with a different story and image.

Subjects did not know *ex ante* which company was in the *Cue* or *NoCue* condition. In fact, subjects did not even know they were being subjected to a within-subjects treatment variation. Instead, we simply instructed them that the events that generate positive and negative news can potentially occur multiple times. For instance, the negative event that a factory burns down can occur multiple times and cause multiple negative news that should all (independently) be taken into account. The instructions emphasized that while the same event can occur multiple times, it can only occur for the same company. Likewise, we emphasized that the same event can only be associated with positive or negative news but never with both. Thus, subjects knew that if the second-period signal for company A is negative because a factory burned down, and the subject remembers having read this story before, then they know that they must have received at least one negative signal about company A in the first period.

We further instructed subjects to treat each piece of news as independent and in an identical fashion, regardless of which events are associated with these news. For instance, we emphasized that the following two signal histories are equally informative: (i) three positive news about a company, all of which are triggered by the same event and (ii) three positive news about a company, each of which is triggered by a different event. We verified subjects' understanding of the intrinsic irrelevance of whether the same event occurs repeatedly through comprehension questions (see Online Appendix D.3.2).

In summary, a within-subjects treatment design is particularly natural in our context because the entire treatment variation boils down to whether, for a given company, a subject receives multiple signals that are triggered by the same event or by multiple different events. The order of companies was randomized at the subject level, such that subjects (unknowingly) repeatedly alternated between *Cue* and *NoCue* companies. Thus, potential order or contrast effects – sometimes a concern in within-subjects designs – are implausible in our context.

Signal histories. On top of the within-subjects-across-company variation in the presence of associations, we also causally identify the role of associative recall by exogenously varying the number (and realizations) of first-period signals at the subject-company level. We leverage this source of exogenous variation to test the predictions derived in Section 2 about how the presence or magnitude of overreaction depends on the number of congruent first-period signals. This layer of randomization is directly built into the design because (i) the number of first-period signals for each company randomly varies between one and four, and (ii) conditional on the number of signals, both first- and second-period signals are randomly generated.

Interpretation of treatment comparison. Our *Cue* condition is admittedly extreme in the sense that signals and contexts are perfectly correlated. We chose this implementation to keep the experimental design as simple and transparent as possible. While in reality people likely do not repeatedly experience the same signals in *exactly* the same context, the *Cue* condition is arguably more reflective of reality than the *NoCue* condition. This is because in many contexts similar signals will be associated with similar contexts. For example, whenever good news prevail in the stock market, people are disproportionately exposed to bulls, upward-sloping trend lines, and good-times stories (see, e.g., Shiller, 2019).

3.3. Incentives

Subjects stated their beliefs about whether a company is good vs. bad using a slider (0%–100%). Beliefs were incentivized using a binarized scoring rule (Hossain and Okui, 2013). Under this scoring rule, subjects could potentially earn a prize of 10 euros. While these stakes are substantially smaller than those present in real financial markets, recent experimental work using very large incentives finds that the presence of belief updating errors is often robust to the stake size employed (Enke et al., 2021a).

The probability of receiving the prize is given by $p = 1 - (b - t)^2$, where b is the belief that a company is good and t the truth.⁷ In order to avoid hedging motives, at the end of the experiment one of the 28 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.

3.4. Serial independence of signals

A key element of the theoretical framework in Section 2 and our experimental design is that signals are conditionally independent. Under serial dependence (positive autocorrelation), it would be “rational” for subjects who forgot the first-period signals to “overreact” to the second-period signal even without any associative recall, simply because they would rationally infer from a positive second-period signal that the first-period signals were likely positive. This would be a potential concern for our design (only) if subjects assumed a higher level of autocorrelation for the *Cue* than the *NoCue* companies.

To address such concerns, we took two steps. First, the instructions used intuitive language to emphasize that the signals are serially (conditionally) independent. We augmented these explanations with a comprehension check question that specifically asked subjects whether a positive signal becomes more likely after a positive signal was drawn.

Second, an account of overreaction that is based on assumed autocorrelation does not generate the additional prediction that overreaction depends in nuanced ways on the signal history. This is because assumed autocorrelation predicts that subjects always infer from a positive second-period signal that the first-period signals were likely also positive, irrespective of the actual realizations of the first-period signals. In contrast, our model predicts that overreaction depends on the company-specific *random realizations* of the first-period signals. It seems implausible that subjects mentally impute (and remember) different degrees of autocorrelation for each company based on the first-period signals, especially given how salient our instructions are about the absence of autocorrelation.

⁷ Danz et al. (2022) provide evidence that the binarized scoring rule can lead to a tendency to state less extreme beliefs. Even if such bias was present in our experiment, it would not confound our causal identification, which holds the belief elicitation constant between treatments. If anything, it would lead to an under-estimation of the effect of associative recall on overreaction.

3.5. Relationship to psychology paradigms

Our experimental design builds on the main ideas of a well-known paradigm in memory research, namely lists of word-pairs (Kahana, 2012). Subjects first sequentially observe word pairs, consisting of a “target” and a “cue”. At a later stage, subjects’ recall of target words is greater when they are provided with the cue word during recall elicitation (Tulving and Thomson, 1973). The analogy to our experimental design is that the signal is the target and the context serves as cue.

The technique we use to generate partial “forgetting” of first-period signals is a variant of the word-pairs paradigm that is called AB/AC in the psychology literature (see chapters 4–5 in Kahana, 2012). Subjects first memorize word pairs (“A” and “B”). Then, in a second step, they memorize new word pairs, some of which involve one of the words from the first set (“A” and “C”). The main finding is that recall of the A-B pair is significantly impaired after subjects learn the A-C pair. This is commonly referred to as “interference”. Building on this paradigm, our experimental design creates partial forgetting of first-period signals through: (i) a time lag (distraction task) and (ii) interference that results from the presence of 14 companies with identical news (“positive” and “negative”). See Online Appendix A for a more detailed discussion of the psychology literature on associative recall.

3.6. Econometric specifications and predictions

Following the theoretical framework in Section 2, for most analyses we transform subjects’ raw beliefs into normalized log posterior odds. Eq. (5) directly suggests the following estimating equation for a potential treatment difference in subject i ’s normalized second-period log posterior odds about whether company f is good:

$$lpo^{i,f} = \beta_1 s_{k+1}^{i,f} + \beta_2 s_{k+1}^{i,f} T^{i,f} + \beta_3 T^{i,f} + \beta_4 (N_p^{i,f} - N_n^{i,f}) + \epsilon^{i,f} \quad (6)$$

where $T^{i,f}$ is a binary treatment indicator that equals one if, for a given subject, company f was in the *Cue* condition. In words, we regress a subject’s normalized log posterior odds on the second-period signal, a treatment indicator, their interaction and the net number of positive first-period signals. We predict that the interaction effect is positive, $\beta_2 > 0$, and that this positive interaction effect is only driven by cases when congruent first-period signals were observed, that is, with $z > 0$.

Furthermore, within the set of *Cue* companies, we test how overreaction depends on the number of cued first-period signals:

$$lpo^{i,f} = \beta_5 s_{k+1}^{i,f} + \beta_6 s_{k+1}^{i,f} z^{i,f} + \beta_7 z^{i,f} + \beta_8 (N_p^{i,f} - N_n^{i,f}) + \epsilon^{i,f} \quad (7)$$

Here, we predict a form of history-dependence, which is that the interaction between the second-period signal and the number of cued first-period signals is positive: $\beta_6 > 0$.

3.7. Procedures and logistics

In total, we implemented three experiments, all of which were conducted at the same time and randomized within experimental sessions. The first experiment is the one described above, which we refer to as *Beliefs* experiment. In addition, we also implemented a *Recall* experiment (to directly elicit which first-period signals subjects remember in the second period) and a *Market* experiment (to study market behavior based on associative recall). We discuss these additional experiments in Sections 5 and 4.4, respectively.

All three experiments were conducted as Zoom online experiments based on the subject pool of the BonnEconLab of the University of Bonn. The experiments were computerized using Qualtrics and lasted up to 90 min. Subjects met with an experimenter via Zoom and received a participation link to the experimental software via Zoom chat. Subjects were told not to use any material (such as pen and paper) during the experiment. Online Appendix D contains the full set of instructions, translated into English. Subjects were given unlimited time to read the

instructions and could ask questions at any point in time using the Zoom chat.

After subjects finished the instructions, they completed computerized comprehension check questions, see Online Appendix D. 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, we exclude subjects that answered more than one comprehension check question incorrectly (5% of potential participants). As we pre-specified, 100 valid completes were collected for the *Beliefs* experiment. Average earnings were 16.50 euros, which includes a participation payment of 10 euros.

All experiments in this paper were pre-registered in the AEA RCT registry, see <https://www.socialscienceregistry.org/trials/9215>. The pre-registration includes the design of all experiments reported in this paper, predictions, sample sizes, and that subjects would be dropped from the sample (and replaced) if they answer more than one comprehension check question incorrectly.

4. Results

4.1. Preliminaries: First-period beliefs

Model equation (2) posits that the normalized log posterior odds in the first period move one-for-one with variation in the first-period signals. Online Appendix Table C1 shows that, in our data, this is indeed the case, for both *Cue* and *NoCue* companies. While our treatment comparisons of second-period beliefs do not hinge on first-period beliefs being close to Bayesian, this piece of information is helpful because it shows that it is largely irrelevant whether our second-period regressions control for first-period log posterior odds or the number of positive minus negative first-period signals (see Eq. (5)).

Our treatment comparison of second-period beliefs is, however, only valid if first-period beliefs do not differ from each other across treatments in a way that would spuriously generate a treatment difference also in second-period beliefs. While the experimental design offers no ex-ante reason for why first-period beliefs should differ across *Cue* and *NoCue* companies, Online Appendix Table C2 formally tests this. Reassuringly, we find that the difference in first-period beliefs across treatments is very small and statistically insignificant.

4.2. Second-period beliefs: A look at the raw data

As derived in Section 2, an immediate implication of associative recall is that second-period beliefs are on average more extreme (further away from 50%) in the presence of associations. Fig. 2 provides a first test of this, by showing kernel density plots of the distribution of second-period beliefs, separately for *Cue* and *NoCue* companies. The left panel shows beliefs following a negative second-period signal, while the right panel shows beliefs following a positive second-period signal. Recall that for each subject-company combination the signal realizations in the first and second period were randomly generated. The belief heterogeneity in Fig. 2 therefore captures a combination of (i) variation in first-period signal realizations and (ii) variation in beliefs across subjects conditional on the same signal realizations.

We see that beliefs in *Cue* are substantially more extreme, following both a positive and a negative second-period signal. While we conduct more sophisticated regression analyses below, we note that this treatment difference in average beliefs is statistically highly significant in both panels, see Online Appendix Table C2.




Fig. 2. Kernel density estimates of second-period beliefs as a function of treatment and second-period signal. The horizontal green line indicates the average Bayesian posterior across all belief formation problems. Kernel is Epanechnikov. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

4.3. Econometric analysis

A main advantage of visualizing the data as in Fig. 2 is that the analysis is very transparent as it does not require any transformations of the raw data. A disadvantage of working with the raw beliefs data, however, is that it does not allow for quantitative analyses of overreaction in which empirical results can be compared against normative benchmarks. The reason is that over- vs. underreaction is typically defined through the Grether (1980) regressions that our theoretical framework also directly motivates.⁸

As shown in Eq. (6), these Grether regressions relate the (normalized) log posterior odds to the second-period signal, controlling for the number of positive minus negative signals in the first period. Fig. 3 visualizes the results of these OLS regressions by displaying the coefficient of the second-period signal. Recall that the second-period posterior log odds are transformed such that the Bayesian benchmark coefficient is one. We conduct this analysis separately by condition and by looking at random variation in the number of first-period signals that are congruent with the second-period signal. The figure shows point estimates along with 95% confidence intervals.

As predicted by the theoretical framework, we observe three patterns. First, when no first-period signal is congruent with the second-period signal ($z = 0$), subjects state identical beliefs across treatments.⁹ Second, for any strictly positive number of congruent first-period signals ($z > 0$), the effect of the second-period signal is significantly larger in *Cue* than in *NoCue*. This documents that associations generate overreaction.

Third, looking within treatment *Cue*, the effect of the second-period signal monotonically increases in the number of congruent signals. Note that – consistent with the results from the *NoCue* condition – in the absence of associative recall the regression coefficient should not at all




Fig. 3. Effect of second-period signal on (normalized) second-period log posterior odds, as a function of the number of congruent first-period signals in *Cue* and *NoCue*. The point estimates are derived from the OLS regression equation (6), which is run separately for each value of z . The figure plots β_1 for *NoCue* and $\beta_1 + \beta_3$ for *Cue*. The figure does not include $z = 4$ because there are very few observations with such a signal history. All regression analyses that are reported in tables include these cases. Whiskers show 95% confidence intervals, computed based on clustering at the subject level.

depend on the first-period signal history. Thus, this third result provides evidence for a form of history dependence that is absent in standard economic models of belief formation without associative recall.

Table 1 provides the regression results. The table notes contain detailed explanations about the construction of each variable, where the construction always follows the logic from the model in Section 2. The regressions again directly correspond to the estimating Eqs. (6) and (7) that our model motivates. We construct the table such that the most relevant independent variables are listed at the top. First, columns (1)–(3) focus on across-treatment differences. Column (1) shows that, in the full sample of 1400 second-period beliefs (100 subjects, 14 companies each), the effect of the second-period signal is indeed significantly larger in *Cue* than in *NoCue*. Columns (2) and (3) decompose this treatment difference into cases with $z = 0$ and $z > 0$, where the model only predicts a treatment difference in the latter case. Consistent with the visual impression from Fig. 3, this is indeed what the regressions show. In terms of quantitative magnitude, column (3) shows that in the theoretically-relevant case with $z > 0$, the effect of the second-period signal is more than 80% larger for the *Cue* companies.

⁸ As is well-known in the literature, a slight challenge in directly estimating Grether regressions on real data is that people occasionally state beliefs of 0% or 100%, which makes them undefined under the log odds transformation. In our data, this is the case for 93 second-period beliefs (6.6% of all data). To avoid a loss of observations, we recode observations of 0% as 1% and 100% as 99%. Online Appendix Table C3 shows that our results are quantitatively virtually identical if we do not replace these observations but instead lose them through the log odds transformation.

⁹ Note that the coefficients in these cases do not differ much from one which suggests that subjects weight the signals on average similar to the Bayesian benchmark. This finding is consistent with Augenblick et al. (2023), who also find updating similar to the Bayesian benchmark for a signal precision of 0.65, which is the signal precision we implemented in our experiment.

Table 1
Overreaction in *Cue* and *NoCue*.

Sample:	Dependent variable: 2nd period normalized log posterior odds						
	<i>Cue</i> vs. <i>NoCue</i>			<i>Cue</i>		<i>Cue</i> vs. <i>NoCue</i>	
	Full	$z = 0$	$z > 0$	Full	Full	Full	$k > z > 0$
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$t = 2$ signal	1.20*** (0.14)	1.20*** (0.15)	1.13*** (0.15)	1.42*** (0.18)	1.40*** (0.14)	1.29*** (0.16)	1.38*** (0.27)
$t = 2$ signal $\times 1$ if <i>Cue</i>	0.61*** (0.15)	−0.017 (0.17)	0.88*** (0.19)				
$t = 2$ signal $\times \#$ of congruent $t=1$ signals				0.32** (0.13)	0.35*** (0.10)		
Sum of congruent $t = 1$ signals $\times 1$ if <i>Cue</i>						0.50*** (0.11)	0.64*** (0.15)
Sum of incongruent $t = 1$ signals $\times 1$ if <i>Cue</i>						0.086 (0.08)	0.19 (0.12)
1 if <i>Cue</i>	−0.097 (0.12)	0.0093 (0.14)	−0.16 (0.12)			−0.10 (0.11)	−0.14 (0.14)
Sum of $t = 1$ signals (pos. minus neg.)	0.42*** (0.05)	0.26*** (0.09)	0.41*** (0.06)	0.40*** (0.08)			
# of congruent $t = 1$ signals				−0.074 (0.06)	−0.068 (0.05)		
$t = 1$ normalized log posterior odds					0.38*** (0.07)		
Sum of congruent $t = 1$ signals						0.26*** (0.07)	0.27** (0.13)
Sum of incongruent $t = 1$ signals						0.28*** (0.09)	0.35*** (0.11)
Observations	1400	418	982	700	700	1400	630
Adjusted R^2	0.41	0.20	0.47	0.49	0.54	0.42	0.46

Notes. OLS estimates, robust standard errors (in parentheses) are clustered at the subject level. Following the estimating equations (6) and (7) that we derived from the model, the estimations do not include a constant. However, we have verified that including a constant delivers almost identical results. Both the first- and the second-period log posterior odds are normalized by the log diagnosticity odds as described by Eq. (2). In column (7), the sample is restricted to signal histories where $k > z > 0$, i.e., (i) with at least one congruent first-period signal and (ii) at least one incongruent first-period signal. Variable labels: “ $t = 2$ signal” equals 1 if signal positive and (−1) if negative. “# of (in)congruent $t = 1$ signals” captures the number of 1st period signals that do (do not) equal the second-period signal. “Sum of (in)congruent $t = 1$ signals” captures the number of positive minus negative 1st period signals that are (in)congruent with the 2nd period signal. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Columns (4) and (5) report regression results that only leverage variation in the number of congruent signals within the *Cue* condition. Here, the coefficient of interest is the interaction between the second-period signal and the number of congruent first-period signals. Consistent with what we saw in Fig. 3, both regression specifications show that the effect of the second-period signal is significantly stronger when there are more congruent signals. This result on history-dependence of beliefs holds both when we control for the first-period signals (column (4)) and when we directly control for the subject's normalized first-period log posterior odds (column (5)).¹⁰

Result 1. Overreaction in beliefs is significantly larger in *Cue* than in *NoCue*. This treatment difference only exists when the number of congruent first-period signals is strictly positive.

Result 2. Within condition *Cue*, overreaction increases significantly in the number of congruent first-period signals.

¹⁰ Our model posits that what drives the magnitude of overreaction is the number of congruent first-period signals, irrespective of the specific order in which signals were received. For example, from the perspective of the model, signal histories of pos-pos-neg and neg-pos-pos are identical. Online Appendix Table C4 provides a tentative analysis of this issue. While these analyses generally suffer from very low power (because there is a large number of distinct possible signal histories), the results are indicative that the order of signals indeed does not affect the magnitude of overreaction.

4.4. Mechanism: Asymmetric recall of Cued signals

The model in Section 2 posits that the overreaction of beliefs reflects that subjects in *Cue* asymmetrically remember those first-period signals that are congruent with the second-period signal. In this subsection, we provide two pieces of causal evidence that the treatment difference between *Cue* and *NoCue* indeed reflects the asymmetric retrieval of specific signals that do/do not get cued by the second-period contexts.

Differential responsiveness of beliefs to congruent and incongruent signals. Re-consider the model in Section 2. Because our model of asymmetric recall focuses on whether or not a first-period signal is congruent with the second-period signal, it is useful to define by $N^z := z s_{k+1}$ the sum of congruent first-period signals and by $N^u := -(k - z) s_{k+1}$ the sum of incongruent first-period signals. Note that one of these quantities is positive, while the other is negative. To take a simple example, suppose that for a given company a subject observed three positive and one negative first-period signals and then a negative second-period signal. In this case, the sum of congruent first-period signals is (−1) and the incongruent sum is three.

Using this notation, the main model equation (5) can equivalently be expressed as

$$lpo_2^{assoc} = s_{k+1} + \sum_{x=1}^k \hat{s}_x = s_{k+1} + [r + (1-r)a\mathbb{I}]N^z + rN^u + \epsilon \quad (8)$$

This alternative expression for the Grether decomposition is helpful because – unlike the regressions reported above – it includes the sum of congruent and the sum of incongruent first-period signals as separate regressors. Crucially, the straightforward implication of associative

recall in our model is that the associations boost is asymmetric and only applies to the congruent-first period signals rather than all first-period signals. This implication of our model is distinctively different from a potential alternative model, according to which a second-period context (if it was also experienced in the first period) cues an improved recall of *all* first-period signals for a company, regardless of whether they are congruent or incongruent.

To test this, column (6) of Table 1 reports the results of a regression that interacts the sums of congruent and incongruent first-period signals separately with a treatment indicator. We find that the interaction effect of our treatment dummy ($= 1$ if *Cue*) with the sum of congruent first-period signals is quantitatively very large (twice as large as the raw coefficient of the congruent signals) and statistically highly significant. Meanwhile, the interaction effect with the sum of incongruent first-period signals is statistically insignificant and the point estimate close to zero. This shows that subjects' beliefs about *Cue* companies overreact to congruent first-period signals, as predicted by Eq. (8). Column (7) of Table 1 shows this result also holds when we restrict attention to cases where at least one first-period signal was congruent and at least one incongruent with the second-period signal ($k > z > 0$). This is a useful robustness check because in these cases the second-period signal definitely acts as a cue for either positive or negative signals (in column (6) we also include cases where potentially no first-period signal gets cued, such that treatment *Cue* cannot have an effect).

In summary, the results from columns (6) and (7) clarify that the entire treatment difference between *Cue* and *NoCue* is driven by an asymmetric responsiveness to congruent first-period signals, rather than an improved overall responsiveness to first-period signals in *Cue*.

Differential recall of congruent and incongruent signals. A second test of our mechanism is to directly gather data on which first-period signals subjects remember. According to the model, subjects should remember more congruent (but not incongruent) signals in *Cue* than in *NoCue*.

To test this, we implemented experiment *Recall*, which again randomized companies to treatments *Recall Cue* and *Recall NoCue* within subject. This experiment was randomized within experimental sessions with the *Beliefs* experiments described above. The experiment was identical to the *Beliefs* experiment, except that after receiving a second-period signal, subjects were asked to directly report the number of positive and negative signals they recall for a company. Subjects answered 28 such questions (recall of positive and negative signals for 14 companies each). For a randomly selected recall question, subjects received 10 euros if their answer was within ± 1 of the truth. Except for the recall component, the experiments and the underlying instructions were identical to those in *Beliefs*. To maximize similarity with the *Beliefs* experiment, the initial instructions only explained the belief elicitation task, and that first-period signals would also be relevant for second-period beliefs. Then, after subjects had concluded the first period as well as the distraction task, the recall task was announced as a surprise. As in the *Beliefs* experiment, subjects received one additional signal for each company and immediately after indicated their recall of positive and negative signals. Online Appendix D provides the experimental instructions. As we pre-registered, 70 subjects participated in this experiment. Average earnings were 18.50 euros, which includes a participation payment of 10 euros.

Fig. 4 summarizes the results by reporting subjects' effective recall of first-period signals as a function of the truth, separately for congruent first-period signals (left panel) and incongruent first-period signals (right panel).¹¹ The left panel shows a large and statistically highly

¹¹ In our experiments, we elicited subjects' *total* recall of signals in the entire experiment, including of those in the second period. For example, in cases in which we elicit recall of positive signals and the subject observed a positive second-period signal, effective recall of first-period signals is given by the reported recall minus one. This corresponds to the arguably very plausible assumption that subjects do not forget the second-period signal that they saw a few seconds ago on the previous screen.

significant treatment difference in the recall of congruent signals: subjects remember substantially more congruent first-period signals for *Cue* companies than for *NoCue* ones. In contrast, the right panel shows that for incongruent first-period signals (those that differ in signal type from the second-period one), this treatment difference is much smaller and not statistically significant. Indeed, Online Appendix Table C5 shows that the relevant difference-in-difference effect (treatment condition times congruent/incongruent signals) is statistically highly significant. These results again show that the associations that are present in *Cue* primarily induce *asymmetric* recall of congruent first-period signals rather than improved recall in general.¹²

Result 3. *Overreaction in second-period beliefs is driven by asymmetric recall of those first-period signals that get cued by the second-period context.*

Because our beliefs experiment and our recall experiment were conducted with different sets of participants, we cannot investigate whether, across individuals, the magnitude of associate recall and the magnitude of overreaction line up. However, using our model – in particular the memory technology in Eq. (4) – we can use a back-of-the-envelope calculation to triangulate between average beliefs and average recall. First, in the beliefs data, notice from Eqs. (5) and (6) that the regression coefficient of the interaction between a treatment dummy and the second-period signal (β_2 in eq. (6)), identifies $(1 - r)az^{i,f}$. Thus, multiplying $s_{k+1}^{i,f}$ in Eq. (6) by $z^{i,f}$ and then estimating the regression allows us to back out the estimated “associations boost” $(1 - p^{beliefs})\hat{a}^{beliefs} = 0.47$.

In the recall data, observe from Eq. (4) that, if the second-period signal equals the first-period one, the “associations boost” can be estimated by regressing reported effective recall (of positive or negative first-period signals) on the corresponding true number and its interaction with a treatment dummy. Here, under the model, the interaction coefficient reveals $(1 - r)a$, see Eq. (4). We estimate this quantity as $(1 - p^{recall})\hat{a}^{recall} = 0.37$, roughly in the same ballpark as the estimates based on the beliefs data.

5. Betting market experiment

5.1. Design

The basic structure of the *Market* experiment is identical to the *Beliefs* experiment, except that the belief elicitation task is embedded in a parimutuel betting market. While the canonical application of parimutuel markets is horse race betting, there are direct analogies to financial markets, where betters bet on mutually exclusive states of the world, such as whether an asset will increase or decrease in value. Indeed, parimutuel betting markets are frequently implemented in laboratory experiments because of their simplicity and appealing resemblance of real-world markets (e.g., Plott et al., 2003; Kendall and Oprea, 2018; Enke et al., 2023).

Parimutuel betting and payoffs. In our implementation, subjects are again asked to state probabilistic beliefs about whether each of 14 hypothetical companies is of a good type, after receiving a series of binary signals. The prior probabilities and signal structure are identical to those in the *Beliefs* experiment. Subjects are matched into groups of three and know that all participants in their market group receive the same public signals. In both part 1 and part 2 of the experiment, after observing signals as in the *Beliefs* experiment, each market participant privately states their subjective percent chance that the company is good. Immediately after, in both parts of the experiment, the three

¹² Online Appendix Figure B3 shows that we replicate this recall pattern also when we restrict attention to cases in which at least one first-period signal is congruent with and at least one incongruent with the second-period signal, analogous to column (7) of Table 1.

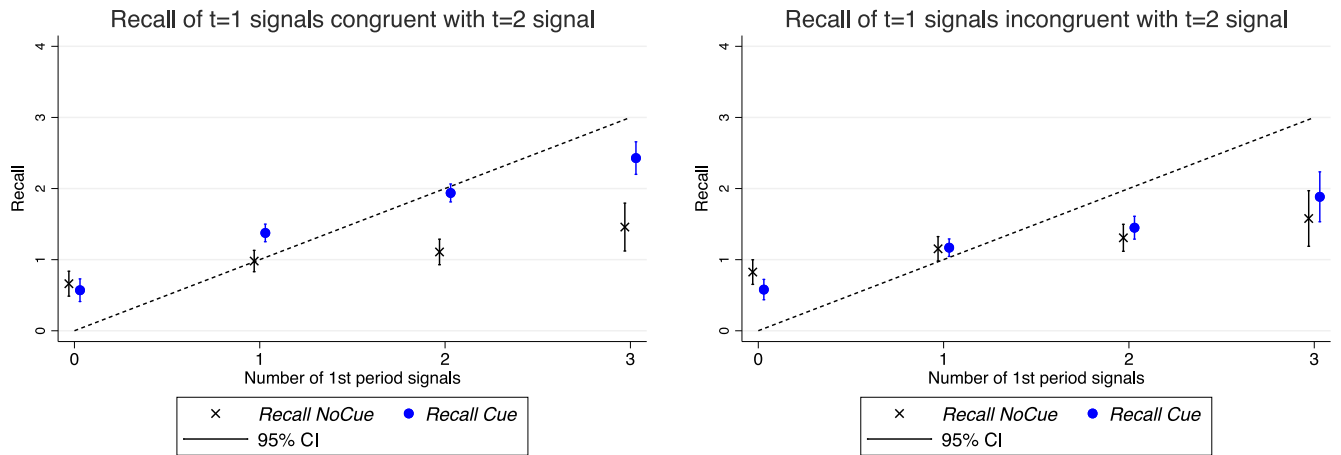


Fig. 4. Effective recall of congruent (left panel) and incongruent (right panel) first-period signals in *Recall Cue* and *Recall NoCue*. Effective recall equals reported recall for signals that differ from the second-period signal, and reported recall minus one for signals that equal it. The point estimates stem from an OLS regression of effective recall on a treatment dummy. The figures plot the coefficient of the constant for *NoCue* and the sum of the coefficients of the constant and the treatment dummy for *Cue*. The figure does not include the case of four first-period signals because there are very few observations with such a signal history. Whiskers show 95% confidence intervals, computed based on clustering at the subject level.

subjects interact in a parimutuel betting market. For each of their 28 betting decisions (14 companies and two periods each), subjects receive a budget of 10 euros. This money can be fully or partly bet on one or both of two propositions: that the company is good and that it is bad. The bets are implemented in two steps (on the same decision screen):

1. *Betting amount*: Subjects state the total amount they want to bet (maximum 10 euros, minimum 0 euros). We denote this amount by m .
2. *Betting proportion*: Subjects state which fraction (denoted by w) of m they bet on the proposition that the company is good. Thus, the total amount bet on the event that the company is good is $w m$ and the total amount bet on the event that the company is bad is $(1 - w)m$.

With subjects $i = 1, 2, 3$, the parimutuel market price in period t in group g for the asset that company f is good is defined as¹³:

$$\omega_t^{g,f} = \frac{\sum_{i=1}^3 w_t^{i,g,f} m_t^{i,g,f}}{\sum_{i=1}^3 m_t^{i,g,f}} \in [0, 1] \quad (9)$$

This has a simple interpretation, according to which the price is given by a weighted average of the betting proportions, where the weights are given by the betting amounts. Thus, the parimutuel price for an asset increases in the fraction of the total money in the market that is bet on the respective state. To intuitively relate this market price back to subjects' beliefs, consider the hypothetical scenario that the betting proportion for the good state is directly given by each subject's belief that the company is good (under expected utility, this will be the case when utility is $u(x) = \ln(x)$, see [Wolfers and Zitzewitz, 2006](#)). Under this scenario, the market price would be given by the weighted average belief, where the weights are given by how much each subject bets.

The payoffs of subject i in period t in market group g for company f are given by

$$\pi_t^{i,g,f} = (10 - m_t^{i,g,f}) + \frac{w_t^{i,g,f} m_t^{i,g,f}}{\omega_t^{g,f}} \mathbb{1}_{\theta=G} + \frac{(1 - w_t^{i,g,f}) m_t^{i,g,f}}{(1 - \omega_t^{g,f})} \mathbb{1}_{\theta=B} \quad (10)$$

In words, the subject keeps the part of the endowment that is not bet. In addition, the subject loses all money bet on the wrong state. Money bet on the right state yields a positive return whose magnitude

depends on the market price (9). Intuitively, the subject earns more money the more the subject bets on the right state and the more other subjects bet on the wrong state. The parimutuel price mechanism in our implementation fully redistributes all money that is bet; there is no efficiency loss or transaction cost.

Treatments and randomization. Analogously to the *Beliefs* experiment, we conduct two treatment conditions (within-subject-across-companies) that exogenously manipulate the presence of associations. For companies in *Market Cue*, each positive/negative signal is again associated with the same context/event. In contrast, for companies in *Market NoCue*, each signal is communicated with a different event. As in the individual belief elicitation treatments, for each subject the computer randomly selected seven companies to be in *Cue* and seven to be in *NoCue*. In addition to the across-treatment variation in the relevance of associative recall, the experiment again features random variation in the number of cued first-period signals (z).

Logistics and payoffs. As in the *Beliefs* experiment, there are a total of 14 hypothetical companies. The number of signals subjects see in part 1, the signal realizations as well as the order in which subjects see the companies is fully randomized across market groups. To avoid hedging, at the end of the experiment, for each subject, one of the two parts of the experiment and one company are randomly selected to be payout-relevant. For the randomly-selected company and part, either the subject's belief or the betting decision are randomly chosen and implemented for payment.¹⁴

The *Market* experiment was randomized within experimental sessions with the *Beliefs* experiment. The procedures and the subject pool (BonnEconLab, conducted over Zoom) were identical. The sample size was 240 subjects (80 groups). Average earnings were 19.80 euros, which includes a participation payment of 10 euros. Subjects remained in the same group throughout the experiment. No feedback was provided at any point. The *Market* experiment was also part of the pre-registration mentioned previously, including the experimental design, predictions and sample size.

5.2. Predictions

As is well-known, depending on assumptions on utility functions, betting market prices need not necessarily reflect average beliefs in

¹³ When none of the subjects in a market group bets money in a given round, the market price is missing. This occurs in only two out of 1120 second-period observations.

¹⁴ Individual beliefs were incentivized with the same binarized scoring rule as in Section 3.

the market.¹⁵ Our main interest here, however, is not in understanding how exactly beliefs aggregate to market prices but, instead, in the comparative statics effects regarding the role of associative recall: whether second-period market prices react more strongly to second-period signals in the presence of associations, and how this “overreaction” varies as a function of the signal history. Accordingly, we here only heuristically discuss our pre-registered predictions for the experiment.

To begin, consider again the definition of the parimutuel market price in Eq. (9). Note that subjects’ betting proportion on the proposition that a company is good (w) will usually increase in their belief that the company is good. Then, if associative recall generates overreaction in beliefs, we might also expect to see overreaction in market prices. To take a particularly simple example, suppose again that each subject’s betting proportion is directly given by his/her belief. Then, the market price is given by the weighted average belief in the market, and the degree of overreaction in market prices will be given by the average amount of overreaction in beliefs, weighted by each subject’s betting amount.

At the same time, there are also reasons to expect that the degree of associations-driven overreaction in market prices may be attenuated relative to overreaction in beliefs. The reason is that – just like almost all real market environments – our betting market entails a strong element of *self-selection*, which is given by the amount of money a subject is willing to bet, captured by m . In particular, it is conceivable that those people who are more susceptible to associative recall have a loose awareness that their beliefs may be biased (but do not know how specifically, such that they cannot correct for it). If this is the case, then these people may be less inclined to bet aggressively on their beliefs and therefore influence the price less. Thus, heterogeneity in betting amounts that reflects heterogeneity in people’s confidence in their belief updating rule “re-weights” individual beliefs as far as the market price is concerned. This re-weighting of beliefs through self-selection is similar to how wealth heterogeneity re-weights individual beliefs in classical models of betting markets (Wolfers and Zitzewitz, 2004).

To illustrate, take the extreme example that out of the three subjects in a market group, one has no memory limitations, while the other two succumb to associative recall. Further suppose that the two associative recall types have sufficient doubts about the rationality of their beliefs that they do not bet at all in the market. Then, average beliefs in the market will exhibit overreaction, but the market price will not, purely as a result of differential self-selection. Of course, by an analogous logic, the market price could also reflect more overreaction than individual beliefs if those subjects that have a stronger tendency for associative recall bet more money in the parimutuel market.

Enke et al. (2023) study this type of self-selection mechanism in betting markets for various cognitive biases, but they do not consider memory. However, this is important to do because while much psychological research has documented the existence of associative recall, much less is known about whether people are willing to actually act on beliefs that are derived from associative recall, such that they become relevant when multiple individuals interact in markets.

Naturally, our main object of interest in the analysis of market prices will not be the *level* of over- or underreaction (as it could be affected by various considerations of how betting markets aggregate beliefs), but instead the causal effects of the random components of our experimental design: (i) the presence of associations; and (ii) the number of congruent first-period signals. For comparability, we analyze the data from the market experiments using the same methodology as the individual beliefs data. We first transform second-period market

prices into normalized log market price odds¹⁶ ($Impo$) in market g for company f following equation (2) and then link these to the second-period signal:

$$Impo^{g,f} = \tilde{\beta}_1 s_{k+1}^{g,f} + \tilde{\beta}_2 s_{k+1}^{g,f} T^{g,f} + \tilde{\beta}_3 T^{g,f} + \tilde{\beta}_4 (N_p^{g,f} - N_n^{g,f}) + \epsilon^{g,f} \quad (11)$$

where $T^{g,f}$ is a binary treatment indicator that equals one if, for a given market group, company f is in the *Market Cue* condition. Here, we again predict and pre-registered that $\tilde{\beta}_2 > 0$, and that this treatment difference is only driven by cases with $z > 0$.

Furthermore, within the set of *Cue* companies, we again test for an interaction effect of the second-period signal with the number of congruent first-period signals:

$$Impo^{g,f} = \tilde{\beta}_5 s_{k+1}^{g,f} + \tilde{\beta}_6 s_{k+1}^{g,f} z^{g,f} + \tilde{\beta}_7 z^{g,f} + \tilde{\beta}_8 (N_p^{g,f} - N_n^{g,f}) + \epsilon^{g,f} \quad (12)$$

We pre-registered the prediction that $\tilde{\beta}_6 > 0$.

5.3. Results

Replication of patterns on beliefs. Because we also elicited subjects’ beliefs in the market experiments, we can use these data to replicate all patterns from the individual belief elicitation treatments. This is done in Online Appendix Table C6. The results are almost identical to those reported above: (i) there is more overreaction for *Cue* than *NoCue* companies; (ii) this treatment difference only exists when the number of congruent first-period signals (z) is strictly positive; (iii) overreaction in *Cue* significantly increases in z ; and (iv) this overreaction reflects asymmetric recall of cued signals in the sense that the stronger responsiveness of second-period beliefs to first-period signals in *Cue* is only present for congruent signals.

Raw market prices data. Fig. 5 shows that, very similarly to the belief elicitation experiments, second-period market prices in *Market Cue* are more extreme than those in *Market NoCue*, following both a positive and a negative second-period signal. Indeed, we see that market prices are typically more extreme than the average Bayesian belief, though we reiterate that our primary interest is the across-treatment comparison rather than the test against the Bayesian point prediction.

Econometric analysis. More formally, we resort to Grether-style regressions in which the dependent variable consists of the (normalized) second-period log market price odds. If the market price reflected Bayesian beliefs, the OLS regression coefficient of the second-period signal would equal one.

Fig. 6 summarizes the results. Similarly to the belief elicitation experiments, there are three main takeaways. First, when $z > 0$, the responsiveness of second-period market prices to the second-period signal is significantly more pronounced in *Market Cue* than in *Market NoCue*. Second, when $z = 0$, this treatment difference disappears. Third, within *Market Cue*, overreaction of second-period beliefs monotonically increases in the number of congruent first-period signals.

Table 2 provides the regression estimates. The results confirm the statistical significance of the three main patterns that were evident from Fig. 6: (i) higher responsiveness of market prices to the second-period signal in *Market Cue* than in *Market NoCue* (column (1)); (ii) this occurs only when $z > 0$ (columns (2)–(3)); and (iii) responsiveness that significantly increases in the number of congruent signals (columns (4)–(5)).

Result 4. *Market prices react significantly more to the second-period signal in Market Cue than in Market NoCue. This treatment difference only exists when the number of congruent first-period signals is strictly positive.*

¹⁵ For discussions of whether and how betting/prediction markets generally aggregate beliefs, see Manski (2006) and Wolfers and Zitzewitz (2004, 2006).

¹⁶ Similarly to the beliefs data, when a market price is 0% or 100% we replace it by 1% and 99%, respectively, to avoid a loss of observations from the log odds definition. This occurs in 11 out of 1118 cases. We have verified that the results are quantitatively almost identical when we instead drop these observations.

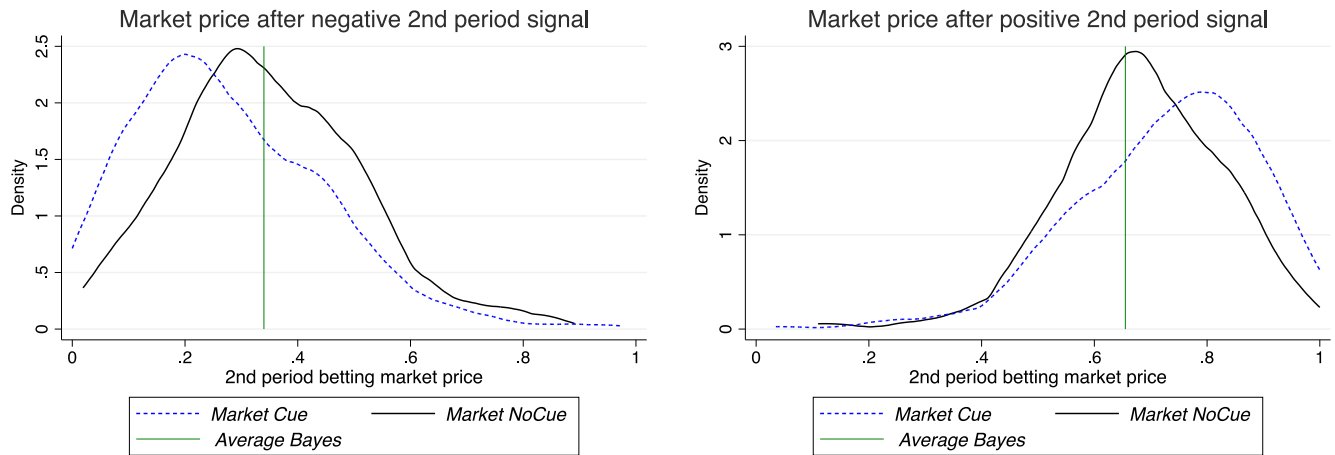


Fig. 5. Kernel density estimates of second-period market prices as a function of treatment and second-period signal. The horizontal green line indicates the average Bayesian posterior across all rounds. Kernel is Epanechnikov. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)




Fig. 6. Effect of second-period signal on (normalized) second-period log market price odds, as a function of the number of congruent first-period signals. The point estimates are derived from regression Eq. (11), which is run separately for each value of z . We do not show $z = 4$ because there are very few observations with such a signal history. The figure plots $\hat{\beta}_1$ for Market NoCue and $\hat{\beta}_1 + \hat{\beta}_3$ for Market Cue. Whiskers show 95% confidence intervals, computed based on clustering at the market group level.

Result 5. Within condition Market Cue, the responsiveness of market prices to the second-period signal increases significantly in the number of congruent first-period signals.

Do markets attenuate associations-based overreaction? In light of the discussion about self-selection in markets potentially attenuating the effect of associative recall on overreaction, it is of interest to compare the quantitative magnitude of associations-driven overreaction in the betting market (Table 2) with that in individual beliefs (Table 1).¹⁷ The relevant quantities of interest here are the causal effects of the treatment and of the number of cued signals, rather than the level of overreaction across treatments. This is because only the causal effects reflect the impacts of associative recall, while the baseline level of overreaction in markets and beliefs may differ for various reasons. Comparing column (1) of Tables 1 and 2, we see that the causal effect of the treatment on overreaction is 0.61 in individual beliefs ($s.e. = 0.15$) and 0.65 in market prices ($s.e. = 0.09$). Similarly, comparing column (4) in Table 1

with column (4) in Table 2, we see that the causal effect of the number of cued signals is 0.32 ($s.e. = 0.10$) in the case of individual beliefs, while it is 0.34 ($s.e. = 0.08$) in the market experiment. These differences are not statistically significant and suggest that subjects who rely on associative recall in forming beliefs do not select out of the market and hence affect market prices. In sum, we believe that these results further underscore the economic relevance of associative memory.

Result 6. Associative-recall-based overreaction in market prices is as large as overreaction in average individual beliefs, despite the scope for self-selection.

6 Discussion

By presenting a set of theory-driven experiments that build a bridge between psychological paradigms on cued recall and structured, quantitative financial decision tasks, this paper has provided a causal analysis of the role of associative memory for belief formation and market behavior. In doing so, we have provided two pieces of evidence that speak to the economic relevance of associative memory. First, associative recall generates systematic overreaction of beliefs when context and news are correlated in a consistent fashion over time, as is likely the case in practice. Second, we have shown that this overreaction in beliefs leads to systematic overreaction of market prices in a betting market environment.

Our experiments are related to an active literature that documents overreaction in survey expectations about financial and macroeconomic variables (e.g., Bordalo et al., 2020a). The result of overreaction in field data is often considered to be a slight puzzle from the perspective of laboratory research on belief formation (Benjamin, 2018). This is because structured laboratory belief updating problems almost always find underreaction. 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 underreaction 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 and “cue-dependent forgetting”.

While our controlled experiments provide clean evidence that associative recall can generate overreaction to news, both at the individual and the market level, an open question is to what extent associative actually matters in real financial markets. Recent work (e.g., Charles, 2022a,b; Jiang et al., 2023) points to an important role of memory biases in financial markets. At the same time, there are plausibly forces

¹⁷ The Beliefs and Market experiments were conducted using within-session, random assignment to experiments.

Table 2
Overreaction in market prices as a function of treatment and signal history.

Sample:	Dependent variable: 2nd period normalized log market price odds				
	Cue vs. NoCue			Cue	
	Full	$z = 0$	$z > 0$	Full	Full
	(1)	(2)	(3)	(4)	(5)
t = 2 signal	1.17*** (0.07)	1.24*** (0.12)	1.10*** (0.08)	1.39*** (0.11)	1.33*** (0.10)
t = 2 signal \times 1 if Cue	0.65*** (0.09)	−0.034 (0.14)	0.94*** (0.11)		
t = 2 signal \times # of congruent t = 1 signals				0.34*** (0.08)	0.41*** (0.07)
1 if Cue	−0.034 (0.06)	−0.063 (0.11)	−0.050 (0.08)		
t = 1 signals (pos. minus neg.)	0.45*** (0.04)	0.33*** (0.09)	0.44*** (0.04)	0.47*** (0.08)	
# of congruent t = 1 signals				−0.026 (0.04)	−0.032 (0.04)
t = 1 normalized log market price odds					0.35*** (0.05)
Observations	1118	330	788	560	559
Adjusted R^2	0.63	0.37	0.69	0.72	0.73

Notes. OLS estimates, robust standard errors (in parentheses) are clustered at the market group level. Following the estimating equations (11) and (12), the estimations do not include a constant. However, we have verified that including a constant delivers almost identical results. Both the first- and the second-period log market price odds are normalized by the log diagnosticity odds as described by Eq. (2). See Table 1 for details on the construction of each variable. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

that may limit the role of memory biases for financial decisions. For instance, past return data can often be looked up and need not be remembered, making memory constraints less relevant. This suggests that forces other than memory biases likely contribute to overreaction to news in financial markets (e.g., Augenblick et al., 2023; Ba et al., 2023).

CRedit authorship contribution statement

Benjamin Enke: Writing – review & editing, Writing – original draft, Formal analysis, Conceptualization. **Frederik Schwerter:** Writing – review & editing, Writing – original draft, Formal analysis, Conceptualization. **Florian Zimmermann:** Writing – review & editing, Writing – original draft, Formal analysis, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

<https://data.mendeley.com/datasets/545wb2gnd7/1>.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.jfineco.2024.103853>.

References

- Abarbanell, Jeffery S., Bernard, Victor L., 1992. Tests of analysts' overreaction/underreaction to earnings information as an explanation for anomalous stock price behavior. *J. Finance* 47 (3), 1181–1207.
- Afrouzi, Hassan, Kwon, Spencer Yongwook, Landier, Augustin, Ma, Yueran, Thesmar, David, 2023. Overreaction in expectations: Evidence and theory. *Q. J. Econ.* 138 (3), 1713–1764.
- Augenblick, Ned, Lazarus, Eben, Thaler, Michael, 2023. Overinference from weak signals and underinference from strong signals. Working Paper.
- Ba, Cuimin, Bohren, Aislinn, Imas, Alex, 2023. Over- and underreaction to information. Working Paper.
- Baddeley, Alan, Eysenck, Michael W., Anderson, Michael C., 2020. *Memory*. Routledge.
- Barrero, Jose Maria, 2022. The micro and macro of managerial beliefs. *J. Financ. Econ.* 143 (2), 640–667.
- Benjamin, Daniel J., 2018. Errors in probabilistic reasoning and judgmental biases. In: *Handbook of Behavioral Economics*.
- Beutel, Johannes, Weber, Michael, 2022. Beliefs and portfolios: Causal evidence. Chicago Booth Research Paper 22–08.
- Bodoh-Creed, Aaron, 2020. Mood, memory, and the evaluation of asset prices. *Rev. Finance* 24 (1), 227–262.
- Bordalo, Pedro, Conlon, John J., Gennaioli, Nicola, Kwon, Spencer Yongwook, Shleifer, Andrei, 2023. Memory and probability. *Q. J. Econ.* 138 (1), 265–311.
- Bordalo, Pedro, Gennaioli, Nicola, Ma, Yueran, Shleifer, Andrei, 2020a. Overreaction in macroeconomic expectations. *Amer. Econ. Rev.* 110 (9), 2748–2782.
- Bordalo, Pedro, Gennaioli, Nicola, Shleifer, Andrei, 2020b. Memory, attention, and choice. *Q. J. Econ.* 135 (3), 1399–1442.
- Bouchaud, Jean-Philippe, Krueger, Philipp, Landier, Augustin, Thesmar, David, 2019. Sticky expectations and the profitability anomaly. *J. Finance* 74 (2), 639–674.
- Brunnermeier, Markus, Farhi, Emmanuel, Kojen, Ralph S.J., Krishnamurthy, Arvind, Ludvigson, Sydney C., Lustig, Hanno, Nagel, Stefan, Piazzesi, Monika, 2021. Perspectives on the future of asset pricing. *Rev. Financ. Stud.* 34 (4), 2126–2160.
- Charles, Constantin, 2022a. Memory and trading. Working Paper.
- Charles, Constantin, 2022b. Memory moves markets. Available at SSRN 4019728.
- Charles, Constantin, Frydman, Cary, Kilic, Mete, 2024. Insensitive investors. *J. Finance*, forthcoming.
- D'Acunto, Francesco, Malmendier, Ulrike, Ospina, Juan, Weber, Michael, 2021. Exposure to grocery prices and inflation expectations. *J. Polit. Econ.* 129 (5), 1615–1639.
- Danz, David, Vesterlund, Lise, Wilson, Alistair, 2022. Belief elicitation and behavioral incentive compatibility. *Amer. Econ. Rev.* 112 (9), 2851–2883.
- Enke, Benjamin, Gneezy, Uri, Hall, Brian, Martin, David, Nelidov, Vadim, Offerman, Theo, van de Ven, Jeroen, 2021a. Cognitive biases: Mistakes or missing stakes? *Rev. Econ. Stat.* 105 (4), 818–832.
- Enke, Benjamin, Graeber, Thomas, 2023. Cognitive uncertainty. *Q. J. Econ.* 138 (4), 2021–2067.
- Enke, Benjamin, Graeber, Thomas, Oprea, Ryan, 2023. Self-selection and bias in the aggregate. *Amer. Econ. Rev.* 113 (7), 1933–1966.
- Fudenberg, Drew, Levine, David K., et al., 2014. Learning with recency bias. *Proc. Natl. Acad. Sci.* 111, 10826–10829.
- Gennaioli, Nicola, Shleifer, Andrei, 2018. *A Crisis of Beliefs: Investor Psychology and Financial Fragility*. Princeton University Press.
- Giglio, Stefano, Maggiori, Matteo, Stroebel, Johannes, Utkus, Stephen, 2021. Five facts about beliefs and portfolios. *Amer. Econ. Rev.* 111 (5), 1481–1522.
- Goetzmann, William, Kim, Dasol, Shiller, Robert, 2022. Crash narratives. NBER Working Paper.

- Graeber, Thomas, Roth, Christopher, Zimmermann, Florian, 2023. Stories, statistics and memory. Working Paper.
- Greenwood, Robin, Shleifer, Andrei, 2014. Expectations of returns and expected returns. *Rev. Financ. Stud.* 27 (3), 714–746.
- Grether, David M., 1980. Bayes rule as a descriptive model: The representativeness heuristic. *Q. J. Econ.* 95, 537–557.
- Hartzmark, Samuel M., Hirshman, Samuel D., Imas, Alex, 2021. Ownership, learning, and beliefs. *Q. J. Econ.* 136 (3), 1665–1717.
- Hossain, Tanjim, Okui, Ryo, 2013. The binarized scoring rule. *Rev. Econ. Stud.* 80 (3), 984–1001.
- Jiang, Zhengyang, Liu, Hongqi, Peng, Cameron, Yan, Hongjun, 2023. Investor memory and biased beliefs: Evidence from the field. Working Paper.
- Kahana, Michael Jacob, 2012. *Foundations of Human Memory*. OUP USA.
- Kendall, Chad, Oprea, Ryan, 2018. Are biased beliefs fit to survive? An experimental test of the market selection hypothesis. *J. Econom. Theory* 176, 342–371.
- Malmendier, Ulrike, 2021. FBBVA lecture 2020 exposure, experience, and expertise: Why personal histories matter in economics. *J. Eur. Econom. Assoc.* 19 (6), 2857–2894.
- Malmendier, Ulrike, Nagel, Stefan, 2015. Learning from inflation experiences. *Q. J. Econ.* 131 (1), 53–87.
- Malmendier, Ulrike, Wachter, Jessica A., 2021. Memory of past experiences and economic decisions. Available at SSRN 4013583.
- Manski, Charles F., 2006. Interpreting the predictions of prediction markets. *Econ. Lett.* 91 (3), 425–429.
- Mullainathan, Sendhil, 2002. A memory-based model of bounded rationality. *Q. J. Econ.* 117 (3), 735–774.
- Plott, Charles R., Wit, Jorgen, Yang, Winston C., 2003. Parimutuel betting markets as information aggregation devices: Experimental results. *Econom. Theory* 22 (2), 311–351.
- Schacter, Daniel L., 2008. *Searching for Memory: The Brain, the Mind, and the Past*. Basic books.
- Shiller, Robert J., 2017. Narrative economics. *Am. Econ. Rev.: Pap. Proc.* 107 (4), 967–1004.
- Shiller, Robert J., 2019. *Narrative Economics: How Stories Go Viral and Drive Major Economic Events*. Princeton University Press.
- Tulving, Endel, Thomson, Donald M., 1973. Encoding specificity and retrieval processes in episodic memory. *Psychol. Rev.* 80 (5), 352.
- Wachter, Jessica A., Kahana, Michael J., 2019. A retrieved-context theory of financial decisions. Available at SSRN.
- Wilson, Andrea, 2014. Bounded memory and biases in information processing. *Econometrica* 82 (6), 2257–2294.
- Wolfers, Justin, Zitzewitz, Eric, 2004. Prediction markets. *J. Econ. Perspect.* 18 (2), 107–126.
- Wolfers, Justin, Zitzewitz, Eric, 2006. *Prediction Markets in Theory and Practice*. National Bureau of Economic Research Cambridge, Mass., USA.