A retrieved-context theory of financial decisions*

Jessica A. Wachter† Michael Jacob Kahana‡

July 10, 2019

Abstract

Studies of human memory indicate that features of an event evoke memories of prior associated contextual states, which in turn become associated with the current event's features. This mechanism allows the remote past to influence the present, even as agents gradually update their beliefs about their environment. We apply a version of retrieved context theory, drawn from the literature on human memory, to four problems in asset pricing and portfolio choice: over-persistence of beliefs, providence of financial crises, price momentum, and the impact of fear on asset allocation. These examples suggest a recasting of neoclassical rational expectations in terms of beliefs as governed by principles of human memory.

*We are grateful to Ernst Fehr, Cary Frydman, Nicola Gennaioli, Nikolai Roussanov, and seminar participants at Wharton, at the NBER Behavioral Finance Meeting, and at the Sloan/Nomis Workshop on the Cognitive Foundations for Human Behavior for helpful comments.

†Department of Finance, The Wharton School, University of Pennsylvania and NBER; Email: jwachter@wharton.upenn.edu

‡Department of Psychology, University of Pennsylvania; Email: kahana@psych.upenn.edu

1 Introduction

Standard decision-making under uncertainty starts with a probability space and an information structure. The information structure implies that the agent associates a value with every subset of the space and then maximizes expected utility. This is the approach of Savage (1954). The difficulty that agents have in forming beliefs over an entire state space has been formulated in the Ellsberg paradox (Ellsberg, 1961), ambiguity aversion formalized by Gilboa and Schmeidler (1989) and in the alternative representations of choice as a probabilistic selection among a small set of alternatives, due to Luce (1959) and McFadden (2001).[1] The set of possible states of nature is impossibly large and ever-changing. Nonetheless, we as individuals do manage to make decisions under uncertainty.

In this paper, we propose a memory-based model of decision-making under uncertainty. A wealth of data support the idea of a human memory system that maintains a record of associations between experiential features of the environment, and underlying contextual states (Kahana, 2012). This record of associations, together with inference about the current contextual state, constitutes a belief system that could potentially affect any kind of choice under uncertainty. This belief system responds to the current environment through retrieved context. The mechanism of retrieved context is how memory "knows" what information is most relevant to bring forward to our attention at any given time. At the same time, any new experience, and the context itself, is then stored again in the memory system (Howard and Kahana, 2002).

This paper applies these concepts to puzzles in asset pricing and portfolio choice that defy the standard Bayesian paradigm. Chief among these are the result that life experience has near-permanent effects on financial decision making (Malmendier and

[1] The problem of determining the underlying state space continues to be a point of contention in recent literature on ambiguity aversion: see, for example, the debate concerning rectangularity of the model set (Epstein and Schneider, 2003; Hansen and Sargent, 2018).

Nagel, 2011, 2016; Malmendier et al., 2017; Malmendier and Shen, 2018), and that an exogenous cue, such as a horror movie, can influence financial decisions (Guiso et al., 2018). More speculatively, we then apply the framework to a broader set of phenomena, such as the sudden onset of a financial crisis, and the momentum effect in the cross-section of asset returns.

When making a decision, an agent is confronted by certain *features* of the environment. For simplicity, we assume that features are perceived as discrete, and that there are a finite number of possible features pertaining to a particular decision. A feature vector, then is an element of $\mathcal{B}^n \subset \mathbb{R}^n$, where $\mathcal{B}^n$ is a set of basis vectors that spans n-dimensional space. For convenience, we assume the standard basis. That is, the time$-t$ features vector f_t has ith element equal to 1 if the ith instance of the feature is realized at time t and all other elements equal to zero. One can think of the features vector as a mathematical representation of objective, verifiable, and most likely transitory, aspects of the environment.

Features are connected over time through *context.* Context is persistent (usually) and endows the agent with an understanding of possibly latent aspects of the environment that are relevant for the decision at hand. Context may also be subjective. Specifically, define the *context space* as the standard simplex $\mathcal{A}_c \subset \mathbb{R}^m$, for $m \leq n$. That is $\mathcal{A}_c = \{c_t = [c_{1t}, \ldots, c_{mt}]^\top \in \mathbb{R}^m \,|\, \iota^\top c_{it} = 1\}$, where ι denotes a conforming vector of ones. We will think of c_t as assigning probabilities to the underlying states of nature, and at that point proceed in a manner similar to the standard economic approach. Indeed, a special case of our framework will be the Bayesian problem under which the agent learns about an unobserved state (context) from observed data (features). Principles of memory, however, can lead context to evolve in ways that are distinctly non-Bayesian.

Whereas many applications of psychological principles to economic decision making have focused on cognitive biases such as loss aversion and narrow framing (see Barberis

(2013) for a review), the literature on human learning and memory offers a different perspective. Three major laws govern the human memory system: similarity, contiguity, and recency: Similarity refers to the priority accorded to information that is similar to the presently active features, contiguity refers to the priority given to features that share a history of co-occurrence with the presently active features, and recency refers to priority given to recently experienced features. All three "laws" exhibit universality across agents, feature types, and memory tasks and thus provide a strong basis for a theory of economic decision-making.

While few economic models explicitly incorporate these laws, there are exceptions. Gilboa and Schmeidler (1995) replace axiomatic expected utility with utility computed using probabilities that incorporate the similarity of the current situation to past situations. Mullainathan (2002) proposes a model in which agents tend to remember those past events which resemble current events, and where a previous recollection increases the likelihood of future recollection. He applies the model to the consumption-savings decision. Nagel and Xu (2018) show that a constant-gain learning rule about growth in dividends can explain a number of asset pricing puzzles; they motivate this learning rule using the memory principle of recency. These models do not employ context-based retrieval, which is the focus of our paper. Bordalo et al. (2017) develop a model based on the geometric similarity of representations in memory. They focus on the the role that similarity in memory representations plays in accounting for the propensity of agents to make large expenditures on housing or durable goods when lower expenditures would appear optimal by standard theory. Their work differs from ours in that we focus on the retrieval of prior contextual states, and we directly model contextual evolution. In their model, as in psychological studies such as Godden and Baddeley (1975), context is embedded in the environment, and thus is static; the feature layer of the environment and the context layer are the same.

The remainder of the paper is organized as follows. Section 2 describes the general form of the memory model we consider, and how we integrate it into a model for decision-making. Section 3 describes the application to over-persistence of memory. Section 4 shows how the *jump back in time* (Howard et al., 2012a), can lead to a model of financial crises. Section 5 shows how the slow adjustment of context leads to a model of momentum. Section 6 shows how the model reproduces a result that a seemingly irrelevant stimulus, such as a movie, can change portfolio choice. Section 7 concludes.

2 Integrating Memory into Decision Making

Unless agents observe all decision-relevant information at the moment of choice, they must use their memory of past experiences to guide their decisions. The question of how past experiences influence present behavior has occupied the attention of experimental psychologists for more than a century (Ebbinghaus, 1913; Müller and Pilzecker, 1900; Jost, 1897; Müller and Schumann, 1894; Ladd and Woodworth, 1911; Carr, 1931). Because memories of recent experiences readily come to mind, early scholars sought to uncover the factors that lead to forgetting. Their experiments quickly challenged the folk assertion that memories decay over time, eventually becoming completely erased. Rather, they found that removing a source of interference, or reinstating the "context" of original learning, readily restored these seemingly forgotten memories (McGeoch, 1932; Underwood, 1948; Estes, 1955). In these early papers, context represents the set of latent (or background) information not specifically related to the present stimulus. Such contextual information could include extrinsic features of the environment, such as the room or setting in which information is studied (Abernathy, 1940; Godden and Baddeley, 1975), but could also include internal states of the agent, including the current set of thoughts, emotions, goals, and concerns that form the cognitive milieu in which new learning occurs (Kahana, 2012). According to modern memory theories, the set of psychological (or neural) features that represent a stimulus enter into association with a mental representation of context, and the database of such associations form the basis for performance in subsequent recall, recognition, and categorization tasks (Howard and Kahana, 2002).

In this paper, we use a dynamic model of contextual coding based closely on that developed by Polyn et al. (2009) to account for data on the dynamics of memory search. Consider an experiment in which the agent studies items, denoted f_i, $i = 1, \ldots, N$, where f_i is a basis vector in $\mathbb{R}^n$, for n large. Memory associates items with (latent) context, $c_i \in \mathbb{R}^m$ via a matrix that sums the outer products of item and context vectors. We can this $m \times n$ the *features-to-context* matrix, and we denote it by $W^{f \to c}$.[2] We will give an explicit form for this matrix in what follows. The model is associative in the sense that "cueing" with context allows the model to recover the items associated with that context, and cueing with an item (a feature) recovers the contexts previously associated with that item.

Context for a current item i depends on context for the previous item and the context associated with the new item presented. That is, context satisfies the recursion

$$c_i = \rho_i c_{i-1} + \zeta c_i^{\text{in}}, \tag{1}$$

where memory retrieves c_i^{in} from the item, based on the prior history of associations between items and context:

$$c_i^{\text{in}} = \frac{W_{i-1}^{f \to c} f_i}{||W_{i-1}^{f \to c} f_i||}. \tag{2}$$

Note that (2) implies that c_i^{in} is scaled so that its length equals one for a given norm $||\cdot||$. According to Equations 1 and 2, context is a recency-weighted sum of presented items. The amount by which element values decay with each presented item is governed by the model parameter ζ. In the Polyn et al. (2009) and related models, it is convenient

[2]Memory models such as Polyn et al. (2009) and Howard and Kahana (2002) make use of a second matrix called the *context-to-features* matrix $W^{c \to f}$. The superscript in $W^{f \to c}$ distinguishes it from $W^{c \to f}$. We do not use $W^{c \to f}$ in this study.

to have $||\cdot||$ be the L^2-norm. In order that c_i lie on the unit circle, a linear relation between c_i, c_{i-1} and c_i^{in} must hold only approximately. For this reason, $\rho_i \approx 1-\zeta$, which is why the coefficient on c_{i-1} has an i subscript.[3] Polyn et al. (2009) close the model by setting $W_0^{f\to c}$ to be the zero matrix and defining its evolution throughout the experiment as:

$$W_i^{f\to c} = W_{i-1}^{f\to c} + c_i f_i^\top. \tag{4}$$

Equation (4) implies an intuitive relation between a feature and the context it retrieves. Specifically:

$$c_i^{\text{in}} = \frac{W_{i-1}^{f\to c} f_i}{||W_{i-1}^{f\to c} f_i||} \propto \sum_{j=0}^{i} (c_j f_j^\top) f_i = \sum_{j=0}^{i} c_j (f_j^\top f_i). \tag{5}$$

Note that $f_j^\top f_i$ is simply the inner product of f_j with f_i, and thus is a scalar. Under our specification with these as orthonormal basis vectors, this value equals zero if $f_j \neq f_i$ and 1 otherwise. Thus item i recalls the context under which the agent last experienced item i. If the subject experiences i under multiple contexts, the subject recalls a weighted average of the contexts, where the weights are given by the number of times the subject experiences f_i. Also, because context is autoregressive, recall of item i calls to mind all of the items that are near to i in the sense that they are also associated with the context.

Before turning to the application of the memory model to financial decision-making, we briefly summarize the psychological and neural evidence for context as an internal state.

---

[3]Specifically,

$$\rho_i = \sqrt{1+\zeta^2[(c_{i-1}\cdot c_i^{\text{in}})^2 - 1]} - \zeta(c_{i-1}\cdot c_i^{\text{in}}). \tag{3}$$

2.1 Psychological and Neural Basis for Contextual Retrieval

In the memory laboratory, researchers create experiences by presenting subjects with lists of easily identifiable items, such as common words or recognizable pictures. Subjects attempt to remember the previously experienced items under varying retrieval conditions. These conditions include *free recall*, in which subjects recall as many items as they can in any order, *cued recall*, in which subjects attempt to recall a particular target item in response to a cue, and *recognition* in which subjects judge whether or not they encountered a test item on a study list. In each of these experimental paradigms, memory obeys the classic "Laws of Association" which appear first in the work of Aristotle, and later in (Hume, 1748). The first of these is *recency*: human subjects exhibit better memory for recent experiences, *semantic similarity*: we remember experiences that are most similar in meaning to those we are currently experiencing, and finally, *temporal contiguity*: we remember items that occurred contiguously in time to recently-recalled items.[4]

A longstanding and persistently active research agenda in experimental psychology seeks to uncover the cognitive and neural mechanisms that could give rise to these regularities. Experimental psychologists have proposed many hypotheses and have (accordingly) refined the tasks above in a number of ways. Some striking findings include the fact that recency and contiguity have similar magnitudes at short and long time scales.[5] Several classic explanations, though successful in many ways, struggled to explain this scale invariance. One highly influential class of explanations posits the

[4]Although quantified in the memory laboratory, each of these phenomena appears robustly in real-world settings, such as recall of autobiographical memories and news events (Moreton and Ward, 2010; Uitvlugt and Healey, 2019).

[5]To measure the effect of contiguity on memory retrieval, researchers examine subjects' tendency to successively recall items experienced in proximate list positions. In free recall tasks, this tendency appears as decreasing probability of successively recalling items f_i and f_{i+lag} as a function of lag, conditional on the availability of that transition (Kahana, 1996). This function reaches its maximum at $lag = \pm 1$, but also exhibits a forward asymmetry in the form of higher probability for positive as compared with negative lags. This indicates the tendency for recollection to move forward in time (Kahana and Caplan, 2002).

existence of a specialized retrieval process for recently-experienced items (short-term memory).[6] A related idea is that associations chain together in the mind of the subject.[7] In contrast, retrieved context theory does not derive contiguity and similarity through direct interitem associations. Rather, they arise because the contextual information retrieved during the recall of an item overlaps with the contextual information associated with similar and neighboring items. Underlying context naturally generates scale invariance.

Figure 1 summarizes some of the evidence supportive of context retrieval. Figure 1A shows that interitem distraction and test delay do not disrupt the temporal contiguity effect (TCE) seen in the relation between transition probability and lag, which speaks directly to time-scale invariance. Figure 1B-D shows that the TCE appears robustly for both younger and older adults, for subjects of varying intellectual ability, and for both naïve and highly practiced subjects. Figure 1E shows that the TCE appears even for transitions between items studied on completely distinct lists, despite these items being separated by many other item presentations. Figure 1F-H shows that the TCE also predicts confusions between different study pairs in a cued recall task, in errors made during probed recall of serial lists, and in tasks that do not depend on inter-item associations at all, such as picture recognition (see caption for details). Finally, long-range contiguity appears in many real-life memory tasks, such as recalling autobiographical memories(Moreton and Ward, 2010) and remembering news events

[6]The view that recency arises from specialized retrieval processes associated with short-term memory rose to prominence in the 1960s. According to these dual-store models, separate short-term and long-term memory stores support retrieval of information experienced at short and long time scales, with short-term (or "working") memory holding a small number of information units through an active rehearsal process and supporting the rapid and accurate retrieval and manipulation of that information. According to these models, retrieval from long-term memory involved a search process guided by interitem associations and context-to-item associations, and subject to interference from similar memories (Kahana, 2012).

[7]Continental philosophers saw contiguity as the result of chained associations (Herbart, 1834) that could be easily disrupted by interfering mental activity (Thorndike, 1932). This idea took form in cognitive models that conceived of associations as being forged in a limited-capacity short-term memory store (Atkinson and Shiffrin, 1968; Raaijmakers and Shiffrin, 1980), perhaps arising as the result of imagery or linguistic mediation (Murdock, 1974).

(Uitvlugt and Healey, 2019).

A second source of data in favor of retrieved context arises from neurobiology. At a neurobiological level, this implies that the brain states representing the context of an original experience reactivate or replay during the subsequent remembering of that experience. Several studies tested this idea using neural recordings. These studies found that in free recall (Manning et al., 2011a), cued recall (Yaffe et al., 2014) and recognition memory (Howard et al., 2012b; Folkerts et al., 2018) brain activity during memory retrieval resembled not only the activity of the original studied item, but also the brain states associated with neighboring items in the study list. Thus, one observes contiguity both at the behavioral and at the neural level, with these effects being strongly correlated (Manning et al., 2011a). Finally, this recursive nature of the contextual retrieval process offers a unified account of many other psychological phenomena including the spacing effect (Lohnas and Kahana, 2014b), the compound cueing effect (Lohnas and Kahana, 2014a), and the phenomena of memory consolidation and reconsolidation (Sederberg et al., 2011).

Memory theory thus indicates that remembering an item involves a jump-back-in-time to the state of mind that obtained when the item was previously experienced. This neural reinstatement, in turn, becomes re-encoded with the new experience and also persists to flavor the encoding of subsequently experienced items. The persistence of the previously retrieved contextual states thus enables memory to carry the distant past into the future, thus allowing the contextual states associated with an old memory to re-enter one's life following a salient cue and associate with subsequent "neutral" memories. While the original memory is retained in association with its encoding context, the retrieval and re-experiencing of that memory forms a new memory in association with the mixture of the prior and retrieved context. Memory theory thus also predicts that multiple recalls of an item will largely appear to the agent as if there were multiple experiences, when in fact there was only perhaps a single experience.

The well-known existence of post-traumatic stress disorder attests to the power of continually recurring danger that is wholly in the mind of an agent.

2.2 Retrieved-context theory and financial decisions

The above theory treats memory as an outcome of a mechanistic process. There is no explicit decision-maker facing an objective function. To map the above framework into financial decisions, we assume a link between memory and the subjective probability the agent assigns to future events. The equations themselves suggest such a link. For example, (2) reflects storage of co-occurrences of features with contexts, while (1) reflects (partial) updating of beliefs based on new information.

We start by making a simple technical change. We replace the normalization of the context vector by the L^2-norm with normalization by the L^1-norm.[8] Because elements of context are positive, this implies that context vectors sum to one; it is natural then to interpret the context vector as a vector of probabilities. That is, we define $c_t \in \mathcal{A}_c$, the simplex in m-dimensional space, to be the agent's context at time t.

Equation 1 becomes

$$c_t = \rho c_{t-1} + \zeta c_t^{\text{in}}, \tag{6}$$

with $\rho = 1 - \zeta$,[9] and where

$$\begin{aligned} c_t^{\text{in}} &= \frac{W_{t-1}^{f \to c} f_t}{||W_{t-1}^{f \to c} f_t||} &(7)\\ &= \frac{1}{\iota^\top W_{t-1}^{f \to c} f_t} W_{t-1}^{f \to c} f_t, &(8) \end{aligned}$$

where ι is a conforming vector of 1s. Note that c_t^{in} is only defined up to a scaling

[8]There is no experimental evidence in favor of one norm versus the other. The L^2-norm is convenient for modeling free recall, which requires retrieving features from context as well as context from features.

[9]Because $c_{t-1}, c_t^{\text{in}} \in \mathcal{A}_c$, $c_t \in \mathcal{A}_c$. Unlike, (1), the relation between $\rho = 1 - \zeta$ is exact.

constant. We therefore can normalize $W^{f\to c}$ with no change to the evolution in context. Specifically, define

$$W_t^{f\to c} = \frac{1}{t+\tau}\sum_{s=-\tau}^{t} c_s f_s^\top, \tag{9}$$

where τ represents the length of the prior sample. Equivalently, we can initialize $W_t^{f\to c}$ at zero, and use the following updating rule, starting at τ:

$$W_t^{f\to c} = \frac{\tau+t-1}{\tau+t} W_{t-1}^{f\to c} + \frac{1}{\tau+t} c_t f_t^\top. \tag{10}$$

While (4) and (10) may appear to be two alternative ways of closing the model, they generate identical implications for context evolution, assuming that we normalize elements of c_t to sum to 1.

In the special case where both context and features are basis vectors, (9) literally represents the joint contingencies of contexts and features.

The resultant model, when combined with standard economic optimization, becomes a memory-driven model of choice under uncertainty. The agent still maximizes utility subject to the usual constraints. However, beliefs come from memory. Figure 2 illustrates the mechanism behind retrieved context theory. The current state of context contains both an autocorrelated component that overlaps with the contexts of recent experiences, and a retrieved context component that overlaps with items experienced close in time to the just-recalled item(s) (below we show how this accounts for a phenomenon called the contiguity effect). The figure illustrates these two effects as spotlights shining down on memories arrayed on the stage of life. Memories are not truly forgotten, but just obscured when they fall outside of the spotlights.

3 Retrieved-context theory and the persistence of beliefs

A classic problem in asset allocation is that of an investor allocating wealth between a risky asset (with unknown return) and a riskless asset with known return (Arrow, 1971; Pratt, 1976). This deceptively simple problem is the subject of a large and sophisticated literature (Wachter, 2010). In a new take on this classic problem, Malmendier and Nagel (2011) report an intriguing pattern in the portfolio choice of investors in the Survey of Consumer Finances. Investors whose lifetime experience includes periods with lower stock returns invested a lower percentage of their wealth in stocks as compared with investors whose lifetime experience includes periods with higher returns. While, on one level this may seem intuitive, it is a puzzle from the point of view of standard asset allocation theories. For one thing, experience should not matter, only objective data on returns. For another, even if investors over-weight their own experience, and under-weight returns outside of their experience, investors in the sample had experiences of sufficiently long length (and the return distribution is sufficiently ergodic) such that their beliefs should quickly converge.

Here, we abstract from many interesting features of the Malmendier and Nagel (2011) study. For example, investors exhibit a recency effect (their portfolio choice depends more on recent observations than on past observations) which we do not emphasize here, but which is very much in the spirit of a memory model. We focus on a qualitative implication of their results, namely, that personal experiences can continue to influence investors' beliefs, even though there are sufficient data (if investors were Bayesian) to over-ride a specific time path of experience. Thus, in this section, we focus simply on the question of persistence of beliefs, abstracting both from many features of memory models, and many features of the portfolio choice data. We do so to highlight the novel implications of the theory.

3.1 The portfolio choice problem

We consider a two-period problem in which an investor chooses a portfolio at time 0 and consumes wealth at time 1. The agent starts time 0 with wealth of 1. We denote wealth at time 1 as $\tilde{X}$. The investor chooses between a riskless bond, earning net return of zero, and a risky stock, earning net return of $\tilde{r}$. The agent also receives risky labor income $\tilde{y}$. To capture the preference of more wealth to less, and risk averse, we assume mean-variance preferences:

$$\max_{\pi} E[\tilde{X}] - \frac{1}{2}\mathrm{Var}(\tilde{X}) \tag{11}$$

where π is the percent allocation to the risky asset, and where the assumptions above imply

$$\tilde{X} = 1 + \pi\tilde{r} + \tilde{y}. \tag{12}$$

Note that the expectation and the variance in (11) are with respect to the agents' subjective preferences. Substituting (12) into (11), and setting the derivative of the objective function with respect to π equal to zero leads to

$$\pi = \frac{E^*\tilde{r} - \mathrm{Cov}^*(\tilde{r}, \tilde{y})}{\mathrm{Var}^*(\tilde{r})}. \tag{13}$$

If stocks deliver a low return in a negative labor income state, that makes them unattractive.

Assume that the risky return takes on two possible values $r(gain) > r(loss)$. Assume labor income $\tilde{y}$ takes on two possible values $y(normal) > y(depression)$. We consider beliefs that take the following form: gain and loss states each occur with probability 1/2, that a gain and a depression cannot co-occur, and that a depression has (unconditional) probability p, for $p \leq 1/2$. The following matrix captures the state space and the probabilities:

$$P = \begin{bmatrix} \text{Prob}(\textit{gain \& normal}) & \text{Prob}(\textit{loss \& normal}) \\ \text{Prob}(\textit{gain \& depression}) & \text{Prob}(\textit{loss \& depression}) \end{bmatrix} = \begin{bmatrix} \frac{1}{2} & \frac{1}{2}-p \\ 0 & p \end{bmatrix}.$$

where $p \in [0, \frac{1}{2}]$.

Further assume $\tilde{r}(\textit{gain}) = 1+\sigma$, $\tilde{r}(\textit{loss}) = 1-\sigma$, for $\sigma > 0$, $\tilde{y}(\textit{normal}) = y > 0$ and $\tilde{y}(\textit{depression}) = 0$. Then. $E\tilde{r} = 1$, $\text{Var}(\tilde{r}) = \sigma^2$. Note that $\tilde{y}$ is a Bernoulli process multiplied by a constant y, so that

$$\begin{aligned} E\tilde{y} &= (1-p)y \\ \text{Var}(\tilde{y}) &= p(1-p)y^2. \end{aligned}$$

Direct calculation implies

$$\text{Cov}(\tilde{r}, \tilde{y}) = E\big[(\tilde{r} - E\tilde{r})\tilde{y}\big] = \frac{1}{2}\sigma y - \left(\frac{1}{2} - p\right)\sigma y = p\sigma y.$$

Then the optimal allocation (13) equals

$$\pi(p) = \frac{1 - py\sigma}{\sigma^2}. \tag{14}$$

The higher the beliefs about the depression, the lower is the portfolio allocation.

3.2 Memory for stock market gains and losses

We identify the feature vector f_t with realizations of the stock market, so that $f_t = [1, 0]^\top$ represents gain, and $f_t = [0, 1]^\top$ represents loss. We assume the context vector c_t represents the unobserved labor market state. $c_t = [1, 0]^\top$ implies 100% probability on no depression in the labor market.

Define the features-to-context matrix as in (9). For τ sufficiently large, if the labor income state were perfectly observed (as well as the stock market state):

$$W_0^{f\to c} = P \tag{15}$$

However, suppose instead that the agent has experienced a biased sample (or otherwise has formed a set of associations) in which the depression is over-represented:

$$W_0^{f\to c} = P^* \equiv \begin{bmatrix} \frac{1}{2} & \frac{1}{2} - p^* \\ 0 & p^* \end{bmatrix}. \tag{16}$$

For a Bayesian agent the effect of a distorted prior disappears relatively quickly.

Consider instead the implications of context retrieval. We apply the model described above, with the restriction (for simplicity) of $\zeta = 1$.[10] We thus use the following recursion for context:

$$c_t = \frac{1}{\iota^\top W_{t-1}^{f\to c} f_t} W_{t-1}^{f\to c} f_t, \tag{17}$$

where $W_t^{f\to c}$ evolves according to (10).

Suppose the agent starts with (16). Consider what happens at $t = 1$. A stock market gain retrieves 100% probability on the normal labor income state:

$$c_1 \propto W_0^{f\to c} \begin{bmatrix} 1 \\ 0 \end{bmatrix} \propto \begin{bmatrix} 1 \\ 0 \end{bmatrix}$$

A stock market loss, on the other hand, retrieves a positive probability of a depression,

[10] Allowing for $\zeta \leq 1$ would not change the inference on the unconditional probability of a depression state, but would alter the covariances, which would affect the quantitative conclusions (though not the qualitative ones). One solution is to assume neutral features on stock market returns and the labor market state that are the most common (see Howard and Kahana (2002)). In subsequent examples, we allow for both $\zeta < 1$ and neutral features.

even if one has not occurred:

$$c_1 \propto W_0^{f\to c}\begin{bmatrix} 0 \\ 1 \end{bmatrix} = \begin{bmatrix} \frac{1}{2} - p^* \\ p^* \end{bmatrix} \propto \begin{bmatrix} 1-2p^* \\ 2p^* \end{bmatrix}$$

That is, the agent recalls the depression. The key difference between this model and a rational updating model is that this act of recollecting implies that there is a new "depression" observation in the agent's mental database.

Consider then what happens to the $W^{f\to c}$ matrix:

$$W_1^{f\to c} = \frac{\tau}{1+\tau}W_0^{f\to c} + \frac{1}{1+\tau}c_1 f_1^\top,$$

where

$$c_1 f_1^\top = \begin{cases} \begin{bmatrix} 1 \\ 0 \end{bmatrix}[1,0] = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} & \text{if } gain \\ \begin{bmatrix} 1-2p^* \\ 2p^* \end{bmatrix}[0,1] = \begin{bmatrix} 0 & 1-2p^* \\ 0 & 2p^* \end{bmatrix} & \text{if } loss \end{cases} \tag{18}$$

Regardless of whether a gain or loss occurs, the columns of $W_1^{f\to c}$ relate to those of $W_0^{f\to c}$ by a constant of proportionality. Thus, at time 2, a stock market gain retrieves $[1,0]^\top$, whereas a stock market loss retrieves $[1-2p^*, 2p^*]^\top$. The agent's probability distribution is the same as before.

A formal induction argument (given in Appendix A) shows that, after t periods, given k periods of gains:

$$W_t^{f\to c} = \begin{bmatrix} \frac{1}{2}\frac{\tau}{\tau+t} + \frac{k}{\tau+t} & (\frac{1}{2} - p^*)\frac{\tau}{\tau+t} + (1-2p^*)\frac{t-k}{\tau+t} \\ 0 & p^*\frac{\tau}{\tau+t} + 2p^*\frac{t-k}{\tau+t} \end{bmatrix}. \tag{19}$$

Then

$$\text{plim}_{t\to\infty} W_t^{f\to c} = P^*.$$

It does not matter how much data the agent observes: probabilities remain distorted.

The original matrix over-associates a downturn with a depression due to prior beliefs. What (17) does is recall the depression with every down realization. The act of recalling the depression reinstates the depression context. Thus high probability of depression remains associated with losses in the mind of the agent.

Interestingly, if the agent happened to arrive at the correct probabilities at the beginning, the updating rule (17) would have produced the correct probabilities P.

Figure 3 contrasts implications for three types of agents: the agent who knows the true probability, the agent who starts with an incorrect prior and learns the true probability according to Bayesian updating, and the agent who starts with the same incorrect prior and whose learning is subject to context retrieval. For the purposes of the figure, $p = 0.02$, $p^* = 0.50$, $\sigma = 1$, and $y = 2$.[11] Figure 3 shows the mean of the posterior distribution for p. The Bayesian agent's beliefs converge quickly to something close to the truth. Thus, while precise convergence to a 2% probability of Depression takes many years, updating is very fast for values of the probability that are far from the truth. Twenty years of data suffice to bring the probability sufficiently close so that the resulting portfolio allocation is virtually indistinguishable from that of the full-information agent. On the other hand, the agent who relies on memory does not learn, and can maintain an incorrect probability even in the face of many years of evidence.[12] The point is that memory itself produces a distorted database because the agent relives his worst fears when a stock market downturn occurs.

[11] The Bayesian investor has prior beliefs on p given by $\mathrm{Beta}(p^*\tau, \tau)$, we set $\tau = 2$, implying that prior beliefs are relatively uninformative. Given the likelihood implied by Bernoulli observations on $\tilde{y}$, the posterior is also Beta. We report the mean of this distribution, which is all that is required to (13), since the covariance is linear in the depression probability. Note that the Bayesian agent who infers the correct probability thus behaves the same as the agent who knows the probability for certain; in this economic problem, parameter uncertainty has no effect.

[12] Recent survey evidence (Goetzmann et al., 2017) indicates irrationally high levels of fear of stock market crashes, and that exogenous events can trigger such fears. The latter point is specifically addressed in the model below.

4 Context and the jump back in time: Application to the financial crisis

The failure of Lehman Brothers is widely recognized as a point of inflection in the 2008 financial crisis.[13]

An open question is: why was the failure of Lehman Brothers so pivotal? A growing line of research answers this question by focusing on the importance of financial intermediation to the overall the economy. Brunnermeier and Sannikov (2014) and He and Krishnamurthy (2013) develop models in which the balance sheets of intermediaries are responsible for business cycle fluctuations. However, while it may be necessary to have specialized institutions trade certain complicated investments, it is not clear why the failure of a financial institution should be followed by a broad-based stock market decline. Common stocks are not intermediated assets: trading costs for common stocks, already quite low for the past half-century, have only gotten lower (Jones, 2002). Another possibility is that Lehman represented a sunspot that caused a run on other intermediaries, and other forms of debt (Allen and Gale, 2009; Gorton and Metrick, 2012). Unanswered is why this should cause the stock market to crash, as it did in the fall of 2008, when most companies have very low leverage and can fund themselves through retained earnings?[14].

A third possibility, emphasized by Gennaioli and Shleifer (2018), is that individuals and banks were excessively optimistic, leading to an (observed) explosion in household debt. Indeed, there is some evidence that rapid growth in credit precedes crises (Mian and Sufi, 2009; Schularick and Taylor, 2012, e.g.).[15] Meanwhile, policymakers, while aware (at least shortly) prior to the failure of Lehman brothers of the risks faced by

[13]See, for example, French et al. (2010).

[14]Kahle and Stulz (2013) argue that firms dependent on bank-lending were not unduly affected by the crisis. Gomes et al. (2019b) argue that fluctuations in borrowing conditions are more likely to be affected by investment opportunities than the other way around.

[15]However, causality is not obvious, see Gomes et al. (2019a).

financial institutions, somehow neglected the risk to the overall economy. The failure of Lehman Brothers itself led agents to recall the risks that they had always faced. This possibility is most in the spirit of the model here. However, while related, the two explanations are distinct. In the model below, the increase in actual risk is illusory. While this may seem like an extreme position, in fact, the Great Recession was nothing like the Great Depression.[16] The realized outcome does not seem commensurate with the panic in the fall of 2008.

Our hypothesis is that the financial crisis was a psychological event caused by the failure of Lehman Brothers. The actual realization of a important financial institution failing in the absence of insurance reminded investors of the Great Depression.[17] Some felt that they had – literally – returned to the Great Depression. Investors experienced what the memory literature refers to as a *jump back in time* (Manning et al., 2011b; Howard et al., 2012a). Once this feeling entered the discourse, it proved hard to shake. Subsequent events showed that in fact there was no Great Depression. This was only revealed, though, over time. Somehow, what emerged from the crisis and recession was not a feeling of relief but rather a renewed emphasis on the fragility of the financial sector and the possibility that a Great Depression might in fact occur. The model below formalizes this intuition.

4.1 Asset prices for a given rare event probability

Consider a representative agent facing a probability p of a rare event. We follow Barro (2006) and assume that the endowment process equals

$$\log C_{t+1} = \log C_t + \mu + u_{t+1} + v_{t+1}, \tag{20}$$

[16]The effect of the financial crisis on aggregate consumption was relatively minor: from the start of 2008 to the end of 2009, aggregate consumption fell by 3%, and consumption began to recover by 2010. In contrast, consumption fell by 16% in the Great Depression.

[17]See, for example, the reporting of *The Guardian* on the day's events: https://www.theguardian.com/business/2008/sep/15/marketturmoil.stockmarkets.

where u_{t+1} and v_{t+1} are independent, $u_{t+1} \sim N(0, \sigma^2)$ and

$$v_{t+1} = \begin{cases} 0 & \text{with prob. } e^{-p} \\ \log(1-\xi) & \text{with prob. } 1 - e^{-p} \end{cases} \tag{21}$$

where ξ is a random variable with support on $[0, 1)$. We consider an asset paying dividends that satisfy

$$\log D_{t+1} = \log D_t + \mu + u_{t+1} + \lambda v_{t+1} \tag{22}$$

with $\lambda > 1$. This assumption captures the fact that dividends fall by more than consumption during financial disasters (Longstaff and Piazzesi, 2004).

We assume that at every period, the agent maximizes utility

$$E_t \sum_{s=1}^{\infty} \beta^s \log C_s$$

If P_t is the claim to the cash flows (41), the first-order conditions of the agent imply

$$P_t = E_t \left[M_{t+1}(P_{t+1} + D_{t+1}) \right], \tag{23}$$

where the intertemporal marginal rate of substitution equals

$$M_{t+1} = \beta \left(\frac{C_{t+1}}{C_t} \right)^{-\gamma},$$

As in the standard representative agent endowment framework (Lucas, 1978), equilibrium requires that optimal consumption satisfy (20) and (21) and that cash flows equal (41). Asset prices adjust to satisfy the first-order conditions for the representative agent.

In Appendix B, we show (23) has solution

$$\frac{P_t}{D_t} = \sum_{n=1}^{\infty} \Phi(p)^n = \frac{\Phi(p)}{1-\Phi(p)}.$$

for

$$\Phi(p) = \beta\left(e^{-p} + (1-e^{-p})E\left[(1-\xi)^{\lambda-\gamma}\right]\right) \tag{24}$$

When $\lambda > 1$, an increase in p (in a comparative statics sense), lowers the price.

The riskfree rate also solves first-order condition, and equals

$$\begin{aligned} 1 + r_f &= E\left[M_{t+1}\right]^{-1} \\ &= \beta^{-1} e^{\mu - \frac{1}{2}\sigma^2}\left(e^{-p} + (1-e^{-p})E\left[(1-\xi)^{-\gamma}\right]\right)^{-1}, \end{aligned}$$

An increase in p lowers the riskfree rate; fear of a disaster leads the investor to want to save for the future.

4.2 Memory for rare events

We identify features with the state of the financial system, so $f_t = [1,0]^\top$ represents normal times and $f_t = [0,1]^\top$ represents a crisis. We identify context with the state of the underlying economy, so $c_t = [1,0]^\top$ represents normal times, and $c_t = [0,1]^\top$ represents a depression state.

Using the interpretation of prior associations as probabilities, we can write

$$W_0^{f\to c} = \begin{bmatrix} \text{Prob}(\textit{no crisis \& no depression}) & \text{Prob}(\textit{crisis \& no depression}) \\ \text{Prob}(\textit{no crisis \& depression}) & \text{Prob}(\textit{crisis \& depression}) \end{bmatrix} \tag{25}$$

$$= \begin{bmatrix} 1-p^c & p^c(1-q) \\ 0 & p^c q \end{bmatrix}, \tag{26}$$

where

$$
\begin{aligned}
p^c &= \text{probability of a financial crisis} \\
q &= \text{probability of an economic disaster, given a financial crisis,}
\end{aligned}
$$

namely, in investors' minds, economic disaster is always accompanied by crisis

We connect context to asset prices by assuming, for simplicity, homogeneous investors aggregating to the representative agent of the previous section. Agents extract probabilities of a disaster from c_t (the probability is the second element of c_t), and view these probabilities (again, for simplicity) as permanent.

We assumed the generalized context evolution (6), with context retrieval (8). We assume that in the recent past, only neutral features have been observed, namely $f_t = [1, 0]^\top$. Assuming that the features-to-context matrix is in the steady state given by (25), neutral features imply the neutral context:

$$
c_t^{\text{in}} \propto W_0^{f \to c} \begin{bmatrix} 1 \\ 0 \end{bmatrix} \propto \begin{bmatrix} 1 \\ 0 \end{bmatrix}. \tag{27}
$$

Given sufficiently many observations of neutral features, context reaches a steady state value $c_t = [1, 0]^\top$, as follows from setting $c_t = c_{t-1}$ in (6).[18] The model thus implies neglected risk (Gennaioli and Shleifer, 2018).

Though agents neglect the depression state is neglected, they have not forgotten it. Representing the failure of Lehman brothers is $f_1 = [0, 1]^\top$, the well-publicized failure

[18] The discussion thus far assumes $W_t^{f \to c}$ remains fixed. However, even if we were to allow for updating this matrix, as we do below, it would not change (27). This is because only the relative weights on the elements would change. (27) does not depend on the weights, but only on the fact that a depression and a crisis cannot co-occur.

of a major financial institution. It follows that

$$c_1^{\text{in}} \propto W_0^{f\to c} \begin{bmatrix} 0 \\ 1 \end{bmatrix} \propto \begin{bmatrix} 1-q^* \\ q^* \end{bmatrix}. \tag{28}$$

Equation 28 represents reinstatement of the depression context. Even though a depression has not occurred, the agent is reminded strongly of a depression because of the financial crisis. The stronger the association between depression and crisis (the higher is q), the more strongly is the depression context reinstated.

Because context is autoregressive, neutral features are still present. The retrieved depression context mixes with the prior neutral context to form

$$c_1 = \rho \begin{bmatrix} 1 \\ 0 \end{bmatrix} + \zeta \begin{bmatrix} 1-q \\ q \end{bmatrix} \tag{29}$$

The probability of a depression changes from zero to ζq, causing an immediate decline in stock prices and in the riskfree rate.

What does the model say about the time path of context, and hence that prices and the riskfree rate, following the event? We discuss in detail one such possible path in which the agent observes only neutral features. To understand the time path of context, we need to understand how features affect memory Because of the assumption of a long prior sample, these affects are small at the beginning, but accumulate over time.

As in the previous section, tet τ be the length of the prior sample. The features-to-context matrix updates as follows:

$$
\begin{aligned}
W_1^{f\to c} &= \frac{\tau}{\tau+1}W_0^{f\to c} + \frac{1}{\tau+1}c_1 f_1^\top \\
&= \frac{\tau}{\tau+1}\begin{bmatrix} 1-p^c & p^c(1-q) \\ 0 & p^c q \end{bmatrix} + \frac{1}{\tau+1}\begin{bmatrix} 0 & \rho+\zeta(1-q) \\ 0 & \zeta q \end{bmatrix}
\end{aligned} \tag{30}
$$

Regardless of whether or not a depression actually occurs, the agent updates memory, represented by $W_1^{f\to c}$, with a partial observation of a depression, co-occurring with crisis (we emphasize this effect in Section 3).

Following the crisis, we assume the agent observed neutral features. These neutral features lead to reinstatement of the neutral context one period after the crisis:

$$
c_2^{\text{in}} \propto W_1^{f\to c}\begin{bmatrix} 1 \\ 0 \end{bmatrix} = \begin{bmatrix} 1 \\ 0 \end{bmatrix} \tag{31}
$$

Because the agent observes only neutral features, the crisis does not affect c_2^{in}, because only entries associated with crisis features have been updated. In particular, the agent does not yet link a depression with neutral features. Neutral features reinstate the neutral context. It follows from (31) and (6) that context at time 2 equals:

$$
c_2 = \rho c_1 + \zeta\begin{bmatrix} 1 \\ 0 \end{bmatrix} = \begin{bmatrix} \rho^2+\rho\zeta(1-q)+\zeta \\ \rho\zeta q \end{bmatrix}. \tag{32}
$$

The (partial) depression context decays, and the new context is a weighted average of this depression context and the reinstated neutral context. Now, however, given (32), depression is mixed with neutral features in semantic memory:

$$\begin{aligned} W_2^{f\to c} &= \frac{\tau+1}{\tau+2}W_1^{f\to c} + \frac{1}{\tau+2}c_2 f_2^\top \\ &= \frac{\tau}{\tau+2}\begin{bmatrix} 1-p^c & p^c(1-q) \\ 0 & p^c q \end{bmatrix} + \frac{1}{\tau+2}\begin{bmatrix} 0 & \rho+\zeta(1-q) \\ 0 & \zeta q \end{bmatrix} \\ &\quad + \frac{1}{\tau+2}\begin{bmatrix} \rho^2+\rho\zeta(1-q)+\zeta & 0 \\ \rho\zeta q & 0 \end{bmatrix}. \end{aligned}$$

where we substitute in for $W_1^{f\to c}$ using (30). The partial depression context appears twice, first combined with the financial crisis and second (and down-weighted by ρ), combined with neutral features. These latter two terms are small assuming a large previous sample, but they are not negligible, and the recursive structure of memory implies that they exert permanent effects.

Thus, there are two opposing effects on memory two years following the financial crisis, assuming the agent continues to observe neutral features. One is that the memory of the crisis decays further. The second is that the presence of the depression context in semantic memory, mixed with neutral features implies that even neutral features do not recall 100% probability of no depression. Consider the context now reinstated by neutral features:

$$c_3^{\text{in}} \propto W_2^{f\to c}\begin{bmatrix} 1 \\ 0 \end{bmatrix} \propto \frac{\tau}{\tau+2}\begin{bmatrix} 1-p^c \\ 0 \end{bmatrix} + \frac{1}{\tau+2}c_2$$

Unlike the reinstated context from the previous period, this context puts some weight on the depression. Context three years following the crisis thus equals

$$c_3 = \rho c_2 + \zeta c_3^{\text{in}}.$$

The crisis has an ever-decreasing influence on beliefs. However, the decay is slower than exponential, as the memory of the crisis continues to exert a pull on the agent through the features-to-context matrix.

This discussion illustrates the forward and backward-looking nature of memory in this model. Going into the crisis, investors extrapolate based on a period where nothing much has happened. However, patterns stored in memory imply that beliefs can respond in unexpected ways to events. In this example, the association of a financial crisis with a depression leads agents to forecast a depression. The memory of a depression then is reinstated, affecting future choices.

As an example calibration, we assume $p^c = .025$, $q = 0.5$, $q = 1$, a decline in aggregate consumption of 20%, risk aversion of 1 (log utility) and leverage of 2. Here, and in the applications that follow, we assume $\rho = 0.8$, in the range of values typically found in the memory literature (Lohnas et al., 2015). We also assume that the agent has observed a prior sample of 100 years, so $\tau = 100$.

Figure 4 shows the time path of the price-dividend ratio and the real riskfree rate under the model described above. The crisis leads to a jump back in time, reinstating the depression context and implying an immediate decline in prices (recall that there is no direct effect of the crisis on stock prices), and an immediate decline in the interest rate due to the precautionary savings motive. Both were of features of the 2008 financial crisis. Prices and interest rates gradually recover as neutral features are observed. Prices and interest rates return to values above the full information case after about 15 years. However, though they exhibit neglected risk, they remain depressed even 25 years following the crisis, due to the memory of the depression.

5 Sticky context as an explanation for price momentum

We show how context theory provides an explanation for the momentum effect. Price momentum is the finding that stocks with the best past price performance outperform, over the next year, stocks with the worst performance, by a wide margin. This stocks with the best performance are known as "winners" with the others as "losers." When sorted into deciles, the WML (winners-minus-losers) portfolio, formed on the extreme deciles, generates an annual return of about 11% (Jegadeesh and Titman, 1993).[19]

We assume a simple neoclassical production model in which productivity, earnings, and dividends all covary up to a scale factor. Firms can experience high or low productivity growth. They can also experience high or low observed earnings. The joint contingency matrix, in which observed earnings growth varies along the columns, and the true growth state varies along the rows is:

$$P = \begin{bmatrix} \Pr\left(\frac{E_{t+1}}{E_t} = g_H \ \&\ g = g_H\right) & \Pr\left(\frac{E_{t+1}}{E_t} = g_L \ \&\ g = g_H\right) \\ \Pr\left(\frac{E_{t+1}}{E_t} = g_H \ \&\ g = g_L\right) & \Pr\left(\frac{E_{t+1}}{E_t} = g_L \ \&\ g = g_L\right) \end{bmatrix},$$

where we use the shorthand E_{t+1}/E_t to denote growth in earnings from one period to the next.

These growth rates are permanent, and, for simplicity, there is no noise, so high observed growth is uniquely associated with high true growth.[20] There is a measure of

[19] The model presented here can be seen as a foundation for under-reaction, which is studied by Daniel et al. (1998) and Barberis et al. (1998), or for slow information diffusion (Hong and Stein, 1999).

[20] Useful generalizations would be to assume a Markov switching model, and/or have an iid shock (or potentially measurement error) create a wedge between observed and true earnings. For our basic qualitative result, this model is sufficient however.

p of high growth firms. Therefore

$$P = \begin{bmatrix} p & 0 \\ 0 & 1-p \end{bmatrix}. \tag{33}$$

We associate context with the long-run productivity growth rate of the firm, and features with high or low observed earnings growth. Investors observe a large sample of firms over a long time period. They store the relation between features and context in a stable features-to-context matrix, in which the high growth context is overwhelmingly associated with high observed growth, and the low growth context with low observed growth. These assumptions will greatly simplify our analysis below.[21] That is

$$W^{f\to c} = \sum_{t=-\tau}^{0} \sum_{i=1}^{N} c_{it} f_{it}^{\top} \approx P. \tag{34}$$

where i indexes firms and t indexes time, and where (as before), without loss of generality, we assume we are currently at time 0.

Now consider a set of firms indexed by j ($j = 1, \ldots, N$) for which investors do not yet know the growth rate. At time t, some of these firms will be assigned the high growth rate, while the complement will be assigned the low growth rate. Those that are assigned the high growth rate are "winners" (in the sense of Jegadeesh and Titman (1993)) whereas those assigned the low growth rate are "losers." This terminology will be justified shortly. The total set of firms is sufficiently small so that the estimation of the $W^{f\to c}$ matrix is unaffected.

Assume that, as in the model in Section 4, investors do not accurately forecast the time path of context. Rather, they price assets using the context vector to assign probabilities to permanent states. Assume risk-neutral investors and an exogenous

[21] Accounting for short-term context dynamics in the $W^{f\to c}$ matrix would complicate the analysis without changing the conclusions.

discount rate r, such that $r > g_H > g_L$. If investors assign subjective probability $\tilde{p}_{jt}$ at time t to firm j having high growth, then the price-dividend ratio for firm j is given by

$$\begin{aligned}
\frac{P_{jt}}{D_{jt}} &= \tilde{p}_{jt} E\left[\sum_{s=1}^{\infty}(1+r)^{-s}\frac{D_{j,t+s}}{D_{jt}} \mid g = g_H\right] + (1-\tilde{p}_{jt}) E\left[\sum_{s=1}^{\infty}(1+r)^{-s}\frac{D_{j,t+s}}{D_{jt}} \mid g = g_L\right] \\
&= \tilde{p}_{jt}\sum_{s=1}^{\infty}\left(\frac{1+g_H}{1+r}\right)^s + (1-\tilde{p}_{jt})\sum_{s=1}^{\infty}\left(\frac{1+g_L}{1+r}\right)^s \\
&= \tilde{p}_{jt}\frac{1+g_H}{r-g_H} + (1-\tilde{p}_{jt})\frac{1+g_L}{r-g_L}. \qquad (35)
\end{aligned}$$

We assume, for all firms j, that

$$c_{j0} = \begin{bmatrix} p \\ 1-p \end{bmatrix}$$

Winning firms have high productivity growth rates, and so experience $f_{j1} = [1, 0]^\top$. Losing firms have low productivity growth rate, and so experience $f_{j1} = [0, 1]^\top$. It follows from the features-to-context matrix (34) that $c_{j1}^{\text{in}} = f_{j1}$ for all j. Applying (6), winners have time-1 context of

$$c_{j1} = \begin{bmatrix} \rho p + 1 - \rho \\ \rho(1-p) \end{bmatrix}.$$

Because $\rho p + 1 - \rho > p$, the subjective probability of high growth has increased. Losers have time-1 context of

$$c_{j1} = \begin{bmatrix} \rho p \\ \rho(1-p) + 1 - \rho \end{bmatrix}.$$

Because $\rho p < p$, the subjective probability of high growth has decreased.

Finally note that for stock j, the price change is given by

$$\frac{P_{j1}}{P_{j0}} = \frac{D_{j1}}{D_{j0}} \frac{P_{j1}/D_{j1}}{P_{j0}/D_{j0}}. \tag{36}$$

It follows from (35), that the price-dividend ratio is an increasing function of $\tilde{p}$. Moreover, high growth firms also have higher realized values of dividend growth $\frac{D_{j1}}{D_{j0}}$. The price change (36) is therefore higher for winners than for losers, justifying the terminology.

The high-growth features of time-1 is what makes the winning stocks "winners" and the losing stocks "losers." It is not surprising that positive earnings surprises should lead investors to update their probabilities of a stock having high long-run earnings growth. The key implication of context dynamics (6) is that this updating is incomplete.

Indeed, in subsequent periods, context puts ever-increasing weight on the high-growth state for winning firms and ever-decreasing weight for losing firms. Eventually the context vectors converge, with convergence being faster for lower values of ρ. When convergence has been reached, returns on the two firms are the same. Before this limit, however, returns on winning firms will always be above returns on losing firms.

Figure 5 shows the prices and returns on winners and losers and constrasts the implications of retrieved-context theory with Bayesian updating (in this case, Bayesian updating is identical to full information). The key assumption in this model that differentiates it from Bayesian updating is the slow evolution of context implied by (6). In the simple model outlined above, a Bayesian investor should update immediately to 100% probability in the high growth state upon observing one observation. More generally, beliefs under Bayesian updating should follow a martingale. One cannot have positive "surprises" systematically following positive events. Yet that is what the data appear to show (Chan et al., 1996).

From a psychological perspective, the reason updating is slow is that c_0, representing a low probability of a high growth state is still in investors' context when they evaluate the new information. Repeated observations are required before c_0 disappears from context.

6 Fear and asset allocation

In this section, we apply retrieved-context theory to explain a change portfolio holdings in response to a stimulus, such as viewing a horror movie. We will need to modify the simple model required to explain the effect of the Lehman Brothers bankruptcy on stock market valuations of the previous section. In the previous section, an stimulus triggered a jump back in time, namely a sudden jump in beliefs regarding the probability of a Great Depression. This stimulus had previously been directly associated with a Great Depression. However, in that case, agents could have believed that the true probability of a Great Depression had changed. In this section, where we seek to explain experimental evidence, subjects were told explicitly that risks had not changed, and yet there was a change in portfolio choice.

Guiso et al. (2018) observe that watching a horror movie influences the risk premium investors require to hold risky assets (see also Cohn et al. (2015)). What is striking about this experiment is that fear alone, as opposed to new information, has a substantial effect on risk taking. We hypothesize that fear operates through the memory channel. As we have shown, the context-retrieval mechanism allows negative associations to have both a short-lived effect (through the autoregressive structure) and a highly persistent effect (through the features to context matrix). It will be the first that that is the focus of this section.

We assume that the feature state can consist of the presence of danger, which may or may not be associated with a financial crisis. Danger is evoked by the kind of movie that Guiso et al. (2018) showed in the experiment. The feature space consists of:

$$f_t = \begin{cases} e_1 & \text{if} \quad \text{no danger \& no crisis} \\ e_2 & \text{if} \quad \text{danger \& no crisis} \\ e_3 & \text{if} \quad \text{danger \& crisis} \end{cases}$$

where e_j is the jth basis vector. We assume a two-dimensional context vector, depending on whether the underlying state represents a high level of risk or a low level of risk. We refer to $f_t = e_1$ as neutral features.

As in Section 3, we consider the portfolio choice problem of an agent investing in a risky asset and a riskless asset. We let $\tilde{r}$ denote the risky asset return, and π the percent allocation to the risky asset. Without loss of generality, we assume the agent starts the period with financial wealth equal to one, so that end-of-period financial wealth equals $1 + \pi\tilde{r}$. Similarly to Section 3, the agent also faces the possibility of a negative labor market outcome, which we denote by $\tilde{y}$. We can think of $\tilde{y}$ as expenditures for health expenses, or inability to meet other financial obligations (such as a mortgage) due to a poor labor income shock. To summarize, the agent solves

$$\max_{\pi} Eu(1 + \pi\tilde{r} + \tilde{y}). \tag{37}$$

We model $\tilde{y}$ as a Bernoulli process:

$$\tilde{y} = \begin{cases} 0 & \text{with probability } 1 - p \\ -\xi & \text{with probability } p, \end{cases}$$

with $\xi \in [0, 1]$. We assume $\tilde{r}$ also takes on two possible outcomes (each with equal probability), and has mean μ and standard deviation σ. Unlike the model in Section 3, $\tilde{y}$ and $\tilde{r}$ are uncorrelated.

In Section 3, agents lowered their allocation to stocks in response to fear about a depression state. As a mechanism, the results relied on the covariance between a negative labor income realization and the stock return. In this section, we hypothesize that agents experience fear of physical danger after watching the horror movie. However, we cannot rely on covariance between physical danger and the risky asset return (which would be implausible) to generate the decreased investment in the risky asset.

To allow $\tilde{y}$ to affect the agent's portfolio choice, we assume CRRA utility, as in Section 4. This is the standard specification in asset pricing more generally. Let

$$u(x) = \frac{x^{1-\gamma} - 1}{1-\gamma} \tag{38}$$

(with $\gamma = 1$ corresponding to log utility). This functional form is equivalent to the specification that the agent has CRRA. We further assume $\gamma \geq 1$ (and will focus on $\gamma = 1$). Fixing an outcome $\tilde{y} = y$, (38) implies decreasing relative risk aversion; the agent is very averse to declines in wealth that are close to ξ. Now allowing $\tilde{y}$ to be variable, the greater is the probability, the more weight the agent places on this possible outcome in his decision-making. Thus this formulation, together with the context dynamics below, endogenizes time-varying risk aversion. It also endogenously produces a role for emotional state in the utility function, as suggested by Loewenstein (2000).

As in previous examples, assume the context vector determines the agent's subjective risk probability p. Assume $c_t = [1, 0]^\top$ corresponds to zero risk probability $c_t = [0, 1]^\top$ is probability 1 of risk. The matrix P representing the joint contingencies equals

$$P = \begin{bmatrix} \Pr(nr, nd, nc) & \Pr(nr, d, nc) & \Pr(nr, d, c) \\ \Pr(r, nd, nc) & \Pr(r, d, nc) & \Pr(r, c, d) \end{bmatrix}, \tag{39}$$

where r stands for risk, d for disaster, c for crisis, and nr, nd, nc the lack of risk, disaster, or crisis respectively. Let p be the unconditional probability of danger, and

q be the probability of crisis given danger (we make the reasonable assumption that crisis is always accompanied by danger, which can include dangers inherent in loosing one's money).[22] We assume, for simplicity, that risk and danger always co-occur, so that

$$P = \begin{bmatrix} 1-p & 0 & 0 \\ 0 & p(1-q) & pq \end{bmatrix} \tag{40}$$

however, our results do not depend on this assumption.

We assume context follows (6) and (8). Without loss of generality, label the time of the experiment as $t = 1$, so $t = 0$ refers to the context prior to the experiment. We assume that the agent has accurately observed and recalled a sufficiently long sample, so that

$$W_0^{f \to c} = P$$

While we could assume that the agent has seen only neutral features for a long time, this seems unattractive in the context of an experiment in which there may be a variety of experiences. We therefore assume that the agent begins with the correct probabilities, namely

$$c_0 = \begin{bmatrix} 1-p \\ p \end{bmatrix}$$

As will be clear, our results are robust to variation in this assumption, as all that is required is that the stimulus introduced in the experiment drives context sufficiently far away from the steady state.

[22]This structure does have the implication that stock returns are uncorrelated with the crisis outcome. A richer model might have two types of risk which share a common component of $\tilde{y}$, but one with an additional component that correlates with stock returns.

The stimulus represents $f_1 = e_2$, namely, danger without crisis.[23] We have

$$c_1^{\text{in}} \propto W_0^{f\to c} f_1 = W_0^{f\to c} e_2$$

which implies

$$c_1^{\text{in}} = \begin{bmatrix} 0 \\ 1 \end{bmatrix}.$$

Therefore, the new context is:

$$c_1 = \rho c_0 + \begin{bmatrix} 0 \\ \zeta \end{bmatrix},$$

so that the subjective probability of the risky state rises from p to $\rho p + \zeta$.[24]

Figure 6 shows expected utility (37) as a function of π, prior to and after the stimulus. We assume the following parameters: $\mu = 4\%$, $\sigma = 20\%$, $p = 2\%$, $\xi = -0.8$, and $\gamma = 1$ (log utility). As elsewhere, $\rho = 0.8$. When the agent has the correct probabilities, the portfolio allocation equals 70%, falling to 30% after the stimulus. Note that the model would imply the same shift for a financial crisis. This accounts for the finding of Guiso et al. (2018) that (a) viewing a horror movie and (b) exposure to a financial crisis increases effective risk aversion.

[23]One might object that a movie is not the same as actual physical danger. The model accommodates this difference, however, in that the movie is purely transitory, so that context (because of the autoregressive term) does not fully shift; actual danger would presumably be more persistent. A related objection is that, the experiment implies that neutral features co-occur with a heightened risk context, and danger occurs when the probability of a risk context is not equal to 1. Should this have occurred in the past, then (40) would no longer represent the features-to-context matrix. By assuming that the agent does not exhibit neglected risk, we implicitly allow for this. While the model does imply that viewing a horror movie sufficiently frequently de-sensitizes the agent, Guiso et al. (2018) specifically avoid this problem by choosing a movie that is both intense and obscure.

[24]Note that $\rho p + \zeta > p$ for $p < 1$, because $\rho p + \zeta$ is a weighted average of p and 1.

7 Conclusion

Our past experiences, and our knowledge about the world, constitute a vast database of information that potentially informs every decision we make. Does the human memory system discard most of this information to abstract a small, and possibly biased, subset? Modern research on human memory supports an alternative view in which much of our past information remains in storage, to be retrieved based on a latent dynamic context (Kahana, 2012). According to this view, context updates recursively; features of the environment evoke past contextual states via associative memory. These associations then are permanently stored to be themselves evoked at later times. Thus past contextual states drive the evolution of context itself.

Here we introduce memory into the decision problem of an economic agent, through a formal model of retrieved context theory. Features represent observed stock prices or exceptionally salient news such as a large bank failure. The associative matrices linking context to features draw out the agent's beliefs given these observations. Our model allows for important deviations from Bayesian updating, such as the influence of events in the distant past, the influence of irrelevant events, and slow updating to new information. We apply retrieved-context theory to four illustrative problems in financial economics: the effects of life experience on choices, the sudden onset of a financial crisis, the appearance of momentum in stock returns, and time-variation in risk aversion due to exogenous factors.

A Proof of stable associations in the persistent belief model

We give a formal induction argument for (19), where k is the number of stock market gains from $s = 1, \ldots, t$. Note that (19) holds for $t = 0$ by definition. Assume (19) holds for t; we show it holds for $t + 1$.

In the case of a stock market gain at $t + 1$:

$$c_{t+1} \propto W_t^{f \to c} \begin{bmatrix} 1 \\ 0 \end{bmatrix} \propto \begin{bmatrix} 1 \\ 0 \end{bmatrix}.$$

In the case of a stock market loss,

$$\begin{aligned} c_{t+1} &\propto W_t^{f \to c} \begin{bmatrix} 0 \\ 1 \end{bmatrix} \\ &= \begin{bmatrix} (\frac{1}{2} - p^*)\frac{\tau}{\tau+t} + (1 - 2p^*)\frac{t-k}{\tau+t} \\ p^* \frac{\tau}{\tau+t} + 2p^* \frac{t-k}{\tau+t} \end{bmatrix} \\ &\propto \begin{bmatrix} (1 - 2p^*)(\tau/2 + t - k) \\ 2p^*(\tau/2 + t - k) \end{bmatrix} \\ &\propto \begin{bmatrix} 1 - 2p^* \\ 2p^* \end{bmatrix} \end{aligned}$$

where the normalization in the last expression ensures $\iota^\top c_{t+1} = 1$. We thus have

$$c_{t+1} f_{t+1}^\top = \begin{cases} \begin{bmatrix} 1 \\ 0 \end{bmatrix} [1, 0] = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} & \text{if } \textit{gain} \\ \begin{bmatrix} 1 - 2p^* \\ 2p^* \end{bmatrix} [0, 1] = \begin{bmatrix} 0 & 1 - 2p^* \\ 0 & 2p^* \end{bmatrix} & \text{if } \textit{loss} \end{cases}$$

Then

$$W_{t+1}^{f\to c} = \frac{t+\tau}{t+\tau+1} W_t^{f\to c} + \frac{1}{t+\tau+1} c_{t+1} f_{t+1}^{\top}$$

The result follows.

B Asset pricing with rare events

The model in this section follows that of Barro (2006). We assume a complete-markets endowment economy similar to Lucas (1978). Assume

$$\log C_{t+1} = \log C_t + \mu + u_{t+1} + v_{t+1},$$

where u_{t+1} and v_{t+1} are independent, $u_{t+1} \sim N(0, \sigma^2)$ and

$$v_{t+1} = \begin{cases} 0 & \text{with prob. } e^{-p} \\ \log(1-\xi) & \text{with prob. } 1 - e^{-p} \end{cases}$$

where ξ is a random variable with support on $[0, 1)$. We further assume that the dividend satisfies

$$\log D_{t+1} = \log D_t + \mu + u_{t+1} + \lambda v_{t+1} \tag{41}$$

with $\lambda > 1$. This assumption captures the fact that dividends fall by more than consumption during crisis (Longstaff and Piazzesi, 2004).

We assume that at every period, the agent maximizes utility

$$E_t \sum_{s=1}^{\infty} \beta^s \frac{C_s^{1-\gamma}}{1-\gamma}.$$

Note that p scales with the time interval. Thus we can make p arbitrarily small (without changing the underlying economics) by considering smaller and smaller time intervals (in effect approximating a Poisson process in discrete time). Note, however, that b is a fixed quantity as the time interval shrinks. Besides the closed-form expressions, we will give simpler formulas using

$$1 - e^{-p} \approx p,$$

and for x close to zero,

$$\log(1+x) \approx x.$$

Define the stochastic discount factor as

$$M_{t+1} = \beta \left(\frac{C_{t+1}}{C_t} \right)^{-\gamma}$$

It follows from the first-order condition of the representative agent that

$$\begin{aligned} 1 + r_f &= E\left[M_{t+1}\right]^{-1} \\ &= \beta^{-1} E\left[\exp\left\{-\gamma(\mu + u_{t+1} + v_{t+1})\right\}\right]^{-1} \\ &= \beta^{-1} e^{\gamma\mu - \frac{1}{2}\gamma^2\sigma^2} \left(e^{-p} + (1-e^{-p}) E\left[(1-\xi)^{-\gamma}\right]\right)^{-1}, \end{aligned}$$

so that

$$\log(1+r_f) \approx -\log\beta + \gamma\mu - \frac{1}{2}\gamma^2\sigma^2 - pE\left[(1-\xi)^{-\gamma} - 1\right].$$

The price P_{1t} of a claim to a one-period equity strip satisfies the equation

$$P_{1t} = E_t\left[M_{t+1} D_{t+1}\right].$$

It is straightforward to solve for this price by using the normalization

$$\begin{aligned}
\frac{P_{1t}}{D_t} &= E_t\left[M_{t+1}\frac{D_{t+1}}{D_t}\right] \\
&= \beta e^{(1-\gamma)\mu+\frac{1}{2}(1-\gamma)^2\sigma^2}\left(e^{-p}+(1-e^{-p})E\left[(1-\xi)^{\lambda-\gamma}\right]\right).
\end{aligned}$$

It is convenient to define the notation

$$\Phi(p) = \beta e^{(1-\gamma)\mu+\frac{1}{2}(1-\gamma)^2\sigma^2}\left(e^{-p}+(1-e^{-p})E\left[(1-\xi)^{\lambda-\gamma}\right]\right) \tag{42}$$

as the price-dividend ratio for the one-period claim.

Taking the log of both sides of (42) gives a convenient approximation

$$\begin{aligned}
\log\Phi(p) &= \log\beta+(1-\gamma)\mu+\frac{1}{2}(1-\gamma)^2\sigma^2+\log\left(e^{-p}+(1-e^{-p})E\left[(1-\xi)^{\lambda-\gamma}\right]\right) \\
&\approx -\log\beta+(1-\gamma)\mu+\frac{1}{2}(1-\gamma)^2\sigma^2-pE\left[1-(1-\xi)^{\lambda-\gamma}\right]. \qquad (43)
\end{aligned}$$

Note that $1-\xi \in (0,1]$. Thus the term inside the expectation in (43) is positive if and only if $\lambda > \gamma$. Under these circumstances, an increase in p lowers prices.

Now consider the claim to a stream of dividends following process (41). Let P_t denote the price of this claim. The condition for equilibrium, applied to this claim implies

$$P_t = E_t\left[M_{t+1}(P_{t+1}+D_{t+1})\right]$$

which in turn implies a recursion for the price-dividend ratio

$$\frac{P_t}{D_t} = E_t\left[M_{t+1}\frac{P_{t+1}/D_{t+1}+1}{P_t/D_t}\frac{D_{t+1}}{D_t}\right].$$

This equation is solved by

$$\frac{P_t}{D_t} = \sum_{n=1}^{\infty} \Phi(p)^n = \frac{\Phi(p)}{1 - \Phi(p)}.$$

The model is calibrated using $\mu = 0$, $\sigma = 2\%$, $\beta = e^{-.03}$, $\gamma = 1$, $\lambda = 3$, $\xi = 0.40$.

References

Abernathy, E. M. (1940). The effect of change in environmental conditions upon the results of college exams. *Journal of Psychology*, 10:293–301.

Allen, F. and Gale, D. (2009). *Understanding Financial Crises*. Clarendon lectures in finance. Oxford University Press.

Arrow, K. J. (1971). *Essays in the Theory of Risk-Bearing*. Markham Publishing Co, Chicago.

Atkinson, R. C. and Shiffrin, R. M. (1968). Human memory: A proposed system and its control processes. In Spence, K. W. and Spence, J. T., editors, *The Psychology of Learning and Motivation*, volume 2, pages 89–105. Academic Press, New York.

Barberis, N., Shleifer, A., and Vishny, R. (1998). A model of investor sentiment. *Journal of Financial Economics*, page 37.

Barberis, N. C. (2013). Thirty years of prospect theory in economics: A review and assessment. *Journal of Economic Perspectives*, 27(1):173–196.

Barro, R. J. (2006). Rare disasters and asset markets in the twentieth century. *Quarterly Journal of Economics*, 121(3):823–866.

Bordalo, P., Gennaioli, N., and Shleifer, A. (2017). Memory, attention, and choice. NBER Working Paper No. 23256.

Brunnermeier, M. K. and Sannikov, Y. (2014). A macroeconomic model with a financial sector. *American Economic Review*, 104(2):379–421.

Carr, H. A. (1931). The laws of association. *Psychological Review*, 38:212–228.

Chan, L. K. C., Jegadeesh, N., and Lakonishok, J. (1996). Momentum strategies. *The Journal of Finance*, 51(5):1681–1713.

Cohn, A., Engelmann, J., Fehr, E., and Marechal, M. A. (2015). Evidence for countercyclical risk aversion: An experiment with financial professionals. *The American Economic Review*, 105(2):860–885.

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

Ebbinghaus, H. (1885/1913). *Memory: a contribution to experimental psychology*. Teachers College, Columbia University, New York.

Ellsberg, D. (1961). Risk, ambiguity, and the savage axioms. *The Quarterly Journal of Economics*, 75(4):643–669.

Epstein, L. G. and Schneider, M. (2003). Recursive multiple-priors. *Journal of Economic Theory*, 113(1):1–31.

Estes, W. K. (1955). Statistical theory of spontaneous recovery and regression. *Psychological Review*, 62:145–154.

Folkerts, S., Rutishauser, U., and Howard, M. (2018). Human episodic memory retrieval is accompanied by a neural contiguity effect. *Journal of Neuroscience*, 38(17):4200–4211.

French, K. R., Baily, M. N., Campbell, J. Y., Cochrane, J. H., Diamond, D. W., Duffie, D., Kashyap, A. K., Mishkin, F. S., Rajan, R. G., Scharfstein, D. S., Shiller, R. J., Shin, H. S., Slaughter, M. J., Stein, J. C., and Stulz, R. M. (2010). *The Squam Lake Report: Fixing the Financial System*. Princeton University Press.

Gennaioli, N. and Shleifer, A. (2018). *A Crisis of Beliefs: Investor Psychology and Financial Fragility*. Princeton University Press, Princeton, NJ.

Gilboa, I. and Schmeidler, D. (1989). Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18(2):141–153.

Gilboa, I. and Schmeidler, D. (1995). Case-based decision theory. *The Quarterly Journal of Economics*, 110(3):605–639.

Godden, D. R. and Baddeley, A. D. (1975). Context-dependent memory in two natural environments: On land and under water. *British Journal of Psychology*, 66:325–331.

Goetzmann, W. N., Kim, D., and Shiller, R. J. (2017). Crash beliefs from investor surveys. NBER Working Paper No. 22143.

Gomes, J. F., Grotteria, M., and Wachter, J. (2019a). Foreseen risks. Working Paper 25277, National Bureau of Economic Research.

Gomes, J. F., Grotteria, M., and Wachter, J. A. (2019b). Cyclical dispersion in expected defaults. *Review of Financial Studies*, 32(4):1275–1308.

Gorton, G. and Metrick, A. (2012). Securitized banking and the run on repo. *Journal of Financial Economics*, 104(3):425–451.

Guiso, L., Sapienza, P., and Zingales, L. (2018). Time varying risk aversion. *Journal of Financial Economics*, 128(3):403–421.

Hansen, L. P. and Sargent, T. J. (2018). Structured uncertainty and model misspecification. Working paper, New York University and University of Chicago.

He, Z. and Krishnamurthy, A. (2013). Intermediary asset pricing. *American Economic Review*, 103(2):732–70.

Herbart, J. F. (1834). Lehrbuch zur Psychologie. In Kehrbach, K., editor, *Sämtliche Werke*, volume 4, pages 376–377. August Wilhelm Unzer, Königsberg, 2nd edition.

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

Howard, M. W. and Kahana, M. J. (2002). A distributed representation of temporal context. *Journal of Mathematical Psychology*, 46(3):269–299.

Howard, M. W., Viskontas, I. V., Shankar, K. H., and Fried, I. (2012a). Ensembles of human MTL neurons "jump back in time" in response to a repeated stimulus. *Hippocampus*, 22:1833–1847.

Howard, M. W., Viskontas, I. V., Shankar, K. H., and Fried, I. (2012b). Ensembles of human MTL neurons "jump back in time" in response to a repeated stimulus. *Hippocampus*, 22:1833–1847.

Hume, D. (1748). *An Enquiry Concerning Human Understanding.* Retrieved from Project Gutenberg: http://www.gutenberg.org/files/9662/9662-h/9662-h.htm.

Jegadeesh, N. and Titman, S. (1993). Returns to buying winners and selling losers: Implications for stock market efficiency. *Journal of Finance*, 48:65–91.

Jones, C. M. (2002). A century of stock market liquidity and trading costs. Working paper, Columbia University.

Jost, A. (1897). Die assoziationsfestigkeit in ihrer abhängigkeit von der verteilung der wiederholungen. *Zeitschrift für Psychologie und Physiologie der Sinnesorgane*, 16:436–472.

Kahana, M. J. (1996). Associative retrieval processes in free recall. *Memory & Cognition*, 24(1):103–109.

Kahana, M. J. (2012). *Foundations of Human Memory.* Oxford University Press, New York, NY.

Kahana, M. J. and Caplan, J. B. (2002). Associative asymmetry in probed recall of serial lists. *Memory & Cognition*, 30(6):841–849.

Kahle, K. M. and Stulz, R. (2013). Access to capital, investment, and the financial crisis. *Journal of Financial Economics*, 110(2):280–299.

Ladd, G. T. and Woodworth, R. S. (1911). *Elements of physiological psychology: A treatise of the activities and nature of the mind from the physical and experimental point of view*. Charles Scribner's Sons, New York, NY.

Loewenstein, G. (2000). Emotions in economic theory and economic behavior. *The American Economic Review*, 90(2):426–432.

Lohnas, L. J. and Kahana, M. J. (2014a). Compound cuing in free recall. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 40(1):12–24.

Lohnas, L. J. and Kahana, M. J. (2014b). A retrieved context account of spacing and repetition effects in free recall. *Journal of Experimental Psychology: Learning Memory and Cognition*, 40(3):755–764.

Lohnas, L. J., Polyn, S. M., and Kahana, M. (2015). Expanding the scope of memory search: Intralist and interlist effects in free recall. *Psychological Review*, 122(2):337–363.

Longstaff, F. A. and Piazzesi, M. (2004). Corporate earnings and the equity premium. *Journal of Financial Economics*, 74:401–421.

Lucas, R. E. (1978). Asset prices in an exchange economy. *Econometrica*, 46:1429–1445.

Luce, R. D. (1959). *Individual choice behavior*. Wiley, New York.

Malmendier, U. and Nagel, S. (2011). Depression babies: Do macroeconomic experiences affect risk taking? *The Quarterly Journal of Economics*, 126(1):373–416.

Malmendier, U. and Nagel, S. (2016). Learning from inflation experiences. *The Quarterly Journal of Economics*, 131(1):53–87.

Malmendier, U., Nagel, S., and Yan, Z. (2017). The making of hawks and doves: Inflation experiences on the FOMC. NBER Working Paper No. 23228.

Malmendier, U. and Shen, L. S. (2018). Scarred consumption. NBER Working Paper No. 24696.

Manning, J. R., Polyn, S. M., Baltuch, G., Litt, B., and Kahana, M. J. (2011a). Oscillatory patterns in temporal lobe reveal context reinstatement during memory search. *Proceedings of the National Academy of Sciences, USA*, 108(31):12893–12897.

Manning, J. R., Polyn, S. M., Baltuch, G., Litt, B., and Kahana, M. J. (2011b). Oscillatory patterns in temporal lobe reveal context reinstatement during memory search. *Proceedings of the National Academy of Sciences, USA*, 108(31):12893–12897.

McFadden, D. (2001). Economic choices. *American Economic Review*, 91(3):351–378.

McGeoch, J. A. (1932). Forgetting and the law of disuse. *Psychological Review*, 39:352–70.

Mian, A. and Sufi, A. (2009). The consequences of mortgage credit expansion: Evidence from the u.s. mortgage default crisis. *Quarterly Journal of Economics*, 124(4):1449–1496.

Moreton, B. J. and Ward, G. (2010). Time scale similarity and long-term memory for autobiographical events. *Psychonomic Bulletin & Review*, 17(4):510–515.

Mullainathan, S. (2002). A memory-based model of bounded rationality. *The Quarterly Journal of Economics*, 117(3):735–774.

Müller, G. E. and Pilzecker, A. (1900). Experimental contributions to memory theory. *Zeitschrift für Psychologie Eganzungsband*, 1:1–300.

Müller, G. E. and Schumann, F. (1894). Experimentelle beiträge zur untersuchung des gedächtnisses. *Zeitschrift für Psychologie und Physiologie der Sinnesorgane*, 6(80-191):257–339.

Murdock, B. B. (1974). *Human memory: Theory and data.* Lawrence Erlbaum and Associates, Potomac, MD.

Nagel, S. and Xu, Z. (2018). Asset pricing with fading memory. Working paper, University of Chicago and University of Michigan.

Polyn, S. M., Norman, K. A., and Kahana, M. J. (2009). A context maintenance and retrieval model of organizational processes in free recall. *Psychological Review*, 116(1):129–156.

Pratt, J. W. (1976). Risk aversion in the small and in the large. *Econometrica*, 44(2):420–420.

Raaijmakers, J. G. W. and Shiffrin, R. M. (1980). SAM: A theory of probabilistic search of associative memory. In Bower, G. H., editor, *The Psychology of Learning and Motivation: Advances in Research and Theory*, volume 14, pages 207–262. Academic Press, New York.

Savage, L. (1954). *The Foundations of Statistics.* Wiley, New York.

Schularick, M. and Taylor, A. M. (2012). Credit booms gone bust: Monetary policy, leverage cycles, and financial crises, 1870-2008. *American Economic Review*, 102(2):1029–1061.

Sederberg, P. B., Gershman, S. J., Polyn, S. M., and Norman, K. A. (2011). Human memory reconsolidation can be explained using the temporal context model. *Psychonomic Bulletin & Review*, 18(3):455–468.

Thorndike, E. L. (1932). *The fundamentals of learning*. Bureau of Publications, Teachers College, New York.

Uitvlugt, M. G. and Healey, M. K. (2019). Temporal proximity links unrelated news events in memory. *Psychological Science*, 30(1):92–104.

Underwood, B. J. (1948). 'Spontaneous recovery' of verbal associations. *Journal of Experimental Psychology*, 38:429–439.

Wachter, J. A. (2010). Asset allocation. *Annual Review of Financial Economics*, 2(1):175–206.

Yaffe, R. B., Kerr, M. S., Damera, S., Sarma, S. V., Inati, S. K., and Zaghloul, K. A. (2014). Reinstatement of distributed cortical oscillations occurs with precise spatiotemporal dynamics during successful memory retrieval. *Proceedings of the National Academy of Sciences*, 111(52):18727–8732.

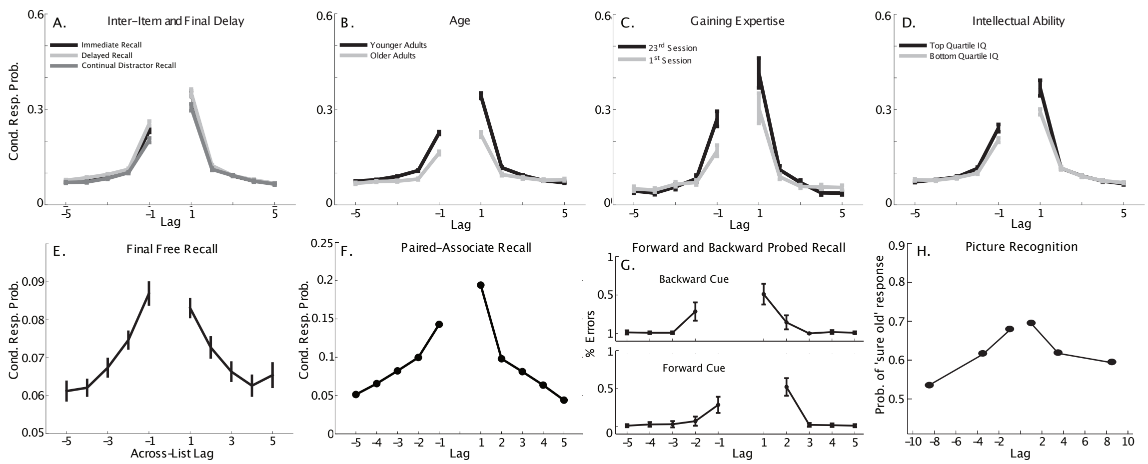

Figure 1: **Universality of Temporal Contiguity. A.** People tend to successively recall items studied in nearby list positions, as measured by the conditional-response probability as a function of the lag, or distance, between studied items (the lag-CRP). This function, which illustrates the temporal contiguity effect (TCE), remains invariant across immediate, delayed and continual distractor conditions of a free recall task. **B.** The TCE is reduced in older adults, indicating impaired contextual retrieval in older adults. **C.** Massive practice recalling items can increase the TCE, as evidenced by the comparison of first session and 23rd session of a multi-session recall experiment. **D.** Higher-IQ subjects exhibit a stronger TCE than individuals with average IQ. **E.** The TCE is not due to inter-item associations as it appears in transitions across different lists, separated by minutes, in a delayed final test given to subjects who studied and recalled many lists. **F.** The TCE appears in conditional error gradients in cued recall, where subjects tend to mistakenly recall items from pairs studied in nearby list positions. **G.** When probed to recall the item that either followed or preceded a cue item, subjects occasionally commit recall errors whose distribution exhibits a TCE both for forward and backward probes. **H.** The TCE also appears when subjects are asked to recognize previously seen travel photos. When successive test items come from nearby positions on the study list, subjects tendency to make high confidence “old” responses exhibits a TCE when the previously tested item was also judged old with high confidence. This effect is not observed for responses made with low confidence.

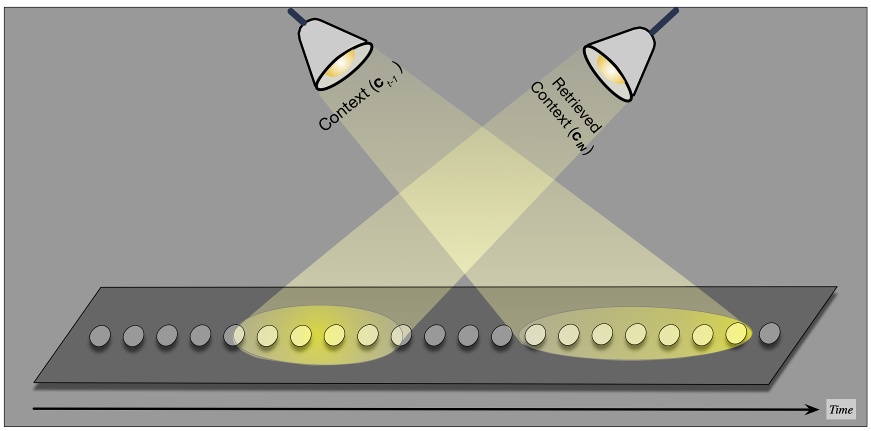

Figure 2: Retrieved Context and the spotlights of memory. In this illustration, memories appear as circles on the stage of life, with time going from left to right. All experiences that enter memory, as gated by perception and attention, take their place upon the stage. Context serves as a set of spotlights, each shining into memory and illuminating its associated features. The prior state of context c_{t-1} illuminates recent memories (rightgoing beam), whereas the context retrieved by the preceding experience, c^{IN}, illuminates temporally and semantically contiguous memories (leftgoing beam). Due to the recursive nature of context and the stochastic nature of retrieval, the lamps can swing over time and illuminate different sets of prior features.

Figure 3: Posterior probability and asset allocation as a function of sample length

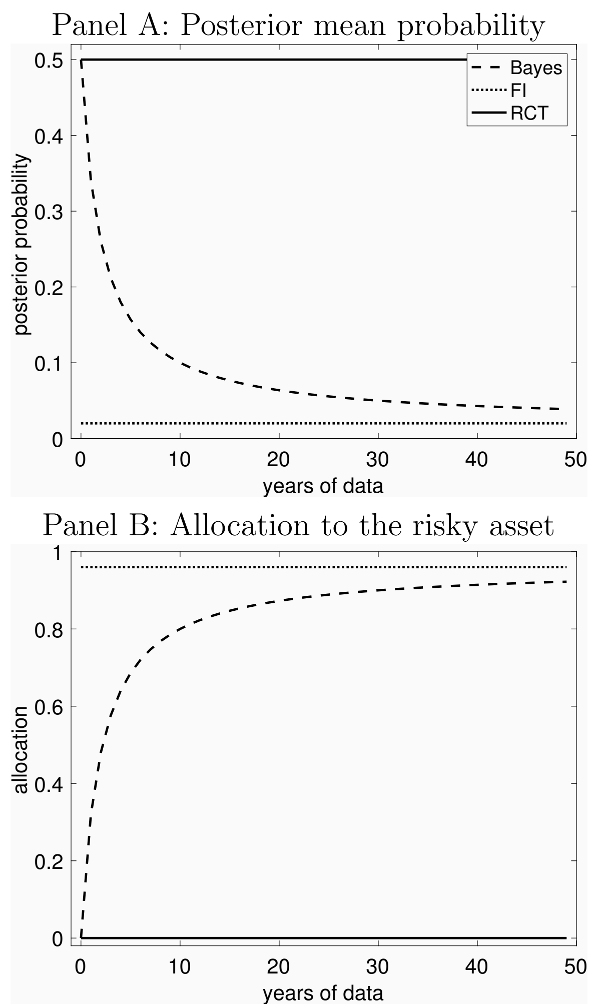

Notes: The figure shows posterior mean of the probability of a depression (Panel A) and the resulting asset allocation (Panel B) for the model presented in Section 3. FI refers to full-information. RCT refers to Retrieved context theory.

Figure 4: A jump back in time

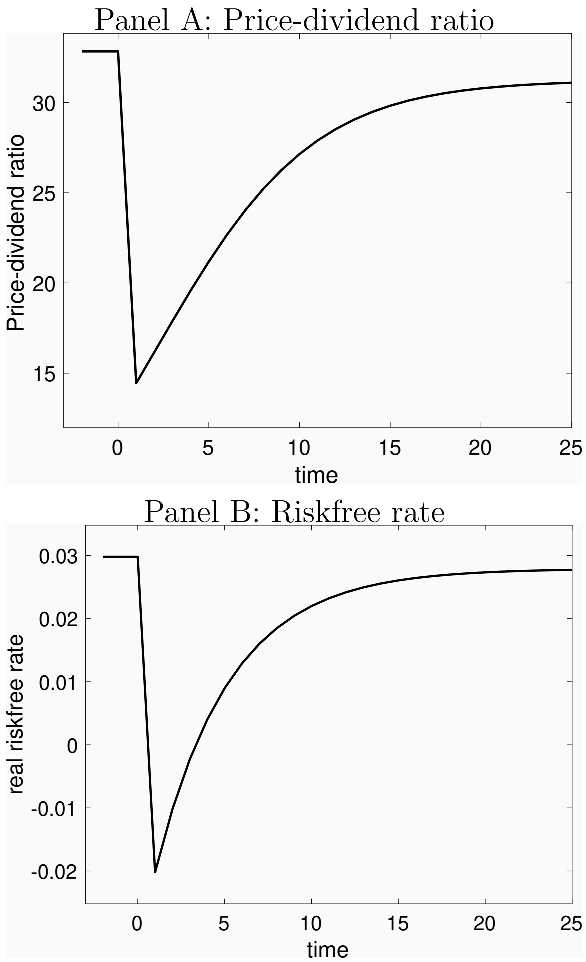

Notes: Panel A shows the aggregate price-dividend ratio pre-and post-banking crisis as a function of time since the crisis (the crisis is assumed to occur at time 1) for the model in Section 4. Panel B shows the real riskfree rate. We assume that neutral features follow the crisis. RCT stands for retrieved context theory. FI stands for full information. Exponential stands for exponential decay of context.

Figure 5: Prices and returns in the model for momentum

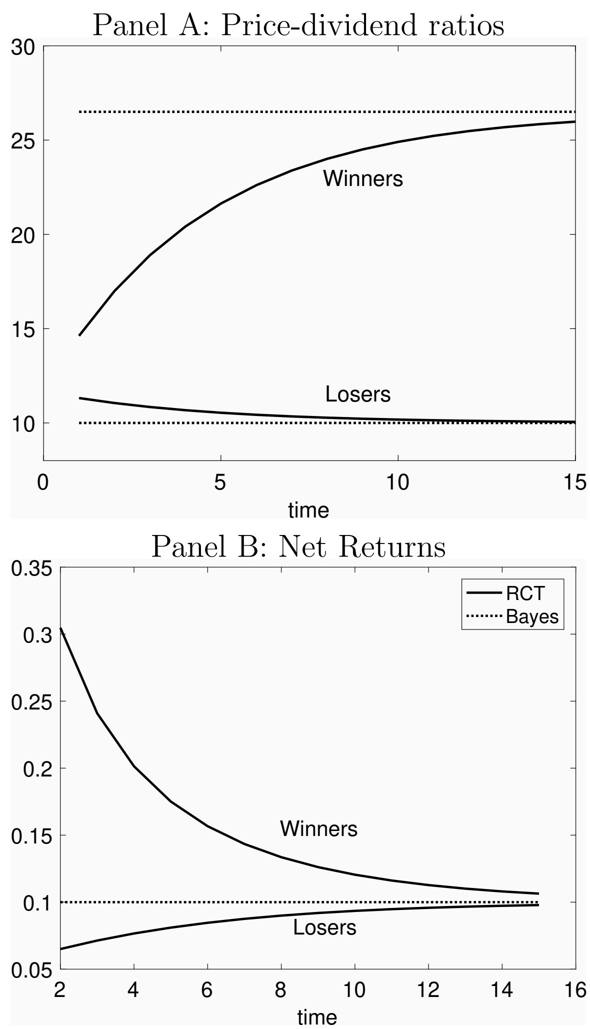

Notes: The figure shows price-dividend ratios (Panel A) and returns (Panel B) for the model presented in Section 5. RCT refers to prices and returns under retrieved context theory. Bayes refers to prices and returns under Bayesian updating (returns on winners and losers are identical).

Figure 6: Expected utility under context manipulation

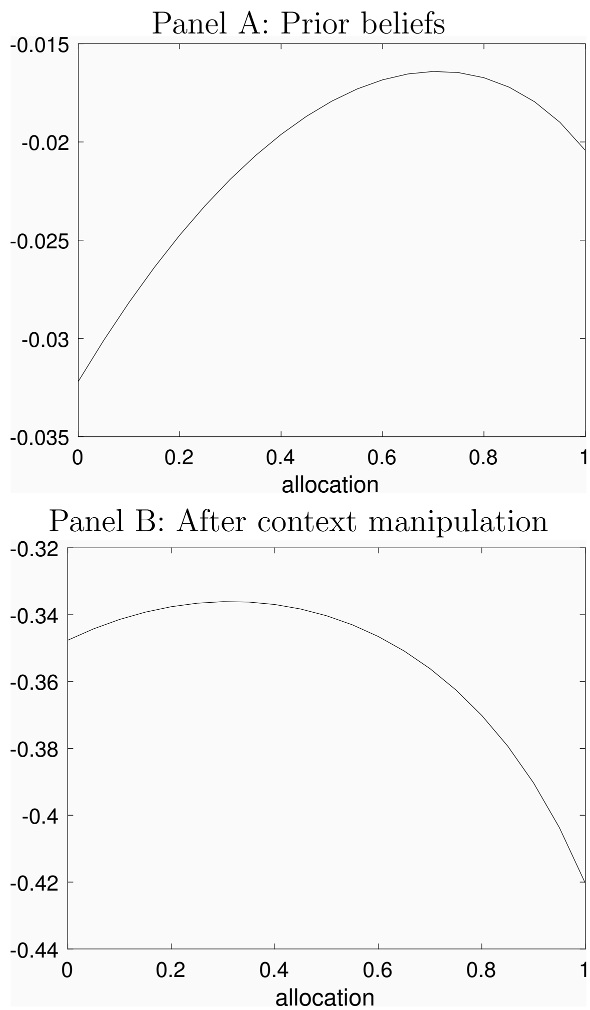

Notes:Expected utility as a function of allocation to the risky asset in the model of Section 6. Panel A shows utility prior to treatment by viewing a horror movie. Panel B shows utility after context has been manipulated by introducing a feature suggestive of danger (specifically, a horror movie).