Memory of
Past Experiences and Economic Decisions

Prepared for the *Oxford Handbook of Human Memory*

Ulrike Malmendier* Jessica A. Wachter†

January 20, 2022

Abstract

In traditional economic models, memories of past experiences affect choices only to the extent that they represent information. We review recent advances in economic research that have introduced a role for long-lasting effects of personal past experiences and the memory thereof into economics. We first document the empirical evidence on long-lasting *experience effects* in finance and economics. We then discuss the main approaches the literature has taken in incorporating psychological theories of long-lasting memories into economics. Our treatment suggests a role for models of memory in accounting not only for micro-level phenomena, but for anomalies within asset pricing and macroeconomics more broadly.

---

*Department of Economics, University of California, Berkeley and NBER; Email: ulrike@econ.berkeley.edu

†Department of Finance, The Wharton School, University of Pennsylvania and the Securities and Exchange Commission; Email: jwachter@wharton.upenn.edu. The Securities and Exchange Commission disclaims responsibility for any private publication or statement of any SEC employee or Commissioner. This chapter expresses the authors' views and does not necessarily reflect those of the Commission, the Commissioners, or other members of the staff. We thank James Paron for valuable resesarch assistance.

I Introduction

Over the last 70 years, at least until Behavioral Economics gained a foothold in economics, the field of economic research had converged toward a neoclassical consensus. This consensus provided economists with a common language, a core set of modeling principles, and some agreement over what questions are important. At the core of the neoclassical consensus is the notion that economic choices result from agents maximizing their individual objective functions subject to a set of constraints.

Consider an individual who aims to make a decision by using the past to infer something about the future. To fix ideas, suppose the agent is deciding how much wealth to invest in the stock market and how much in a riskfree bond. The agent has a set of beliefs about what the future might look like, for example, what future stock returns will be. In economic theory, this optimization problem is commonly expressed as

$$\max_{\pi} \mathbb{E} u(W(1 + \pi \tilde{r})), \tag{1}$$

where $\tilde{r}$ denotes the (random) net stock return, and $\mathbb{E}$ denotes the agent's subjective expectations. The function u measures the agent's happiness or "utility" from wealth. We can think of utility in literal terms, as "usefulness." Because future wealth is unknown, the agent computes the *expectation*, or *expected utility*, namely the probability-weighted sum of possible future utilities. The general approach of maximizing the probability-weighted sum of utility from possible future payoffs (as represented by the expectations operator) is a foundation of modern economics.[1]

One might ask, though: What are the sources of the beliefs? Do they come from data, or do they come from "priors"? If data, then which data, and if priors, what is the source of the priors? Given that memory intuitively seems central to decision-

[1] Savage (1954) describes the foundations of this general approach, while Arrow (1970) famously analyzes the portfolio choice problem.

making problems such as (1), why do so few economic models make reference to the memory literature?

Implicitly, there is always a model of memory. In the neoclassical economic model, agents can perfectly retrieve data, and not only some of the data, but *all* of the data relevant to the problem at hand. One variant is that agents possess full information, which is to say that they have access to a sufficiently large amount of data so that there is minimal (if any) lack of knowledge about probabilities. A second variant is that their database comes equipped with a prior and a likelihood that permits them to learn the correct population probabilities in a way that is perfectly consistent with Bayesian updating. Both are versions of an influential framework known as "rational expectations." We can think of rational expectations as taking the original Savage (1954) approach and combining it with Bayesian updating, to turn it into a model that allows for dynamic learning. (In what follows, we will use the terminology "Savage paradigm," "rational expectations," and "Bayesian updating" interchangeably.)

The rational expectations framework is disciplined, and it takes seriously the idea that agents might have different information. Some might have better information than others, and they can learn from observing the outcome of others' actions.[2] It is also self-contained. It apparently has no need of a more explicit model for memory.

But does it? Challenges to the Savage (1954) paradigm, not to mention the later paradigm of rational expectations are nearly as old as the paradigm itself. Luce (1959) provided an early alternative to Savage (1954), in which an individual probabilistically makes choices within a narrowly defined set. This approach is carried forward to the present in economics by McFadden (2001) and is highly influential in empirical research. Modigliani (1977) was an early critic of rational expectations. Savage himself, as well as Gilboa and Schmeidler (1995) illustrate the weakness of the approach using examples

[2]It also leaves room for investors facing costs in acquiring information (Sims et al., 2010). Such a model is still within the utility maximization framework above, but with a broader view of the utility function and the constraints.

in which it is impractical for agents to formulate the various alternative states of the world that might arise, never mind the probabilities.

Interestingly, these early critiques, however powerful, did not do much to dent the status of utility theory, and rational expectations, as a dominant paradigm. Rather, the newer, "behavioral" models of belief formation have gained tractions because of empirical evidence against Bayesian updating and rational expectations. In applied microeconomics, researchers with access to improved data on individual investment, savings, and consumption behavior became increasingly concerned about the rational-expectations hypothesis, as they were encountering evidence on, for example, self-control problems and loss aversion.[3] Yet it would take two economic crises: the stagflation of the 1970s, and the financial crisis and Great Recession 2008, to bring behavioral approaches into the mainstream; the reason being that there simply was no alternative.[4]

In the remainder of this chapter, we first synthesize some of the empirical literature that has analyzed these crises and other macroeconomic realizations. Based on the evidence on the lasting impact of these events on those who personally experienced them and in particular the "four stylized features" of long-term scarring that emerge, we turn to new models of experience-based learning, contextual retrieval, and memory, which this evidence necessitates.

II Empirical Motivation: The Role of Past Experiences and Memory in Financial Decisions

When testing for the determinants of human decision-making and beliefs, financial markets are a fruitful laboratory. Few other markets offer as rich a data set on prices,

[3]For reviews of the literatures see Frederick et al. (2002); Ericson and Laibson (2019) on limited self-control and O'Donoghue and Sprenger (2018); Barberis (2018) on loss aversion.

[4]The clash of model and reality within the field of macroeconomics is a fascinating topic in and of itself and is partially recounted in Mankiw and Romer (1991), Woodford (2013), Cochrane (2018), and Akerlof (2019).

buy and sell decisions, and the timing of information releases. Compared to the typical laboratory study on the role of memory, field data from financial markets offer less control and randomization; but a key advantage is that they provide insight into the role of past experiences and memory thereof over significantly longer horizons and for economically more significant decisions. They allow to establish a link between memory and persistent, even generational differences in economic decision-making, and to link those differences to the shared experiences of members of those generations.

Our overview of the motivating empirical facts focuses on distilling four key features that have emerged from studying how past experiences affect financial decisions:

1. **Long-lasting effects.** Personal experiences of economic and other shocks shape individual beliefs and choices for years and decades to come.

2. **Recency bias.** More recent experiences have a stronger impact on individual expectations and risk-taking than experiences made earlier in life, though big enough shocks have a detectable impact on individuals decades later.

3. **Context dependence.** Experience effects are specific to the domain, or context. They do not alter risk attitudes in general, nor do they alter beliefs in a way that would seem consistent with physical correlational structures. This is contrast to neo-classical economic models implying a transfer of experience and learning across domains. As a result, the agent may appear to have a set of beliefs that are inconsistent with traditional theories of rationality.

4. **Robustness to expert knowledge.** Scarring effects of personal experiences are observed not only among untrained individual investors or consumers, but also among highly educated and specialized individuals, even in their area of expertise.

We provide a selective overview of the empirical evidence that has generated these four key features.

The starting point of the research agenda on experience effects was a study that analyzed the generation of "Depression Babies," i. e., those who lived through the Great Depression at young age. The Depression experience famously turned an entire generation away from the stock market and other risky investment for decades to come. Malmendier and Nagel (2011) provide evidence of precisely this relation between stock-market experiences and stock-market investment, not only for the generation coming of age during the Great Depression, but also for later generations who experienced bad stock-market realizations. Their analysis starts from one of the main puzzles economists wrestle with, the so-called equity premium puzzle. Given the high returns to stock-market investment, why is stock-market participation so low? Why doesn't the entire population invest at least some fraction of their savings in the stock market? While the increased availability of broadly diversified, low-fee funds has helped increase participation from about 20% post-World War II to 50% today (Vissing-Jorgensen, 1998; Li, 2014), economists have struggled to explain the absence of the remaining population.

To test for the influence of past experiences, Malmendier and Nagel (2011) collected about five decades of stock-market participation data from the *Survey of Consumer Finances* and its precursor, obtained from the Inter-university Consortium for Political and Social Research at the University of Michigan. The data include individuals born as early as the beginning of the 20th century all the way up to recent times. The data reveal that individuals' past exposure to stock-market realizations has significant explanatory power for the decision to invest in the stock market, both on the extensive margin—whether individuals participate at all in the stock market—and on the intensive margin—how much they invest conditional on participating. Consider a person over whose lifetime so far the stock market performed quite well, say at the 90th percentile of the sample, which amounts to a weighted average of 11.9% (—see Figure 1 and more explanations below for how individuals weight their lifetime experiences—), and contrast that person with another investor over whose lifetime so far the stock market performed only at the 10th percentile of the sample (6.2%). The authors estimate that, controlling for age effects, year effects, wealth, and income, the inter-decile difference between these two experiences predicts a difference of 10.2 percentage points (pp) in stock-market participation. This effect is very large relative to the average in-sample participation rate of 34.2%. It is orthogonal to any time trends or aggregate effects (such as time-varying aggregate risk aversion or any mechanical positive relation between recent stock returns and households' stock allocation); all of those effects are already removed by the inclusion of year fixed effects. Hence, we are seeing evidence of lingering effects of past experiences with positive or negative stock-market events on individuals' willingness to participate in the stock market years and decades later.

The influence of past experiences is similarly large on the intensive margin. Conditional on stock-market participation, past life-time experiences help to explain what *amount* of their liquid assets investors are willing to invest in publicly traded stock, either directly or indirectly (through funds). Here, a comparison of investors at the 10th versus the 90th percentile of experienced stock returns predicts a difference in the allocation of liquid assets to stocks of about 7.9 pp. This finding is particularly remarkable since prior finance research on household portfolio choice had failed to identify household characteristics that predict the risky-asset share within the sample to stock-market participants (Curcuru et al., 2009). Hence, experienced stock returns emerge as a major determinant of households' willingness to bear stock-market risk in the long-run, and are a first example of key feature 1 (Long-lasting effects).

The authors find a very similar relationship between investment in bonds and bond-related past experiences, but—and this is particularly relevant in the context of memory—no cross-fertilization. That is, while stock-related experiences predict stock-market investment and bond-related experiences predict bond-market investment, there are no (positive or negative) cross-effects between these two asset markets. The latter result is hard to grapple with if we do not allow for experiences and memories to play a role. It indicates a "domain specificity" (cf. key feature 3) that is well known in the psychology literature (Slovic et al., 1986; MacCrimmon and Wehrung, 1990; Bromiley and Curley, 1992; Goldstein and Weber, 1995; Weber et al., 2002), but hard to reconcile with standard economic modeling of risk preferences.

Moreover, the results of Malmendier and Nagel (2011) do point to a beliefs-based channel through which experiences affect risk taking: As additional UBS/Gallup data reveal, a 1 pp increase in experienced stock returns is associated with a 0.5-0.6 pp increase in the stock-market return expected for next year. That is, these experiences do not necessarily change the utility u, but the probabilities underlying the expectation that we referred to in (1). Needless to say the distinction between the beliefs channel ($\mathbb{E}$) and the preference channel (u) is not necessarily as clear-cut as model (1) makes it appear. If our stock-market experiences alter our "taste" for stock-market investment and make us think more negatively about it, this negative thinking could tie both aspects together – a pessimistic view of future stock-market returns might make an investor dislike the stock market more, above and beyond the shift in probabilities. As we will explore in our review of the theoretical work on memory and retrieval, it is indeed the case that the traditional distinction in economics between belief- and preference-based channels is too simplistic. Regardless of the categorization, though, the important insight is that memories of past experiences are a strong determinant of behavior in the long-run, and also affect how individuals assess the distribution of future realizations in similar markets or domains.

So far, we have seen that the longer-horizon analysis of stock-market data helps to reveal what we dubbed key feature 1 (Long-lasting effects) and key feature 3 (Domain specificity). Another advantage of the longer-horizon studies in economics, relative to the laboratory studies on memory, is that they identify which past experiences exert

Figure 1: Experience Weighting

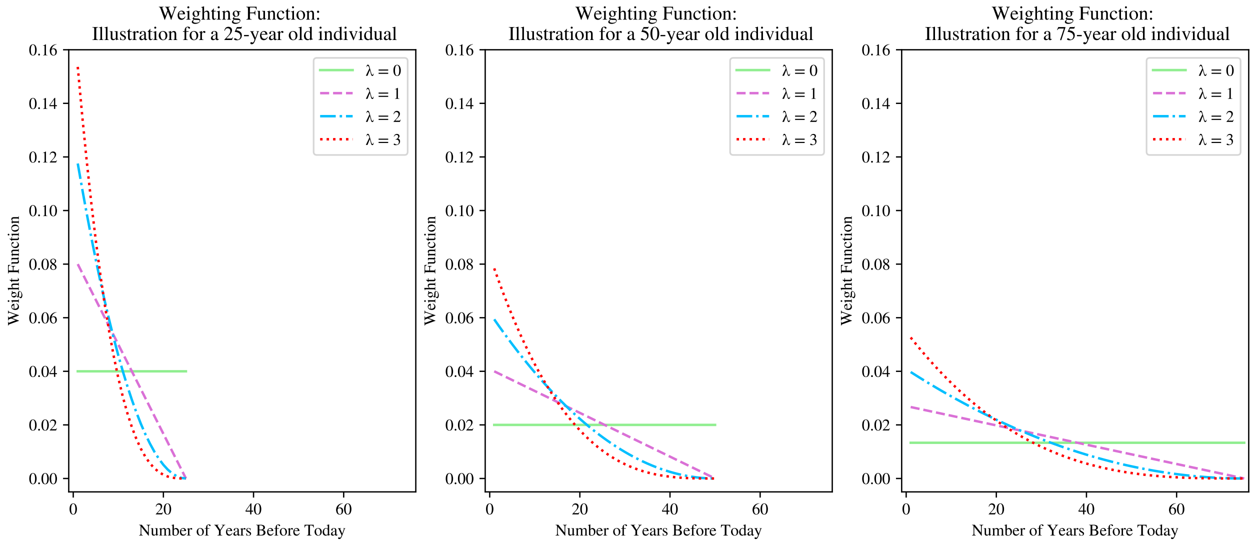

Figure 1 illustrates the shapes of the weighting function for positive λ values 1, 2, and 3 (in line with empirical estimates) and the case of $\lambda = 0$ (equal-weighting, for comparison). The weights on the left are for a 25-year old investor, those in the middle for a 50-year old investor, and those on the right for a 75-year old investor.

stronger or weaker influences on which generations (cohorts) at what point in time. That is, the field data allows to estimate a weighting function that determines how important a realization from k years ago is for an individual of age a at time t.

The study above as well as later studies (discussed below) have shown that a weighting function of the type displayed in Figure 1 for $\lambda > 0$ (and typically $\lambda \approx 1$ or slightly higher) appears to perform best in fitting the data. As the figure shows, individuals put extra weight on all experiences that have occurred over their lifetimes so far. Hence, positive weights are assigned to realizations from the past 25 years for a 25-year old (left graph), to realizations from the past 50 years for a 50-year old (middle graph), and to realizations from the past 75 years for a 75-year old (right graph).

For the range of empirical estimates in prior work around $\lambda = 1$, these weights are decaying in the time lag. That is, for positive values of λ as shown by the dashed and dotted lines, most weight is assigned to the most recent experiences and least weight to the earliest experiences in one's life. Hence, individuals exhibit recency bias (key feature 2) in the type of weighting they apply to their personal past experiences. Note that recency-biased weighting is not necessarily implied by experience-based learning. For example, the data could have identified negative values of λ as generating the best fit, in which case experiences from childhood would have left the strongest mark on beliefs for the rest of the individual's life ("first impressions matter"). Instead, the field data indicates that individuals continue to update throughout life and continue to assign extra weight to their most recent experiences.

The juxtaposition of the weighting functions for three differently-aged investors in Figure 1 also highlights, though, that younger generations generally put *more* weight on recent experiences than older generations: The weights towards the left of the plots (recent past) are shifted up more for the 50-year old than the 75-year old investor, and even further up for the 25-year old investor. In other words, individuals with a shorter lifespan so far are more heavily influenced by recent macro-finance realizations.

The first-mentioned phenomenon of recency bias is consistent with the literature on recency bias in (social and cognitive) psychology (Luchins, 1958; Tversky and Kahneman, 1973; Greene, 1986; Hogarth and Einhorn, 1992), and had been picked up by several economic theory approaches as well (Fuster et al., 2010, 2012; Glaeser and Nathanson, 2017; Bordalo et al., 2018). The latter phenomenon, however, is unique to experience-based learning, and crucial for the empirical identification: older generations react differently from younger generations with shorter life histories, and as a result can exert more risk aversion and at other times less risk aversion than the other cohort, depending on the difference between recent events and those further in the past.

An intriguing finding emerging from the broader experience-effect literature is that the above patterns are robust to different market settings. That is, experience-based belief formation and risk taking has been found not only in the stock market, but also in several other market settings in the macro-finance realm, such as the context of inflation experiences, and an almost identical weighting function performs well in explaining expectation formation in these other markets.

Let's consider the inflation example in a bit more details – also since it will provide us with evidence regarding key feature 4 (Robustness to expert knowledge). The wide cross-sectional differences in people's beliefs about future price increases have long been a puzzle to economists. Typical macro models aim to fit the aggregate inflation forecast (e. g., the median; Orphanides and Williams (2005); Milani (2007)), and several allow for some cross-section stochasticity (Mankiw and Reis, 2002; Carroll, 2003); but the mainstream models largely fail to explain which individuals would be particularly optimistic and which ones more pessimistic regarding future price increases.

Past inflation experiences are able to do exactly that. Malmendier and Nagel (2016) show that experience-based learning predicts when, say, older generations will expect particularly high inflation, when the reverse will be the case and younger generations are more pessimistic, and when the relative beliefs of different generations will cross. This is exemplified in Figure 2. The figure displays both the raw data of stated inflation expectations by age groups (below 40, 40-60, above 60) relative to the population mean, and the experience-based model fit (for the same age groups). The raw data show how far apart different parts of the population can be in their economic beliefs (up to 3 pp in 1980) and how their relative position can change over time. This heterogeneity in inflation expectations by cohort is predicted by lifetime inflation experiences, weighted in a similar fashion.

Importantly, experience-based beliefs about inflation (and, similarly, about future interest rates, cf. Botsch and Malmendier (2019)) translate into major economic decisions, specifically about housing. The most important financial decisions many households make over their entire lifetime is whether to buy a house and whether (and

Figure 2: Disagreement about future inflation 1960 - 2020

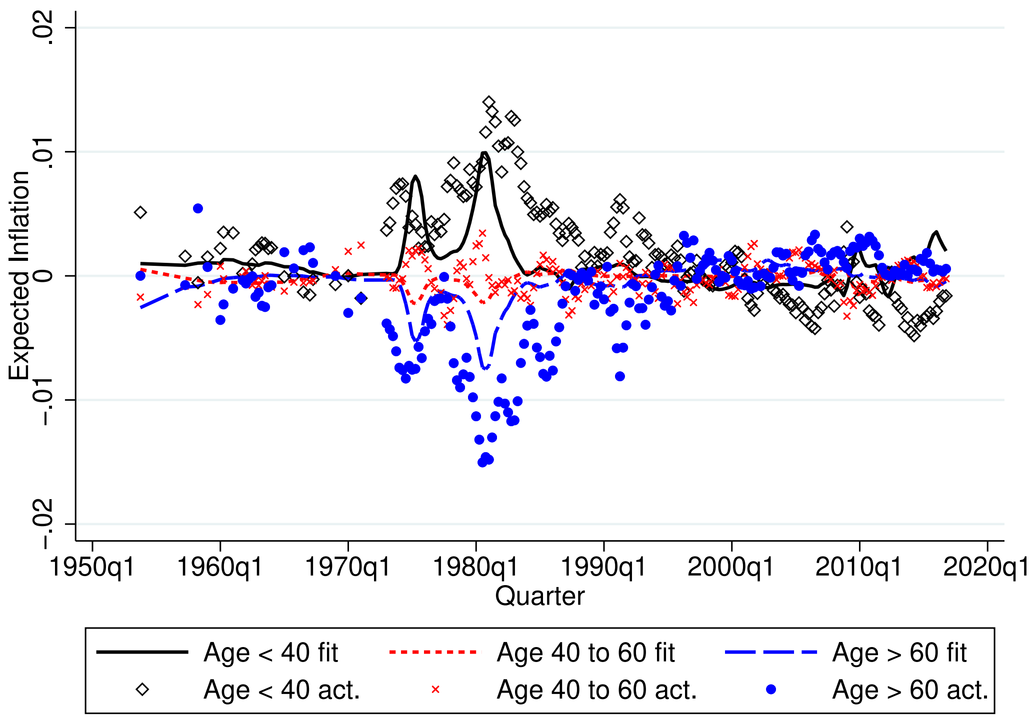

Figure 2 is updated from Malmendier and Nagel (2016) and based on Michigan Survey of Consumers data. The figure demonstrates the discrepancy between inflation expectation of different age groups. We first compute the four-quarter moving averages of mean one-year inflation expectations of young individuals (below 40), mid-aged individuals (between 40 and 60), and old individuals (above 60). We then compute deviations of each age from the cross-sectional mean expectation of all age groups. In most periods percentage forecasts are available from the survey. In some periods they are imputed from categorical responses or not available. The graph provides a comparison of four-quarter moving averages of actual and fitted one-year inflation expectations for individuals of the three age groups, shown as deviations from the cross-sectional mean expectation.

which) mortgage to take out to finance the home purchase. Malmendier and Nagel (2016) show that past exposure to high inflation predicts fixed-rate mortgage (FRM) balances. That is, consumers who have lived through high inflation periods seek to invest their money in real estate and consider fixed-rate mortgages a relatively cheap form of financing. They also overestimate future (nominal) interest rates and choose fixed rates over variable rates, even at times when adjustable-rate mortgages (ARMs) are advantageous (Botsch and Malmendier, 2019). These influences were particularly strong in the wake of the Great Inflation: Given the relative costs of fixed- versus variable-rate mortgages, the generation of Baby Boomers should have taken out one million fewer FRMs in the late 1980s, and still half a million fewer in the late 1990s. The costs of these deviations are large—about \$14bn in excess payments in the late 1980s, and still almost \$9bn in the late 1990s.

We can also compare inflation realizations across countries and show that the national memory of high versus low inflation shapes housing markets. As Malmendier and Steiny (2017) point out, there are vast discrepancies in homeownership rates within Europe. For example, less than half of all households in Germany and Austria own their home, compared to over 80% in Slovakia, Hungary, and Spain. Only 57% of households own their home in France, but 83% do in neighboring Spain. There is also variation in the composition within countries. In Italy, 49% of 30-year-olds own their home compared to 80% of 60-year-olds. In the Netherlands, we see the opposite pattern with 63% homeownership among 30-year-olds and 59% among 60-year-olds. Differences in individuals' lifetime exposure to inflation help explain these differences in homeownership.

Malmendier and Steiny (2017) show that past exposure to inflation helps explain these striking cross-country differences in homeownership rates. They use the European Central Bank's Household Finance and Consumption Survey (HFCS) to estimate the relation between homeownership rates and historical inflation. Figure 3 illustrates the estimated relationship. The left graph shows the annual inflation for countries in the top quartile of homeownership in the HFCS (where homeownership rates are around 80%), and the right graph shows inflation in the bottom homeownership quartile (where homeownership rates are around 50%). Contrasting the high inflation rates over the last 40 years in the left panel with the much lower rates in the right panel illuminates the basic relationship: Countries that experienced higher historical inflation now feature significantly higher homeownership, consistent with high inflation often being quotes

Figure 3: Inflation history in the top and bottom quartiles of homeownership rate

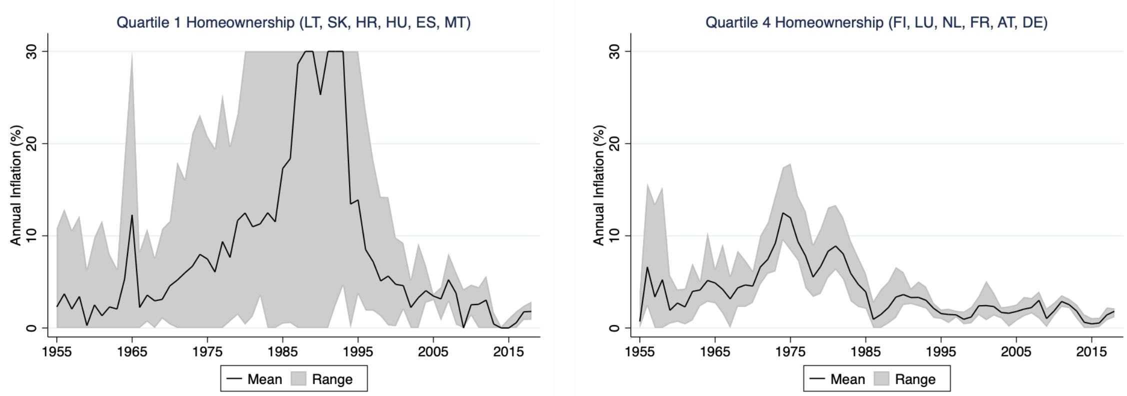

Figure 3 plots the mean and range of inflation across countries in quartile 1 and quartile 4 of homeownership rates from a list of twenty countries. Quartile 1 includes countries with the highest homeownership rates across all available survey waves of the European Central Bank's Household Finance and Consumption Survey (HFCS), and quartile 4 the lowest. The figure updates from inflation data sources described in Malmendier and Steiny (2017). In addition, inflation for chart is capped above at 30% and below at 0%.

as an important factor contributing to the motivation to purchase a home as it may serve as an inflation hedge. The individual-level estimations in Malmendier and Steiny (2017) also confirms that the "national memory" of past inflationary periods affects different cohorts differently, as implied by the weighting functions shown in Figure 1, and help to explain why homeownership is more predominant among older generations in some countries and among younger generations in other countries.

As already alluded to above, the example of inflation experiences and their lasting influence on beliefs and decision-making is of particular importance within the literature on experience effects for yet another reason: it provides a setting to show key feature 4, namely, the robustness of experience-based learning to expert knowledge. The members of the Federal Reserve Bank's Federal Open Market Committee (FOMC) certainly count among such experts on the topic of inflation. The FOMC consists of the Federal Reserve governors and regional presidents, and meets eight times per year to determine monetary policy. Needless to say that FOMC members have all available inflation data and models at their finger tips, and are experts of monetary policy. And yet, it turns out that their own views of future inflation are heavily colored by their personal past experiences.

Malmendier et al. (2021) analyze the semi-annual inflation forecasts that the FOMC member submit to Congress in their Monetary Policy Reports. They document a strong positive relation ship between personal experiences and forecasts, as shown in Figure 4 (replicated based on Malmendier et al. (2021), normalizing by staff forecasts to account for time effects). The past experiences of central bankers translates into their actual decision-making. Prior lifetime exposure to relatively high inflation predicts a significantly more hawkish policy stance, which can be detected in these members' dissenting votes (at the FOMC meetings) towards higher interest rates. Vice versa, experiences of lower inflation predict a more dovish stance, which again can be detected in their dissenting votes, in that case in favor of lower rates. In fact, Malmendier et al. (2021) show that, accounting for for FOMC experiences yields better predictions of the federal funds rate, i. e., the rate rate at which depository institutions (commercial banks) lend reserve balances to each other overnight. Setting the target federal funds is a key policy tool of the FOMC, and the empirical evidence suggests that the personal histories of the central bankers in charge have colored their decision-making.

Both settings we have discussed so far, stock returns and inflation, have in common that they allow researchers to identify experience effects out birth years. At any point in time t, cohorts differ in their stock-market and inflation experiences depending on the realizations over their individual lifetimes so far, and these cohort differences change over time as new realizations accumulate. In other words, individuals of different ages behave in systematically different ways, but not as a result of life-cycle determinants and aging, but as a result of their past experiences. At another point in time, individuals of those age groups will behave differently, as a result of the experiences those cohorts have made by that point in time.

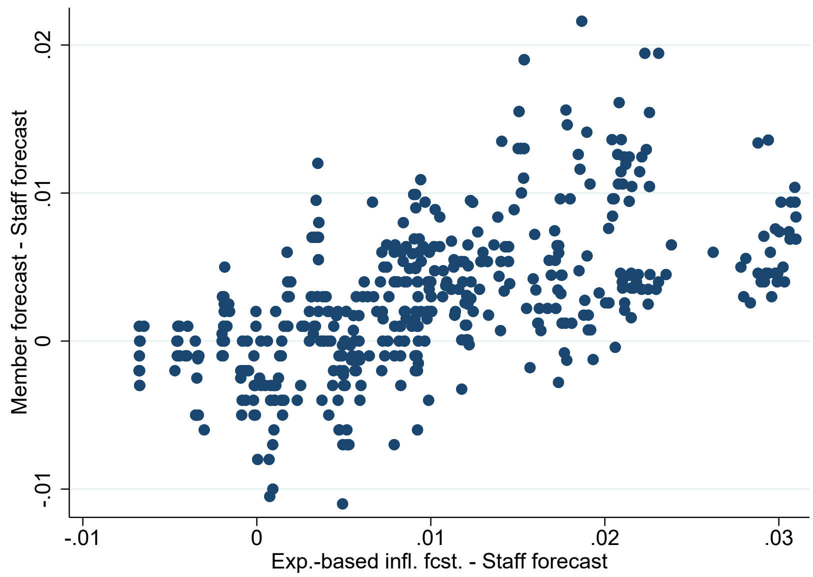

Figure 4: **FOMC Members' Inflation Forecasts and Experience-based Forecast** Member forecast from semi-annual Monetary Policy Report to Congress, 1992 - 2007. Staff forecast: Greenbook forecast. Experience-based forecast: AR(1) model forecast estimated based on weighted life-time inflation data for each FOMC member.

Birth year is not the only determinant of experiences, though. When we step away from aggregate, macro-level experiences, and turn to more micro-level variables, two individuals born in the same year likely have different exposures and make different experiences throughout their lifes. One example is lifetime unemployment experiences. Past unemployment experiences "scar" consumers' beliefs about their employment and earnings prospects. Using data from *Panel Study of Income Dynamics* (PSID), Malmendier and Shen (2018) find consumers who have lived through times of personal unemployment (but also, going back to cohort effects, times of high unemployment rates) exhibit persistent pessimism about their future financial situation and spend significantly less, controlling for the standard life-cycle consumption factors. This is in spite of the finding that their actual future income is uncorrelated with those past experiences (after including the usual controls such as income). Also, due to their experience-induced frugality, scarred consumers build up more wealth.

Figure 5: Average Life-Cycle Consumption

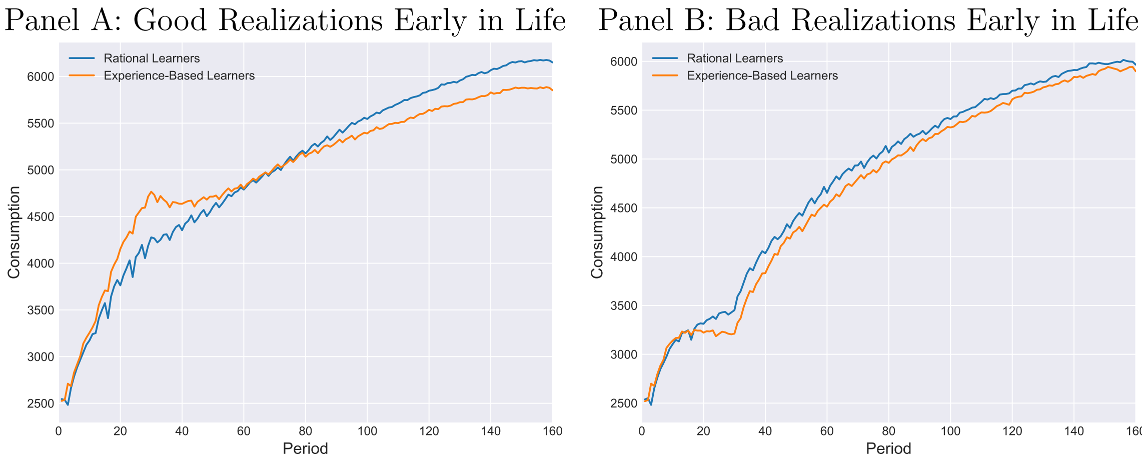

Based on simulated data from the model of Malmendier and Shen (2018). Average consumption for rational learners and experience-based learners (with $\lambda = 1$) in the low-education group, based on 10,000 lifetime simulations for each type and then restricted to those simulations where agents have, or in the rational case would have, a believed delta of 0.025 or less at period 30 in Panel A and a believed delta of 0.1 or greater at period 3 in Panel B.

Panel A of Figure 5 illustrates the long-lasting effects of how past experiences of unemployment in simulations of a lifecycle consumption model adapting the model from Low et al. (2010). First, consider consumers who where "lucky" early in life (no income/job loss). In that scenario, experience-based learners (in orange) believe their risk of job loss to be small (0.025 or less) in period 30 of their lives, even though the true probability of them losing their job is 0.049. As a result, they engage in pronounced overconsumption in the early periods, relative to the rational benchmark (in blue), for which they have to pay later in life. In contrast, Panel B of Figure 5 illustrates the opposite scenario. Here, the simulation considers only consumers who had bad luck early in life and, at period 30, believe their risk of job loss to be 0.1 or greater, instead of the true probability of 0.049. In this "unlucky" group, experience-based learners consistently consume less than their rational counterparts for almost their entire lives and build up significant wealth reserves due to excess frugality.

All of the above studies have helped to establish that personal lifetime experiences have a longlasting influence on individuals' beliefs and decision-making. We also saw some of the characteristics of those longlasting influences—the embedded recency bias, their domain specificity, and their pervasiveness even among experts.

The next question is how to test whether experience-induced choices are linked to memories of those experiences. At this point, there is little direct evidence on that link. It would be interesting to apply some of the techniques eliciting "retrieval" from the laboratory studies on memory to individuals exposed to measurable experiences from years and decades ago as explored in the field studies.

A first step in this direction, albeit still indirect, comes from some recent studies that suggest "synaptic tagging" and, in particular, "emotional tagging" of past experiences as the underlying neurobiological mechanism of experience effects (cf. Laudenbach et al. (2019a) for a first high-level introduction). These papers ask whether the emotional context of an experience determine how strongly and how positively or negatively an experience is engrained (anchored) in someone's memory and in what direction it influences their behavior.

Specifically, Laudenbach et al. (2019b) study a sample of East Germans in the area of the formerly communist GDR and analyze the long-term effects of living under communism and its anticapitalist doctrine on households' attitudes towards capital markets years and decades after the 1990 reunification. By exploiting heterogeneity in how negatively (or positively) different individuals experienced the GDR and its oppression, they are able to tease out the effects of negative and positive encoding of past experiences, pointing to the emotional component of anchoring in memory.

The study starts from the baseline finding that, decades after the German Reunification, East Germans are still significantly less likely to participate in the stock market than West Germans, with a 10 pp gap between East and West Germany. In contrast to the Depression Babies research discussed above, though, this difference cannot be linked to personal negative stock-market experiences since there was no stock market in the GDR. East Germans were exposed only to the communist doctrine and its strong negative views of stock markets as "the root of all evil." Yet, this "experience" sufficed to generate a lasting aversion to stock-market investment. It is also interesting to see that there are some stocks East Germans are *more* likely to invest in, namely, consistent with the communist friends-and-foes propaganda, shares of companies from communist countries (China, Russia, Vietnam) and also of state-owned companies. Vice versa, East Germans are particularly unlikely to invest in American companies and the financial industry.

The core of the analysis in Laudenbach et al. (2019b), though, is not about the exposure to communism, but on *how* an individual has experienced the communist system. According to the emotional-tagging concept, emotionally arousing events are not only remembered better (since emotionally dependent information is associated with enhanced memory (Richter-Levin and Akirav, 2003)), but it also matters whether an experience is tagged with positive or negative emotions, as the affective system determines which components from the collection of processed information are preserved in memory (Bergado et al., 2011). Correspondingly, Laudenbach et al. (2019b) test whether citizens with plausibly more negative experiences under the communist system exhibit weaker exposure effects (and embrace capitalism and the stock market), while those with plausibly more positive experiences exhibit stronger exposure effects.

One proxy for a more negative experience of living under communism is environmental pollution. Some areas in the GDR had the highest levels of dust and sulfur dioxide emissions across all European countries (Petschow et al., 1990), directly contradicting the claim of the communist regime to protect the environment in the interest of peoples' wellbeing. Another proxy is religious oppression. As common in communist systems, religious life was suppressed in the GDR (Müller et al., 2013), and East Germans living in particularly religious areas might hold more positive views about Western countries, which honor the freedom of religion. A third proxy is variation in the access to West German TV, which was a major source of entertainment for East

Germans. As has been documented, the lack of access to it resulted in lower satisfaction with the GDR and hence a higher resistance to the political system (Kern and Hainmueller, 2009). Access to West TV depended on the area of living; for example, TV signals from the West could not be received in some low valleys.

All three proxies of negative emotional tagging work as predicted. Living in areas with high pollution, higher religiosity, and no access to Western television predicts a significantly lower stock-market participation gap. For example, living in an area which had high pollution or no access to Western television closes the gap by 6.8 pp and 9.4,pp respectively relative to other East Germans.

Vice versa, the participation gap is higher among individuals whose experience with the communist system was likely tagged with positive emotions. Examples include East Germans who lived in one of the GDR's celebrated communist "showcase" cities, and East Germans who lived in the municipality of Olympic champions, whom the GDR celebrated as national heroes and symbols of the superiority of socialism over the capitalist system. For these and other proxies, the data reveal an even stronger aversion to capital markets and even lower stock-market participation. For example, East Germans living in a showcase city are 13 pp less likely to participate in the stock-market than other East Germans. Note that this is similar in magnitude to the baseline effect of being an East German.

These results highlight the importance of prior lifetime experiences and their longlasting influence on attitudes, beliefs, and choices. Experience-based learning can strongly influence beliefs about key outcomes, such as inflation or unemployment. The intensity and direction of such influences appears to be predicted by their emotional coloring, which points to a role for memory since we know that emotions are essential in anchoring past experiences in memory. This, in turn, implies that memory might play an essential role for economic choices and behavior in the long-run, and thus beyond the relatively short horizon that characterize much of the laboratory evidence reviewed in this Handbook.

III Theories of Memory and Economic Decisions

How can we incorporate a role for human memory into economic decision-making? In this section, we discuss some of the most prominent theoretical approaches that allow for more psychological realism regarding a role for the past and for memory.

We will organize the presentation by the four key features of experience effects that we have teased out in the prior section. After sketching the traditional economic "benchmark model" in Section III.A, we introduce a simple model with overlapping generations (OLG) in Section III.B that is closely motivated by the empirical facts and directly embeds all features (1) to (4) in the theoretical set-up (Malmendier et al., 2020). This (somewhat brute-force) approach allows us to see how the empirical features can be captured and to derive further implications. The following three sets of theoretical approaches, instead, work towards the deeper conceptual goal of teasing out the underlying mechanisms of memory formation. They vary, however, in their ability to capture the empirical facts (1)-(4). Case-based decision-making, presented in Section III.C, introduces a first notion of the concept of "similarity," and as such the idea that the inferences we make based on new data might not apply to all statistically correlated assets (or "domains"), i. e., a version of feature 3 (Domain specificity). The two-state modeling framework shown in Section III.D additionally incorporate feature 2 (Recency bias). And finally some of the models presented in Section III.E, which introduce the concepts of "context" and retrieval, also incorporate key feature 1 (Long-lasting effects), in addition to recency and a role for domain or context. Note that feature 4, the robustness of experience effects and memory to expert knowledge, is embedded in all modeling approaches as they aim to present psychologically more realistic models of belief formation and choice, which apply to all humans.

Another shared feature of all approaches presented in this section is the distinction of experience- and memory-based decision-making from informational frictions. Agents are not "uninformed" but choose to operate with a "flawed database" in the sense of emphasizing or deemphasizing potential future outcomes in deviation from known statistical probabilities. This feature reflects precisely the empirical evidence discussed above: individual decision makers are influenced by their personal exposure and lifetime experiences, even if they are well aware of the statistical distributions, such as FOMC members are aware of inflation data and still overweigh their personal past experiences.

We note that memories allow for a role of perceptual errors (Woodford, 2019) and may reflect utility from beliefs (Brunnermeier and Parker, 2005). In that sense, calling the database "flawed" may very well be incorrect. In fact, memory may be optimized (for related discussion see chapter 7.3/Gershman), given constraints and the true uncertainty in nature. Dissecting these foundations are important avenues to pursue and will require careful synthesis of economics and psychology outside of the scope of this review.

We preface the theoretical discussion with two final remarks: First, we only consider models that involve an economic decision-maker. We thus leave aside mechanistic models that characterize the psychology literature (Matter et al., 2019). Second, in what follows, models will often specify evolution in agents' beliefs. However, we refrain from taking a stand on what the agent thinks about him or her future self. In a dynamic model, a rational agent would correctly predict her own beliefs. We leave the important questions of the agents' view of her own future mental states, as well as the mental states of other agents' in the economy to future research.

III.A The benchmark model

In this section, we briefly describe the benchmark neoclassical theory of decision making under uncertainty. This theory assumes that there is a state of the world that will realize in some future period. We represent this future state of the world with the random variable Z, whose distribution depends on the parameter θ. Agents have a utility function over the outcomes of Z, which we denote u. This function is assumed to have some basic properties, such as being increasing and concave. Agents choose an action out of a set of possible action to maximize *expected utility*:

$$U(\theta, a) = E_{\theta,a}[u(Z)]. \tag{2}$$

In the introduction, we give an example of where the agent allocates wealth between a risky and riskless asset in which the stock return is uncertain. The decision depends on the unknown parameter θ. One approach to modeling the agent's decisions under this incomplete information is to assume that the agent's beliefs follow the laws of probability Savage (1954). Suppose the agent comes equipped with a prior and a likelihood function, and views data x. Bayes' rule implies a posterior $\pi(\theta|x)$. The agent maximizes expected utility, this time integrating out over the uncertainty in θ:[5]

$$\rho(\pi(\theta|x), a) = \int_{\Theta} U(\theta, a) dF^{\pi(\theta|x)}(\theta) \tag{3}$$

Here, $\pi(\theta|x)$ is the posterior distribution and $F^{\pi(\theta|x)}$ is the associated cumulative distribution function.

An important special case is one of full information, in which the agent treats θ as known. One could justify this as a limiting case where the amount of data goes to infinity. A more narrow definition of rationality is that of rational expectations, in which (2) holds, and the agents' prior and likelihood are such that the posterior π converges to the true distribution.[6]

[5] See (Berger, 1985, Chapter 4) for a formal justification.

[6] The economic literature makes various uses of the term rational expectations. One notion corresponds to a "communism of models" across researcher, source of uncertainty, and agents: The researcher uses the model that he or she believes determines the (stochastic) outcomes of economic

This rational expectations benchmark allows for the modeling of complicated systems from micro-economic foundations of decision-making. Being both tractable and flexible, it is no wonder it became the dominant paradigm, alas, at the cost of ignoring some of the most important empirical economic phenomena detailed above.

In what follows, we examine various economic models that *do* accommodate some or all of those phenomena, including the four stylized features.

III.B Models of Experience-Based Learning

The most straightforward approach to incorporating the field evidence reported in Section II, and to introducing a role for long-lasting memories into economic decisions, is to mirror the empirical findings about long-term experience effects and the tilt towards more recent experiences directly in theoretical assumptions about belief formation. Several papers have done so utilizing so-called overlapping generations (OLG) models, in which individual agents live for finite amounts of time and their live spans "overlap" with the those of earlier- and later-born generations. The dynamic OLG framework lends itself to studying how experiences and memory influence decisions in economic settings as different generations naturally have different lifetime experiences and memories, given past realizations.

Papers that employ the OLG setting to capture experience effects include Collin-Dufresne et al. (2017), Schraeder (2015), and Malmendier et al. (2020).[7] Here, we briefly sketch the OLG equilibrium framework of Malmendier et al. (2020), who also

processes as the model for uncertainty, and the agents have (or can learn about) the same model (cf. Hall and Sargent (2018)). Under an alternative definition, agents follow the principles of probability, but their posterior distribution may differ permanently from the true one. The posterior and the true distribution agree on what is possible (technically, they agree on the zero-probability sets), but the probabilities are different. This type of rational expectations is closely linked to the no-arbitrage ("no money pump") assumption and is common in finance.

[7] Also related, although not an OLG model, is Nagel and Xu (2019), who use a representative agent model with constant-gain learning which assumes that investors form beliefs about future dividend growth based on exponentially declining weights on observed dividends.

tease out implications for market dynamics.

In this OLG setting, a new generation is born at each point in time $t \in \mathbb{Z}$ and lives for q periods. Agents' objective is to maximize their per-period utility.[8] The choice they have to make, given their level of risk aversion, is whether to invest their current wealth in a risk-free asset with gross return $R > 1$ or in a risky asset, which pays a random dividend $d_t \sim N(\theta, \sigma^2)$ at time t. Crucially, agents do not know the true mean of dividends θ and use past observations to estimate it.

Experienced-based learning (EBL) is modeled as agents over-weighing realizations observed during their lifetimes when forecasting dividends (key feature 1 from Section II), and that they tilt the excess weights towards the most recent observations (key feature 2). The model captures domain specificity (key feature 3) in that EBL applies to specific stochastic process whose realizations the agent experienced, and is robust to learned knowledge (key feature 4) in that agents *choose* to under-weight information that they have not personally experienced. That is, even in a *full-information* setting where agents observe the entire history of dividends, prices do not add any additional information. Importantly, these models do not aim to model "forgetting" of other not-personally experienced outcomes.[9] Rather, we can think of the excess weight put on personally experienced outcomes, and especially on more recently experienced outcomes, as capturing the notion of availability bias in Tversky and Kahneman (1974) and, theoretically speaking, agents agreeing to disagree.

How do experience-based beliefs formally look like in this setting? Consider an agent endowed with a prior belief m about mean dividends.[10] As the agent experiences

[8] Agents' focus on per-period, rather than lifetime utility reflects that they are myopic. This assumption is common in the finance literature as it simplifies the maximization problem considerably while still highlighting the main determinants of choice – here, those generated by past experiences.

[9] It might seem that the forgetting is a primary memory-related phenomenon to be incorporated into economic models (Dow, 1991). Yet the idea that experiences are permanently forgotten was challenged very early in the memory laboratory and quickly disproven (McGeoch, 1932; Underwood, 1948; Estes, 1955).

[10] We can think of m as analogous to θ in the previous section: it is an unknown parameter that determines the distribution of interest to the agent.

dividend realizations, these experiences shape the new (posterior) beliefs about the mean of dividends. Namely, if the prior m is Gaussian with mean θ (the true mean of dividends), then the posterior θ_t^n, held at time t by an agent born at time n, can be expressed as a convex combination of the prior m with weight $1 - \omega_{age}$, and the agent's average lifetime experience of dividend realizations $(d_t, ..., d_{t-k}, ..., d_{t-n})$ so far with weight ω_{age},[11]

$$\theta_t^n \equiv (1 - \omega_{age}) \cdot m + \omega_{age} \cdot \sum_{k=0}^{age} w(k, \lambda, age) d_{t-k}, \tag{4}$$

where $age = t - n$, and where the lifetime average of experienced dividend realizations is formed with a weighting function

$$w(k, \lambda, age) = \frac{(age + 1 - k)^\lambda}{\sum_{k'=0}^{age} (age + 1 - k')^\lambda} \tag{5}$$

for the weight an agent aged age assigns to the dividend observed k periods earlier for all $k \leq age$, and $w(k, \lambda, age) \equiv 0$ for all $k > age$. This formula directly mirrors the empirically estimated model of Malmendier and Nagel (2011). Under (5), more recent observations receive relatively more weight if $\lambda > 0$, whereas the opposite holds if $\lambda < 0$. The empirical evidence in Malmendier and Nagel (2011) points to approximately linearly declining weights ($\lambda = 1$). That is, the empirically observed behavior indicates that agents put less weight on experiences further in the past and the weight on past experiences changes every j periods by roughly

$$w(k + j, 1, age) - w(k, 1, age) = -\frac{j}{\sum_{k'=0}^{age} (age + 1 - k')}.$$

for any $0 \leq k, k + j \leq age$ under the estimation that maximizes model model fit.

[11] The weight ω_{age} that agents assign to their experience beliefs amounts to $\omega_{age} \equiv \frac{age+1}{\tau + age + 1}$, which increases with age and decreases with the relative importance agents assign to their prior beliefs, regulated by parameter $\tau \in [0, \infty)$.

A crucial ingredient of the set up is that, by construction, beliefs θ_t^n do not necessarily converge to the truth as $t \to \infty$; it depends on how fast the weights for "old" observations decay to zero (i. e., how small λ is). When agents have finite lives, convergence will not occur. This is a sharp difference from *Full Bayesian Learning*, where agents use all available observations "since the beginning of time" to form beliefs, and differences in personal experiences do not play a role. There, the posterior mean converges (almost surely) to the true mean and the implications of learning vanish as time goes to infinity. Moreover, there is no heterogeneity in beliefs, i. e., all generations alive in a given period have the same belief about mean dividends. Under EBL, instead, agents learn from their own experiences, our model generates learning dynamics even as time diverges.

We see now how the model is directly built around key features 1 and 2, long-term experience effects and recency bias. As discussed, it is also consistent with key features 3 and 4, domain specificity and "imperviousness" to learned knowledge, as experiences are specific to the asset (or stochastic process) for which the agent has accumulated personal experiences, and as the learning process is applicable to all agents. However, the model would need to be extended to multiple assets and agents with different levels of information to fully tease out the implications of the latter two key facts.

The OLG set-up allows to derive several additional testable predictions.[12] First, it implies that asset prices depend on the history of dividends observed by the *oldest generation* still active in the market. Hence, variation in the number of generations participating in the stock market implies variation in how far back past events still affect stock trading today.

Second, the model reveals that the stronger dependence of prices on more recent

[12]In addition to the novel predictions discusse below, the EBL model produces a range of asset pricing implications, including known puzzles such as the predictability of stock returns (Fama and French, 1988), the predictability of the dividend-price ratio, and the excess volatility puzzle (LeRoy (2005) and Shiller (1981)).

dividends (recency bias) is increasing in the fraction of young agents in the market. This result reflects the fact that the dividends at time t are observed by all generations whereas past dividends are only observed by older generations.

More generally, the model establishes a novel link between demographics and the factors influencing asset prices: Each generation's past experiences and beliefs have predictive power for asset prices, and the relative magnitudes of the weights assigned to past dividend experiences (in predicting demand and prices) depend on the *number of cohorts* represented in the market and on the *fraction of each cohort* in the market – very much in line with the lingering, long-lasting "mood of pessimism" generated by adverse outcomes that was observed by (Friedman and Schwartz, 1963) after the Great Depression.

Another cross-sectional implication is that positive shocks (booms) induce a larger representation of younger investors in the market, while negative shocks (crashes) have the opposite effect. In the case of booms, for example, all cohorts become more optimistic, but the effect is particularly strong for the younger generations: Those born and growing up in periods of increasing dividends have a relatively higher demand for the risky asset than the older generations. This induces them to be over-represented in the market for the risky asset. The reverse holds for "depression babies," i. e., generations born during times of contraction. All cohorts of investors living through such times of depression will become relatively pessimistic about future returns. However, the effect is strongest on the younger generation, who will exhibit a particularly low willingness to invest in the risky asset, relative to older generations born during those times, and as a result they will be underrepresented in the stock market.

Some of these novel predictions of the EBL model have already been taken back to the data, including the latter prediction on the composition of market participants. Malmendier et al. (2020) provide and update (and refinement) of the evidence from Malmendier and Nagel (2011) that the EBL model predictions regarding the link between demographic composition and the cross-section of asset holdings are in line with evidence from micro-level data in the Survey of Consumer Finances (SCF) and the Center for Research in Security Prices (CRSP).

For example, they relate the differences in lifetime experiences between older and younger cohorts (i. e., those above 60 and those below 40 years of age) to the differences in their stock-market investment in the modern, triennial SCF since 1983. To extend the analysis of (Malmendier and Nagel, 2011) to more recent data and different experience measures, they construct an indicator of stock-market participation (which equals 1 when a household holds more than zero dollars worth of stocks) for the extensive margin and calculate the fraction of liquid assets invested in stocks for the intensive margin. They then calculate alternative life-time experience measures based on dividends, earnings, and GDP.

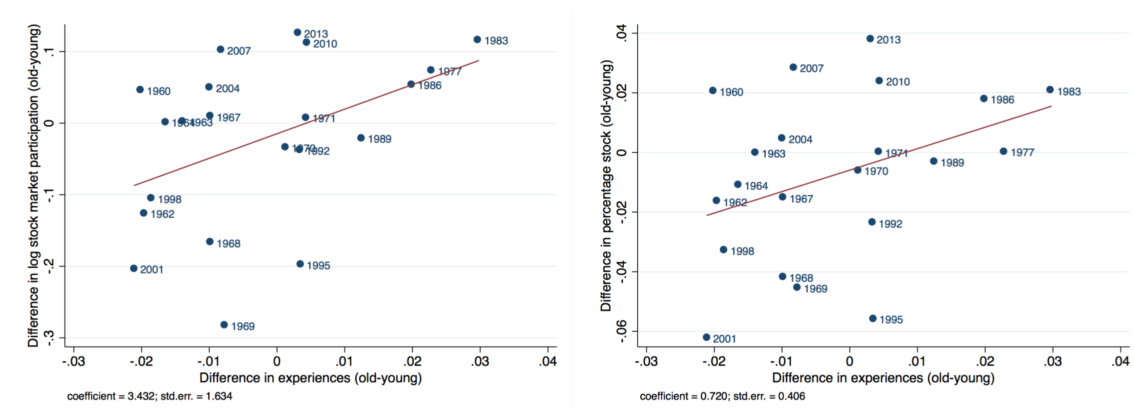

(a) Stock-market participation ($\lambda = 1$) (b) Fraction invested in stock ($\lambda = 1$)

Figure 6: Experienced Returns and Stock Holdings

Difference in experienced returns is calculated as the lifetime average experienced returns of the S&P500 Index as given on Robert Shiller's website, using declining weights with $\lambda = 1$ as in equation (5). *Stock-market participation* is measured as the fraction of households in the respective age groups that hold at least \$1 of stock ownership, held either directly or indirectly. *Fraction invested in stock* is the fraction of liquid assets stock-market participants invest in the stock market. We classify households whose head is above 60 years of age as "old," and households whose head is below 40 years of age as "young." Difference in stock holdings, the y-axis in graph (a), is calculated as the difference between the logs of the fractions of stock holders among the old and among the young age group. Percentage stock, the y-axis in graph (b), is the difference in the fraction of liquid assets invested in stock. The red line depicts the linear fit.

The results for all four performance measures and both for the extensive and intensive margin are in line with the predictions of the EBL model: The older age-group is more likely to hold stock, compared to the younger age-group, when they have experienced higher stock-market returns, dividends, earnings, or GDP in their lives so far. The opposite holds when the experiences of the younger generations are better than those of the older generations. Figure 6 shows the results for experienced returns with $\lambda = 1$. In Panel (a), the slope coefficient of the linear line of fit is significant at 5%. The steepness of the weighting function, and hence the extent of imposed weight on recent data points, makes little difference, as a comparison with calculations for $\lambda = 3$ reveals. The analysis of the intensive margin of stock-market investment yields the same conclusion, as shown in Panel (b). Here, the slope coefficient is significant at 10%. The corresponding analyses for experienced dividends, earnings, and GDP all reveal a positive relation of differences in experienced performance and stock investments between the young and the old.

Experience-based learning also has implications for the trade volume: When the change in a cohort's beliefs is different from the average change in beliefs, trade volume increases. That is, trade volume increases in the dispersion of changes in beliefs. From the analysis of demands, it follows that an increase (decrease) in dividends induces trade when it makes young agents more optimistic (pessimistic) than old agents. This mechanism is solely due to the presence of EBL, since it is essential that each generation reacts differently to the same dividend.

As shown in Figure 7, the data also confirm this prediction. Here the authors measure the dispersion of changes in disagreement among investors as the cross-cohort standard deviation of the change in experienced performance between the current year and the previous year. And to measure the abnormal trade volume, they calculate the deviation of the turnover ratio from its trend, first computing firm-level turnover ratios (number of shares traded over the number of shares outstanding) on a monthly basis,

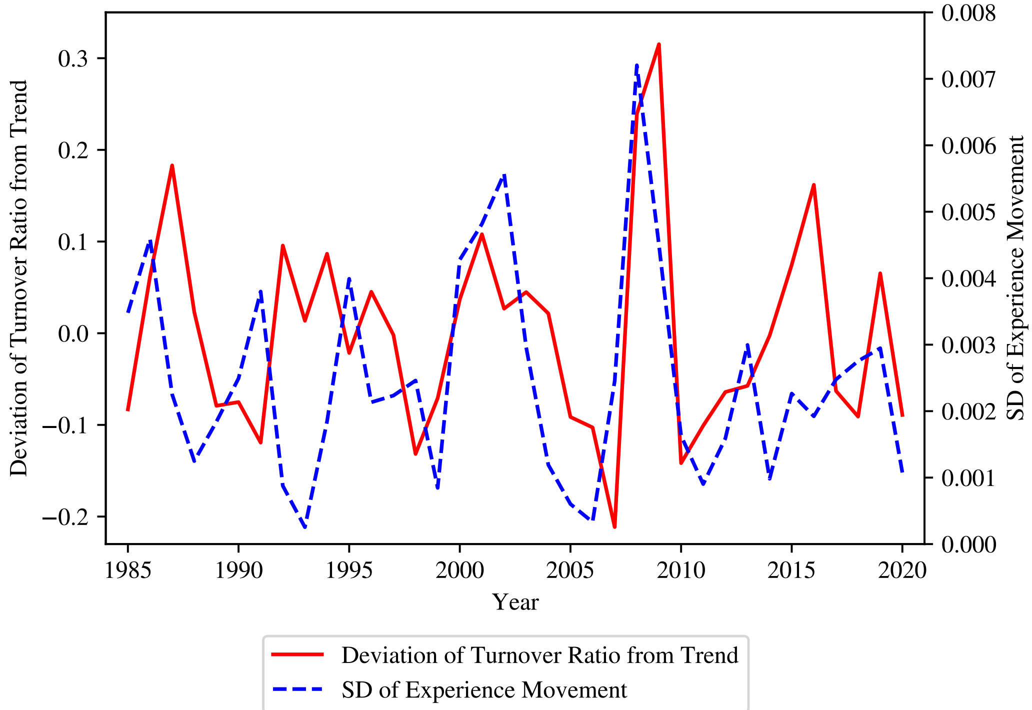

Figure 7: Trading Volume and Standard Deviation of Changes in Experienced Returns
Trading volume, shown in solid-red, is calculated as the market-capitalization weighted average monthly turnover ratio (shares traded divided by shares outstanding) across all firms in January and in December of the preceding year. The variable is logarithmized, linearly detrended, and Christiano Fitzgerald (CF) band pass-filtered to obtain the deviation of turnover ratio from the trend. Returns, shown in dashed-blue are defined as inflation-adjusted change in price from the prior year divided by inflation-adjusted price in the prior year. Returns are linearly detrended and CF filtered. The measure of experience-based disagreement uses the change in experienced returns for individuals of a given age from the experience of those individuals in the prior year, and is equal to the current-year age-cohort population weighted standard deviation of this difference variable for each year.

and then aggregating these numbers into a market-wide turnover ratio, weighting firms by their market capitalization.[13]

Figure 7 displays the trade volume in dark (blue) color, and changes in the experience-based disagreement about returns between cohorts in light (orange) color over time, starting in 1985. (Trading cost were significantly higher up to the mid-1980s, making it less likely that (individual) investors trade repeatedly based on experienced perfor-

[13] Given the cyclical component in the time series, the authors apply a filter, or decomposition, to obtain a smoothed representation of the underlying trend. Since the turnover ratio is non-stationary, they apply the (Christiano and Fitzgerald, 2003) band to the logarithm of the turnover ratio series, keeping frequencies between 2 and 8 years.

mance.) The figure displays a clear co-movement between disagreement among cohorts and trading volume, which is statistically significant at 1%.

In summary, the direct modeling of experience-based learning presented in this section helps illustrate that the demographic structure of an economy, and in particular the cross-sectional composition of investors, affect equilibrium prices, demand, and trade volume in a predictable direction. Since experiences vary across cohorts, a given dividend shock will have different long-term consequences depending on the demographic structure of the economy at the time of the shock. The EBL model predicts the direction of such differences. In addition to providing an explanation for several well-known asset pricing puzzles, that have been observed in the data, it also produces new testable predictions about the relation between demographics, prices trading behavior, and the cross-section of asset holdings, which are in line with the data.

By design, the modeling approach presented in this subsection captures, or is consistent with, all of the key features. It does not allow, however, to explore possible channels that induce the over-weighting of personal experiences or explain the role of "domains" or "context." In the next three subsections, we present different approaches that aim to capture the deeper underlying mechanism: case-based decision theory (in Section III.C), two-state models of memory (in Section III.D), and models of context and retrieval (in Section III.E). These approaches vary however in the extent to which they capture all of the stylized features, with case-based theories providing insights into only (3), domain specificity/similarity; two-state models capturing (3) and also (2), recency; and some of the context/retrieval models capturing (3), (2), and also (1), long-term effects. All of them are modeled as applicable to any decision-maker, including fully informed agents, and as such capture feature (4).

III.C Case-based decision theory

The earliest theoretical development towards a neuroscientifically more realistic model of choice might be embedded in the body of work on case-based decision theory. As we discussed in the prior subsection, the empirically motivated models of experience-based learning (EBL) aimed to convey the notion that, when making a decision, agents will are drawn to prior situations, or cases, where they faced similar circumstances, and might default to a decision that would have proven beneficial under those circumstances. Precisely this notion is at the root of a decision-theoretic approach proposed by Gilboa and Schmeidler (1995). They conjecture that, when agents face the problem of choosing an optimal action, they tend to rely on similar cases they have seen in the past.[14] In the words of Hume (1748), whom Gilboa and Schmeidler (1995) quote: "In reality, all arguments from experience are founded on the similarity which we discover among natural objects, and by which we are induced to expect effects similar to those which we have found to follow from such objects ... From causes which appear *similar* we expect similar effects. This is the sum of all our experimental conclusions." Indeed, modern psychological evidence supports this conjecture (Kahana and Wagner, 2019).

Consider an example that ties closely to the evidence on experience effects above: Suppose one wanted to rescue the economy from a once-in-century financial crisis? Here, Bernanke (2015) describes the decision process at the Federal Reserve:

> It was clear that the Fed needed to do more – but what? In response to current events, people often reach for historical analogies, and this occasion was no exception. The trick is to choose the right analogy. In August 2007 the analogies that came to mind – both inside and outside the Fed – were October 1987...and August 1998... With help from the Fed, markets had rebounded each time with little evident damage to the economy.

[14]PLACEHOLDER: possible cross-link to Chapter 4.4.

In this quote, we see the main elements of case-based decision theory – the reaching for similar cases, the evaluation of the closeness of the cases, the evaluation of the result. Past experiences and memory play a role: The current problem calls to mind similar problems in the past, the chosen actions, and the results. Even though Gilboa and Schmeidler (1995) do not have a theory of memory, the fact that a problem can call to mind similar problems is a central theme in memory research and one that will be explored in the sections that follow.

Gilboa and Schmeidler (1995) develop the formal case-based decision theory as follows. Let P be a finite set of *problems*, A a set of *actions*, and R a set of *results*, with r_0 (normalized to 0) being the result of not choosing an action. A *case* is an element of $P \times A \times R$. A memory is a set of cases $M \subset P \times A \times R$ such that

1. For every pair of problems and actions in the memory, there is a unique result.

2. For every problem, there is a unique action for which the result is not r_0.

Thus M maps from the set of problems and actions to results. Note that each problem is associated with a single action (the agent never experiences the same problem twice). The space $H(M)$ (H stands for history) denotes the problems the agent has experienced.[15] Though the agent experiences a problem only once, and hence has a unique action associated with the problem, Gilboa and Schmeidler (1995) think of M as having as domain the entire space $H(M) \times A$. The existence of r_0 allows them to do this. Every action not associated with the problem is assigned result r_0. The agent has a utility function u over results r.

According to the formal definition of case-based decision theory, an agent, when confronted with a new problem p, ranks actions in A. The ranking will depend on memory M and of course on the problem. Gilboa and Schmeidler (1995) show that, given a set of axioms over the agents' preferences, the agent maximizes utility

[15]More precisely, given a set M, $H(M) = \{q \in P \,|\, \exists a \in A, r \in R, \text{ such that } (q, a, r) \in M\}$ is the projection of M on P.

In this problem, there is concept of probabilities or states of nature. There are only possible actions, and their results in terms of utility. Gilboa and Schmeidler (1995) argue based on first principles that the agent should choose the action based on the result of the action under problems similar to the current one. More precisely, they formulate axioms such that there exists a similarity function s from $P \times P$ to $[0, 1]$, such that the agent chooses the action to maximize the similarity-weighted outcome:

$$U_{p,M}(a) \equiv \sum_{(q,a,r) \in M} s(p,q)u(r).$$

In other words, when confronted with a decision problem p, the agent chooses the action a that, given her memory M, maximizes her utility, and that utility maximization process will be determined by all the problems q she had been confronted with in the past and the action and results they generated, weighted by the similarity between those past problems and the problem at hand.

Note the stark difference to the standard modeling approach introduced in Section III.A. Here, the preference ordering over actions is not determined by the posterior expected loss as in (3). Rather, given a problem and a history of past problems, the agent ranks actions. In fact, case-based decision theory dispenses entirely with the notion of beliefs, which are central to the benchmark model in Section III.A. As such, it is a sharp departure from the Savage (1954) framework itself (and thus the later assumptions of rational expectations that build on it). The theory is fundamentally non-Bayesian – the agent need not have subjective beliefs over states of nature and indeed such states of nature need not even be defined. In the words of Gilboa and Schmeidler (1995), "no hypothetical thinking is assumed." This is a crucial philosophical difference between case-based decision theory and the standard model.

Most importantly to our main theme, case-based theory makes formal use of the notion of memory (memory is an abstract set) and similarity (the axioms leave the form of similarity unrestricted; it need not represent similarity the way that we think of it). However, in its use of these concepts (even just as abstractions) it previews later work and paves the path for "past experiences" to exert a long-lasting influence on decision-making and the the "domain" or "context" of those past experiences to be a relevant determinant via the notion of similarity.

One noteworthy feature is that there is no built-in notion of recency, our key feature 2. Problems all appear equally in memory regardless of when they occurred, waiting to be called up by similar problems. On the other hand, while not explicitly discussed, it embeds the notion of long-lasting effects (key feature 1).

III.D Two-State Models

One of the earliest approaches to introducing memory directly into economic decision-making is the model of Mullainathan (2002). His hypothesis is that agents predicting future realizations of a stochastic variable combine "hard" information, readily at hand, and "soft" information, available only from memory and therefore flawed.

Mullainathan (2002) uses the context of beliefs about future income and consumption choices to introduce his model. A perennially important question for economic policy is how agents perceive their future income. Income evolves over time and is subject to shocks; a key question is how individuals infer the parameters of this process from observations. If income goes up today, does this mean that future income will be higher, or is the change merely temporary? The consumption response to a given change in income is called the *marginal propensity to consume* (MPC). Rational expectations implies a striking finding for the MPC – individuals should only consume out of "permanent" changes to their income, not out of temporary ones (see Deaton (1992) for a review). As we have seen above, in the discussion of the findings in Malmendier and Shen (2018), this is not the case.

Mullainathan (2002) models an agent estimating income at time t, denoted y_t. The agent will form a forecast at time $t-1$ based on previous history, and based on a new observation, namely, an "event," which will convey a noisy signal of y_t. The resultant estimate is called $\hat{y}_t$. In the language of Section III.A, the agent has a loss function based on the distance between the true y_t and the forecast $\hat{y}_t$. The true y_t is the state of nature (θ). Hence, the Mullainathan (2002) model nests the case of Bayesian updating.

Mullainathan (2002) assumes a standard model for the evolution of income. Income is subject to permanent and transitory shocks. That is

$$y_t = \sum_{k=1}^{t} \nu_k + \epsilon_t,$$

where ϵ_t and ν_t are independent normal shocks. Note that shocks ν_t have a permanent effect whereas shocks ϵ_t have a transitory effect. Clearly the decomposition of observed y_t into its permanent and transitory component is important for the estimation of future y_t: the permanent component should enter the expectation whereas the transitory component should be ignored. In this sense, forecasting y_t is closely connected to deviations from the permanent income hypothesis.

In the benchmark Bayesian framework, the agent observes ν_t with error:

$$\nu_t = \xi_t + z_t,$$

with ξ_t observed and z_t normally distributed and independent of all other shocks. The standard Bayesian solution equals

$$\hat{y}_t = \xi_t + \sum_{k=1}^{t-1} (\lambda^{t-k} \xi_k + (1-\lambda^{t-k}) \Delta y_k),$$

where

$$\lambda = \frac{\sigma_\epsilon^2}{\sigma_*^2 + \sigma_\epsilon^2}$$
$$\sigma_*^2 = \frac{1}{2}\left(\sigma_v^2 + \sqrt{\sigma_v^2(\sigma_v^2 + 4\sigma_\epsilon^2)}\right).$$

That is, the agent estimates income at time t to be the sum of the most recent signal ξ_t and the sum of the weighted averages of past realizations of ξ (i. e., all past signals of the permanent shocks to income) and the actual differences (over time) in past realizations of income, where the weighting of past realizations of signals and income differences is shifted more and more towards the actual past income realizations the further we go back in time. The noisier the transitory shocks to income are (the higher σ_ϵ^2), the higher are the weights assigned to the observed signal ξ of the permanent income shocks, as opposed to the actual difference in income realization.

Mullainathan (2002) incorporates memory as follows. Income y_t is "hard" information ("readily available in records"), whereas ξ_t is soft information that cannot be recorded.[16] Rather than viewing ξ_t, the agent sees instead an event e_t:

$$e_t = (\xi_t, n_t) \sim N(0, \Sigma); \qquad \Sigma = \begin{bmatrix} \sigma_\xi^2 & \sigma_{\xi n} \\ \sigma_{\xi n} & \sigma_n^2 \end{bmatrix} \tag{6}$$

The event contains ξ_t, and irrelevant information n_t. It would seem that this is no less information than in the Bayesian case, and that is true for the current realization of ξ_t when forming beliefs at time t. However, a key difference is that the agent does not perfectly recall previous values of ξ_t. Specifically, at time t, the agent possesses a recalled history h_t^R that consists of the events and of the previous hard information.

[16]One might object that there is no inherent reason why ξ_t could not be written down as well. The model is silent on the friction that prevents the recording of ξ_t. This is a gap found in many current formulations of bounded rationality.

The true history equals

$$h_t = (e_1, e_2, \ldots, e_{t-1}, y_1, y_2, \ldots, y_{t-1}).$$

In the recalled history, the agent imperfectly remembers events. That is, previous events are multiplied by a random variable R that equals zero or 1. If it equals zero, the event is not recalled at time t. Thus recalled history equals:

$$h_t^R = (e_1 R_{1t}, e_2 R_{2t}, \ldots, e_{t-1} R_{t-1,t}, y_1, y_2, \ldots, y_{t-1}),$$

where, for $k < t$,

$$R_{kt} = \begin{cases} 1 & \text{with probability } r_{kt} \\ 0 & \text{otherwise} \end{cases}$$

and conditional on r_{kt}, R_{kt} is independent of the history up to time t, and all other time-t shocks.

An important aspect of the model is the specification of the recall probabilities r_{kt}. This specification draws on memory theory, and we will discuss it shortly. First, though, let's solve for the beliefs of the agent. Given the recalled history h_t^R, the agent substitutes, for ξ_k, $k < t$, the recalled ξ_t:

$$\hat{y}^R(h_t^R, \epsilon_t) = \xi_t + \sum_{k=1}^{t-1} (R_{kt} \lambda^{t-k} \xi_k + (1 - \lambda^{t-k}) \Delta y_k). \tag{7}$$

The agent's flawed memory replaces ξ_k with its remembered counterpart. Notice a key feature of this model: at every point in time a *different* history (a distinct flawed database) is available to the agent. Moreover, in this new history, events are either fully remembered or not. An event that is accessible to the agent in one period may become inaccessible in the next.

Turning to the probability of recalling an event at time $k < t-1$, the model specifies

$$r_{kt} = \underline{m} + \rho R_{k,t-1} + \varphi a_{kt}, \tag{8}$$

with

$$a_{kt} = \frac{1}{2}\left(e^{-(\xi_t - \xi_k)^2} + e^{-(n_t - n_k)^2}\right)$$

and with $\underline{m}+\rho+\varphi < 1$, and with $r_{t-1,t} = 1$. First note that $\underline{m}$ presents the base rate of recall; it is the probability of recollecting an event if the event were not recollected last period, and if the event is sufficiently dis-similar from the current event ($a_{kt} \approx 0$). Next, note the presence of a_{kt}. The Mullainathan (2002) model incorporates the principle of semantic similarity. The agent recalls the event with greater probability if it is similar to the current event. The similarity may be along the dimension of interest (ξ_t), or along an "irrelevant" dimension n_t (this is the reason for n_t in Equation 6).[17] Finally, and perhaps of greatest interest, an agent is more likely to recall an event if he or she recalled the event in the previous period. This is captured by $\rho R_{k,t-1}$. Because the agent recalls the most recent event perfectly, this *rehearsal* mechanism creates a recency effect, and creates it through a psychologically plausible channel. (See also Chapters 11.1 and 11.3.)

In fact, (8) is a member of a well-studied and highly influential group of models within the memory literature: models of "primary" and "secondary" memory (more colloquially, short and long-term memory). The agent recalls just-experienced events with probability 1. Events that were just recalled, receive a boost in memory. Thus recent events are more likely to be recalled, with the advantage declining exponentially in ρ. Once an item falls out of primary memory, it is in "secondary" memory. It

[17]Given that n_t is a normally distributed variable, it may be a stretch to think of it in terms of semantic similarity, in which two items are similar because they share the same meaning. Perhaps a better interpretation of n_t is that it represents an exogenous context. This interpretation will be developed further in the next section.

is subject to the baseline recall probability $\underline{m}$, or to whether a similar event should randomly occur through associativeness. If it happens to be recalled, it is back to primary memory again. One can think of the model as isomorphic to a two-state Markov chain. Events can be in the good state (namely recently recalled); as long as they stay in that recently-recalled state, their probability of recall is high. With some probability, events fall out into a low-recall state. Similar models in the memory literature include those of Restle (1965) and Murdock (1967).

As the previous discussion indicates, two-state models as in the one in Mullainathan (2002) are highly successful in accounting for recency. They do not, however, account for temporal contiguity: subjects' tendency to recall items that occurred contiguously in time with recently-recalled items. Equation (8) implies that experiencing an event will imply recall of semantically similar events; those semantically similar events are more likely to be continued to be recalled because they have re-entered primary memory. However, there is no effect on recall of events experienced nearby in time.

Let's explore some of the properties of memory-based forecasts in Mullainathan (2002). A first property is over-reaction in the sense that

$$\mathrm{Cov}(y_t - y_t^R, \xi_t) < 0. \tag{9}$$

When ξ_t is positive, $y_t^R > y_t$ (all variables are mean zero). That is, the agent over-reacts, relative to the truth, to information contained in ξ_t. This is because of associativeness. An observation ξ_t brings to mind previous ξ_k observations with higher probability if they resemble ξ_t. Thus y_t^R will be biased toward ξ_t.

Second, while equation (9) shows that the agent overreacts to news about ξ_t, the reaction to changes in y_t is indeterminate and could be either negative or positive:

$$\mathrm{Cov}(\hat{y}_{t+1}^R - \hat{y}_t^R, y_t - y_{t-1}) \neq 0. \tag{10}$$

A negative sign is consistent with the over-reaction effect discussed above; a positive sign could occur because the agent is slow to learn new information.

In both of these measures of over- and under-reaction, the expectation is taken with respect to the population measure of income (namely from the point of view of a statistician with sufficiently long data). From this statistician's point of view, the agent is irrational because forecast errors depend on time-t variables (9), and revisions in forecasts are predictable (10). In contrast, the optimal Bayesian forecast error must be uncorrelated with time-t variables, and changes in forecasts must be unpredictable. Of course, the agent's probability distribution for y_t is not the Bayesian one.

Thus the memory model in Mullainathan (2002) can account for the slow response of consumption to changes in income. It also provides a potential explanation for why the marginal propensity to consume might be too high, even in the absence of borrowing costs. The agent will over-react to certain changes to permanent income because they will bring to mind other changes in permanent income, thus leading the agent to believe he has more data than he or she in fact does. Irrelevant information on the income stream can also matter, in that it brings to mind income observations when such information was previously observed.

This classic paper effectively integrates insights from memory to offer a potential solution to deep puzzles in macroeconomics. However, more work appears to be required to connect the empirical results (say, the tendency to consume out of windfalls) to these aspects of the model.

More broadly, Mullainathan (2002) provides us with a framework that is effective in capturing both a role for recency (key feature 2) and a role for similarity (akin to what we dubbed domain specificity in key feature 3). And of course, by framing the model as a theoretical conceptualization of memory it is designed to apply to everybody, even the well-trained statistician (key feature 4). What is still missing is a discussion of the long-lasting influence of past events (key feature 1): Are there features that make events more likely to remain in primary memory for a long time, even when no similar events have occurred?

III.E Context and Retrieval

The last set of models presented in this chapter make significant progress in laying out and specifying a role for "context" and addressing when which past experiences will be retrieved and utilized in the decision-making at hand. As we will see, this approach allows to capture the empirically observed features of experience effects and memory-based choices from the field studies in economics.

A first step towards incorporating a role for context to affect economic choices was proposed by Bordalo et al. (2019b). Their notion of context is that it is embedded in the "current environment" akin to Godden and Baddeley (1975). That is, agents evaluate attributes of a choice, such as whether or not to buy a product, based on current conditions, even if these conditions are clearly temporary and the product is a durable one. Examples include an increase in the likelihood of buying (i) cold-weather items on cold-weather days which are subsequently returned (Conlin et al., 2007); (ii) convertibles on warm-weather days (Busse et al., 2015); (iii) health insurance on days (in China) with high air pollution (Chang et al., 2018), policies which are then cancelled if air quality improves during the free cancellation period.

Bordalo et al. build on a classic framework within the memory literature: geometric similarity (Kahana, 2012). According to this model, the memory system calculates a distance between a cue and a prior experience. Suppose that the agent is cued by context x_t.[18] Similarity relative to some other context in memory, x, equals

$$S(x, x_t) = \exp\{-\delta(x_t - x)^2\}. \tag{11}$$

Now suppose the agent is attempting to evaluate the quality of an object (for example, a coat, a vehicle, an insurance policy). The agent does not observe "true" quality, and thus infers quality based on experience as follows. The agent considers current context x_t and calculates (11) for all prior experiences (quality observations) and weights them according to their similarity with the current context to form an estimate of quality:

$$q^n(x_t) = \int \left(w(x, x_t) \int q dF(q \,|\, x) \right) dF(x), \tag{12}$$

where weights $w(x, x_t)$ equal $\frac{S(x,x_t)}{\int S(\tilde{x},x_t)dF(\tilde{x})}$, and where $F(q \,|\, x)$ and $F(x)$ denote the conditional and marginal cumulative distribution functions respectively. Bordalo et al. (2019b) call q^n a quality "norm," but it can also be thought of as a quality estimate. Given a price p_t, the agent decides to make a purchase as long as $q^n(x_t) - p_t > 0$. The agent draws on memory to evaluate quality, and memory depends on context.

How might this explain the data? The context x_t might, for instance, be a cold-weather day. The agent retrieves higher quality norms, because these were experienced on cold-weather days. That is, $\mathbb{E}[q \,|\, x_t]$ is unusually high, leading to higher $q^n(x_t)$. The agent is confused: while it is true that putting on a coat now might be a wise action, the decision to order the coat from the catalogue should be independent of any particular day's temperature. The agent should be calculating unconditional quality

[18] Note that (11) is not exactly the same as the similarity functions in the memory literature. Geometric similarity models are, as their name implies, based on geometric distance. The distance between a cue and a memory is not, however, the sum of squared distances but rather the square root of the sum of squared distances. However, under the assumption of normality, (11) is formally equivalent to Bayesian updating and leads to convenient analytical solutions.

and yet is unable to do so. In relation to the theory in Section III.A, the action is whether to buy an item, and where the unknown state of nature is the quality. The agent uses a biased database (incorrectly informed by context) to estimate the quality.

Bordalo et al. (2019b) also aim to explain a second type of data that relates to how a price norm is constructed when the agent knows the true quality of an object, but memory makes her account the "usual" price in another location. For example, renters who move to a less expensive city spend more on housing than locals, whereas renters who move to a more expensive city spend less (Simonsohn and Loewenstein, 2006). In related experimental findings, Simonson and Tversky (1992) show that willingness to pay in a second stage of an experiment depends on prices presented in the first stage.[19]

To fold in this second type of evidence, Bordalo et al. (2019b) introduce an altogether different role for memory, namely to construct a price norm based on past prices (e.g., rental prices in a prior location), yet at the cost of abandoning the prior concept of context: physical context is exactly what does not work to construct the "norm" in this example. A move from San Francisco to Pittsburgh involves a change in physical context.[20]

To summarize, Bordalo et al. (2019b) propose an elegant and parsimonious explanation for why the environment might exert a significant influence on economic decision-making, and as such conceptualize what we dubbed the "domain specificity" of experience effects (key feature 3). As in Mullainathan (2002), their explanation clearly draws from insights in memory theory to understand a phenomenon outside the reach of standard economic models.

Differently from Mullainathan (2002) (as well as Malmendier et al. (2020)), however, they ignore the time dimension of memory-driven choices. Hence there is no role for

[19] Nakamura and Steinsson (2008) show indirect evidence of this effect in that stores offer infrequent large sales in which they display the "usual" price.

[20] The fix that Bordalo et al. (2019b) propose is to call "renting an apartment" the context. Given that previous examples of context pertained to the physical environment, it appears to ascribe perhaps too much flexibility to the notion of an external context.

recency effects (key feature 2), nor room to tease out the long-term effects of past experiences (key feature 1). While one might attempt to capture the time dimension under the umbrella of "similarity" (with more recently experienced contexts being more similar), the existing evidence on brain functioning and neuroplasticity (synapse formation and decay) would instead point to a separate role.

This is where the model of contextual encoding and retrieval by Wachter and Kahana (2021) comes in. Wachter and Kahana provide a theoretical framework of context and retrieval that incorporates the temporal dimension of memory and in particular a role for very-long-term experiences. They directly pose the question of why, based on the existing evidence from memory research, early-life experience should have such long-lived effect? How does one reconcile the long-lived nature of memory with the greater weight of recent experiences?

In addition, Wachter and Kahana (2021) generalize the conceptualization of "context" beyond the physical features of the environment. They argue that the existing evidence pushes toward a notion of context that is more abstract than simply a physical environment (see also Chapter 5.12). Reading about a stock market crash or viewing a horror movie do not affect the physical environment, and yet choices are affected by these experimental manipulations (Cohn et al., 2015; Guiso et al., 2018). As they point out, it is also the case that models of recency, of semantic similarity, and of physical context miss the temporal contiguity effect originally stated by Aristotle and found in experimental data.

The modeling framework of Wachter and Kahana (2021) builds on Howard and Kahana (2002) to propose a model of temporal context. In contrast to, say, the notion of context as environment in Godden and Baddeley (1975), used in Bordalo et al. (2019b), context is a fully latent and endogenous state. The Wachter and Kahana (2021) model incorporates this idea of context into decision-making and uses it to explain evidence on the effects of early-life experience, as well as the effects of irrelevant cues such as scenes from a movie.

The temporal context model, as its name implies, relies on the notion of time, and the state of the world potentially evolving through time. As in Section III.A, assume an underlying state of nature, represented by the stochastic process Z_t, taking on potentially finitely many values $\{z_1, \ldots, z_m\}$. However, to give the model a time component, assume the state is persistent, and define a probability of transitioning from one state to another as $p^Z_{ik} \equiv \text{Prob}(Z_{t+1} = z_k \,|\, Z_t = z_i)$. The agent may not perfectly observe outcomes Z_t, but does observe a related random variable Y_t that takes on finitely many values $\{y_1, \ldots, y_n\}$. Nature provides a mapping from Z_t to Y_t: let $p(y_j|z)$ denote the probability that $Y_t = y_j$ conditional on $Z_t = z$ for $j = 1, \ldots, n$.

The observable state Y_t is represented by the *features* of the environment. Features are generally characterized by an $n \times 1$ column vector f_t such that

$$f_t(j) = \begin{cases} 1 & \text{if} \quad Y_t = y_j \\ 0 & \text{otherwise.} \end{cases} \tag{13}$$

That is, $f_t = e_j$, the jth standard basis vector in n-dimensional space, where j corresponds to the state of Y_t. Given features, the following recursive equations (together with initial conditions) define the latent context:

$$x_t = (1-\zeta)x_{t-1} + \zeta x^{\text{in}}_t, \tag{14}$$

$$x^{\text{in}}_t = \frac{M_{t-1} f_t}{||M_{t-1} f_t||}, \tag{15}$$

$$M_t = M_{t-1} + x_t f_t^\top. \tag{16}$$

Here, x_t and x^{in}_t are $m \times 1$ vectors. The matrix M, called the “memory matrix,” is $m \times n$. Note that (15) implies that x^{in} (where the “in” stands for input) is of unit length. Equation (14) implies that current context x_t is a weighted average of past context x_{t-1} and retrieved context x_t^{in}.[21] The parameter $\zeta \in [0, 1]$ governs the rate of decay of context, with ζ closer to one implying a faster rate of decay.

The model is associative in the sense that "cueing" with a feature vector recovers the contexts previously associated with those features

$$\begin{aligned} x_t^{\text{in}} &= \frac{M_{t-1} f_t}{||M_{t-1} f_t||} \\ &\propto M_0 f_t + \sum_{s=0}^{t} (x_s f_s^\top) f_t \\ &= M_0 f_t + \sum_{s=0}^{t} x_s (f_s^\top f_t). \end{aligned}$$

Note that for orthonormal basis vectors, $f_s^\top f_t$ is either 0 or 1, depending on whether $f_s = f_t$ or $f_s \neq f_t$. If the agent experiences a feature under multiple contexts, then the feature recalls a weighted average of the contexts. In a more general model, features need not be represented by basis vectors, and the interpretation still holds.[22]

To motivate the link between this memory model, which makes no mention about beliefs, and beliefs and ultimately choice for the agent, consider what M_t would represent should Y and Z both be observed. Equation 16 implies that M_t represents the counts of co-occurrences of Y and Z-states in the agent's experience. Because all that matters is the relative values of M_t (due to the scaling in (15), M_t encodes Bayesian posterior probabilities assuming a multinomial model and a flat prior. Furthermore, (15) implies that $x_t^{\text{in}}(i)$ is the conditional probability of Z_t given Y_t. This analogy, while helpful, is also limited, in part because it says nothing about the dynamics of

[21] Wachter and Kahana (2021) scale x^{in} by the sum of its elements, which implies (14) holds exactly. The memory literature, e. g., Polyn et al. (2009), typically uses the L^2-norm, with $x_t = \rho_t x_{t-1} + \zeta x_t^{\text{in}}$ and $\rho_t \approx 1 - \zeta$, to maintain x_t on the unit circle.

[22] Consider how this model might account for the finding of Simonsohn and Loewenstein (2006) discussed above. The search for an apartment is a feature of the environment, bringing to mind the internal mental state of previous searches, making it easier to remember the rental prices. As Bordalo et al. (2019a) report, for the subsample of renters who move twice, only those who have not experienced similar rents in the past exhibit the effect of the city of origin. In this, we see again the powerful effects of early experience.

beliefs, and in part because it does not yet describe how the agent functions in a world where Z_t is unobserved. To understand the model's implications, we first note some immediate properties, and then go to several examples.

Wachter and Kahana (2021) takes context $x_t(i)$ to be the agent's beliefs about each latent state. Equation (14) implies that context is persistent, but eventually decays. It is a weighted average (with exponentially decaying weights) of previous contexts retrieved from experiences. Because of this exponential decay, more recent experiences hold more weight in the agent's mind, suggesting a channel for the importance of recency in forming beliefs. A second immediate implication is that events that share certain features will retrieve a similar context and similar beliefs, even if the underlying data generating process producing the features is quite different. If the agent observes some feature vector f_j with $f_j^\top f_t \approx 1$, conditional probabilities at time t are as if f_t was observed.[23]

These implications pertain to how agents form associations in memory and use them to understand latent states and their associations with observables. Ultimately, agents must translate probabilities on underlying states Z (given by context) to those on outcomes Y. While nature supplies $p(y|z)$, they must learn this distribution from observations. Just as context is retrieved from features, there is an additional step by which features are retrieved from context. Specifically, given a basis context vector $\hat{e}_i$, we define retrieved features as follows:

Just as, in the memory literature, it is necessary to link retrieval of features to context, so it is that Wachter and Kahana (2021) must link beliefs about the unobservable

[23]More formally, for a series of features vectors $\{f_t^n\}$ with $\lim_{n\to\infty} f_t^n = f_t$, we have that $\lim_{n\to\infty} x_t^{\text{in},n} = x_t^{\text{in}}$, with a rate of convergence independent of t.

Table 1: Features corresponding to gains, losses, and depressions

Basis Vector	Features	Outcome for wealth
e_1	gain	$1 + \pi(\mu + \sigma) + \ell$
e_2	loss	$1 + \pi(\mu - \sigma) + \ell$
e_3	depression	$1 + \pi(\mu - \sigma)$

state Z to the observable state Y. This is done through traditional features retrieval:

$$\begin{aligned} f_t^{\text{in}} &= \sum_{i=1}^{m} x_t(i) f_{i,t}^{\text{in}}, \\ f_{i,t}^{\text{in}} &= \frac{M_{t-1}^{\top} \hat{e}_i}{||M_{t-1}^{\top} \hat{e}_i||}. \end{aligned}$$

The retrieved feature vector f_t^{in} represents a subjective probability distribution over features.[24] This retrieval process shares with context an associative property: the agent retrieves features experienced under similar contexts.

Taking features retrieval into account allows the model to produce a psychologically plausible theory of disagreement, and the long-run persistence in beliefs observed in the literature. To illustrate consider a problem similar to the portfolio choice problem outlined at the start of this chapter. Suppose the latent state Z_t corresponds to normal economic conditions ($Z_t = z_1$) or adverse conditions ($Z_t = z_2$). The latter occurs with unconditional probability $p < 1/2$. The agent decides to invest wealth between a stock with risky return $\tilde{r}$ and a bond with zero net return. Let π denote the percent allocation to the stock. The stock can experience either a gain or a loss with equal probability: $\tilde{r} \in \{\mu+\sigma, \mu-\sigma\}$ for $\sigma > \mu > 0$. The investor also receives risky labor income $\tilde{\ell}$, which is either positive ($\tilde{\ell} = \ell > 0$) or zero. We interpret the state in which a stock market loss co-occurs with zero labor income as a depression; assume that a stock market gain cannot co-occur with zero labor income. Thus, there are three observable states (outcomes of Y_t), summarized in Table 1.

[24]Note that $||f_t^{\text{in}}|| = 1$ because it is a weighted average of unit vectors $f_{i,t}^{\text{in}}$.

Under a standard specification for utility in which preferences are an increasing function of the mean and a decreasing function of variance, the optimal portfolio equals

$$\pi(p) = \frac{\mu - p\sigma\ell}{\sigma^2}, \tag{17}$$

a decreasing function of the probability of an adverse state.

Fix a time t for the agent to make a decision. Wachter and Kahana (2021) assume that, prior to time t, the agent has a record of associations:

$$M_{t-1} \propto \begin{bmatrix} \frac{1}{2} & 0 & 0 \\ 0 & \frac{1}{2} - p^* & p^* \end{bmatrix}, \tag{18}$$

so that loss and depression states share a context. This can arise if, for example, a stock market loss were to have occurred close in time to a depression. The special case of $p^* = p$ implies that the agent's associations correspond to the correct joint population frequencies of labor market and stock market states.

To focus on the novel implications of the model, consider the simplified case of $\zeta = 1$.[25] This implies the following recursion for context:

$$x_t = x_t^{\text{in}} \propto M_{t-1} f_t. \tag{19}$$

A stock market gain ($f_t = e_1$) retrieves

$$x_t \propto M_{t-1} e_1 \propto \begin{bmatrix} 1 \\ 0 \end{bmatrix},$$

[25] Wachter and Kahana (2021) extend these results to $\zeta < 1$.

which in turn retrieves features

$$f_t^{\text{in}} \propto M_{t-1}^\top \begin{bmatrix} 1 \\ 0 \end{bmatrix} \propto e_1.$$

A stock market loss ($f_t \in \{e_2, e_3\}$), on the other hand, retrieves context

$$x_t \propto M_{t-1} e_2 \propto \begin{bmatrix} 0 \\ 1 \end{bmatrix},$$

which retrieves features

$$f_t^{\text{in}} \propto M_{t-1}^\top \begin{bmatrix} 0 \\ 1 \end{bmatrix} \propto \begin{bmatrix} 0 \\ 1 - 2p^* \\ 2p^* \end{bmatrix}.$$

That is, the agent recalls a positive probability of a depression, *regardless of whether one has occurred*. The key difference between this model and the Bayesian benchmark where the agent observes both labor income and stock return realizations is that this act of recollecting implies that there is a new "depression" observation in the agent's mental database.

Suppose the agent encodes this experience into memory using the retrieved features. Consider then what happens to the M matrix:

$$M_t = M_{t-1} + x_t f_t^\top,$$

where

$$x_t f_t^\top = \begin{cases} \begin{bmatrix} 1 \\ 0 \end{bmatrix} e_1^\top & = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} & \text{if } gain \\ \begin{bmatrix} 0 \\ 1 \end{bmatrix} [\,0 \;\; 1-2p^* \;\; 2p^*\,] & = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 1-2p^* & 2p^* \end{bmatrix} & \text{if } loss. \end{cases}$$

It follows that, after τ periods of which k are gains,

$$M_{t+\tau} = \begin{bmatrix} \frac{1}{2}t + k & 0 & 0 \\ 0 & (\frac{1}{2} - p^*)t + (1-2p^*)(\tau - k) & p^*t + 2p^*(\tau - k) \end{bmatrix}. \tag{20}$$

Hence, in the limit as τ approaches infinity, the relative probability of losses and depressions, $M(2,2)/M(2,3)$, remains the same, regardless of how much experience the agent accumulates.[26] It does not matter how much data the agent observes: the probability of a depression remains distorted.[27] This illustrates the claim of long-run belief stability discussed above.

Why, intuitively, does the agent fail to update his or her probabilities? The reason is that the agent's memory over-associates a stock market loss with a depression. The appearance of a stock market loss then reinstates the depression context. This act of recalling the depression context is similar to experiencing the depression. Thus, a 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 (19) would have reached the correct probability p.

As a second example, consider an application of the similarity effect discussed above. Guiso et al. (2018) find that professional investors in Italy required double the premium to accept a lottery following the financial crisis than before. To rule out

[26]Note that $((\frac{1}{2} - p^*)t + (1-2p^*)(\tau - k))/(p^*t + 2p^*(\tau - k)) = (\frac{1}{2} - p^*)/p^*$ regardless of τ or k.
[27]The agent does, however, learn the correct probabilities of stock market gains and losses.

the explanation that the apparent shift in risk aversion followed from a decrease in wealth (as it would if, say, agents exhibited decreasing relative risk aversion), they conducted an experiment with undergraduate students as subjects. The students were assigned at random to watch, or not watch, a scene from a horror movie. The authors found that the students who watched the scene required a 50% greater premium to accept the lottery. Similarly, Cohn et al. (2015) report results from an experiment on financial professionals, in which some viewed a fictive chart of a booming stock market, while others viewed a chart with a market crash. In both cases, professionals answered questions about their trading strategies during the event in question. They then performed an investment task. Investors in the boom condition invested 17 pp more in the risky asset than did those in the bust condition. The results from these experiments are striking in that fear alone, as opposed to new information, has a substantial effect on risk taking. Here, we apply retrieved-context theory to explain how an emotional experience can change portfolio holdings.

We hypothesize that fear operates through the memory channel (see Chapter 3.7). 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 is the focus of this section. The setup is similar to that of the previous example. The agent chooses an investment π into a 50/50 lottery with return $\tilde{r} \in \{\mu + \sigma, \mu - \sigma\}$. Wealth is equal to $1 + \pi\tilde{r} - \tilde{\ell}$, where we now interpret $\tilde{\ell}$ as a human capital shock taking values

$$\tilde{\ell} = \begin{cases} 0 & \text{prob. } 1 - p \\ \delta & \text{prob. } p, \end{cases}$$

for unknown probability p. The agent will determine p from memory. Unlike in the previous example, we do not rely on an unknown correlation between $\tilde{r}$ and $\tilde{\ell}$.

Just as in the previous example, the first features vector corresponds to a normal state, whereas the next two correspond to the negative outcomes. We assume that there are also a large number of "neutral" features, which are neither positive nor negative. In particular, the memory matrix

$$\begin{array}{cc} & \begin{array}{ccccc} \text{normal} & \text{fin. crisis} & \text{danger} & \text{other associations} & \end{array} \\ M_{t-1} \propto & \begin{bmatrix} 1-p_1-p_2 & 0 & 0 & 0 & \cdots \\ 0 & p_1 & p_2 & 0 & \cdots \\ 0 & 0 & 0 & \hat{M}_{t-1} & \\ \vdots & \vdots & \vdots & & \end{bmatrix}. \end{array}$$

The second column represents a negative economic event (a financial crisis) and the third represents a danger signal; these observable states share the context of the negative event $\tilde{\ell} = \delta$ occurring. Let $\tilde{q}_t = x_t(2)$ denote the subjective probability of this negative event, and let $\tilde{q}_{t-1} = p_1 + p_2 = p$ be the agent's belief prior to the experiment.

We can represent the viewing of a horror movie as the observation of features similar to, but not exactly the same as, danger: $f_t \approx e_3$. In line with the similarity result discussed above, retrieved context will be similar to what would occur under actual danger. That is, $x_t^{\text{in}} \approx \hat{e}_2$, so $\tilde{q}_t \approx (1-\zeta)\tilde{q}_{t-1} + \zeta$. Therefore, retrieved features equal

$$f_t^{\text{in}} \approx (1-\tilde{q}_t)e_1 + \tilde{q}_t \propto (p_1+p_2)^{-1}(p_1 e_2 + p_2 e_3),$$

so that the subjective probability of danger or depression equals $\tilde{q}_t$.

Wachter and Kahana (2021) assume that the agent chooses π to maximize expected log utility over wealth:

$$\mathbb{E}^{\ell}\left[\frac{1}{2}\log(1+\pi(\mu+\sigma)-\tilde{\ell}) + \frac{1}{2}\log(1+\pi(\mu-\sigma)-\tilde{\ell})\right],$$

where, again, the outcome $\tilde{\ell} = \delta$ occurs with probability $\tilde{q}_t$ and is 0 otherwise. They calibrate the example with an excess return $\mu = 4\%$, a standard deviation $\sigma = 20\%$, a prior probability of the negative labor market outcome $p = 2\%$, and a percent decline $\delta = 0.8$, should the outcome occur. They also assume $\zeta = 0.35$. They find that, when the agent has the correct probabilities, the portfolio allocation equals 70%; after the stimulus, this falls to 30%. 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.

The horror movie changes the beliefs of the agent about the risk the agent might face. It is as if the movie reminds the agent that the world is a risky place, and one thus should not take risks with one's financial wealth. Our setup could either be interpreted as the one in which the agent is reminded of why wealth is necessary (that is the literal interpretation above), or is reminded of how painful (through time-varying risk aversion) low-wealth states are.

The response of the agent to the experiment cannot be Bayesian: a movie has not changed anything about the outside world. In that sense, the response of risk-taking to viewing a horror movie is a good test of our theory. The experiment shows that financial decisions in one context do not resemble financial decisions in another, even though the financial decision in both cases is materially the same. Context "should" be irrelevant, and yet it is not. The agent may know, intellectually, that nothing has changed, and yet the powerful pull of context implies that choices change anyway.

In conclusion, Wachter and Kahana (2021) provides us with the first model of memory in economics where the long-term effects of past experiences (key feature 1) are front and center of the model setup. The notion of decay, in turn, relates to recency bias (key feature 2). In addition, they propose a more general model of context that is not necessarily tied to the physical environment. Instead, their temporal context model allows to directly link to the notion of time, tying context specificity (akin to key feature 3) to the temporal features 1 and 2. Finally, their model is presented as applicable to every human, and in fact repeatedly refers to the decision-making of professionals, consistent with key feature 4.

IV Conclusion

Our review of the empirical and theoretical literature on the influence of past experiences and memory on decision-making has emphasized the urgent need for economic research to incorporate personal histories and the memory thereof into theoretical model of decision-making. We have highlighted four empirical features – long-term effects of past experiences, recency bias, domain specificity, and robustness to learned knowledge – which should guide the development of further theoretical frameworks. In particular, we have emphasized that memory appears to play an essential role for economic choices and behavior in the long-run, and thus beyond the relatively short horizon that characterize much of the laboratory evidence reviewed in other chapters of this Handbook.

We have introduced several sets of models that achieve some or all of these objectives. The existing theoretical approaches span a wide spectrum from simple overlapping generations (OLG) models that directly embeds all features (1) to (4) to more abstract theoretical developments including case-based decision-making, two-state memory models, and models of context and retrieval. While some models, such as Wachter and Kahana (2021) are able to capture all of the empirical key features, many concepts and details remain to be ironed out. For example, there is the question of what exactly is and is not a "context" or "feature" or, as we put it in describing key feature 3, a "domain." Further work is also needed in conceptualizing emotions, such as fear and joy, and understanding their role in anchoring experiences in memory as well as the retrieval process.

References

Akerlof, G. A. (2019). What They Were Thinking Then: The Consequences for Macroeconomics during the Past 60 Years. *Journal of Economic Perspectives*, 33(4):171–186.

Arrow, K. J. (1970). *Essays in the theory of risk-bearing*. North Holland.

Barberis, N. (2018). *Chapter 2 - Psychology-Based Models of Asset Prices and Trading Volume*, volume 1. Elsevier.

Bergado, J. A., Lucas, M., and Richter-Levin, G. (2011). Emotional tagging - a simple hypothesis in a complex reality. *Progress in Neurobiology*, 94:64–76.

Berger, J. O. (1985). *Statistical decision theory and Bayesian analysis*. Springer-Verlag, New York, 2nd edition.

Bernanke, B. S. (2015). *The Courage to Act*. W. W. Norton & Company, New York.

Bordalo, P., Gennaioli, N., and Shleifer, A. (2018). Diagnostic expectations and credit cycles. *The Journal of Finance*, 73(1):199–227.

Bordalo, P., Gennaioli, N., and Shleifer, A. (2019a). Memory and reference prices: An application to rental choice. *AEA Papers and Proceedings*, 109:572–76.

Bordalo, P., Gennaioli, N., and Shleifer, A. (2019b). Memory, attention, and choice. NBER Working Paper No. 23256.

Botsch, M. and Malmendier, U. (2019). The long shadows of the great inflation: Evidence from residential mortgages. Working paper, University of California, Berkeley.

Bromiley, P. and Curley, S. (1992). Individual differences in risk taking. In Yates, J. F., editor, *Risk-taking behavior*, pages 87–132. John Wiley, New York.

Brunnermeier, M. K. and Parker, J. A. (2005). Optimal expectations. *The American Economic Review*, 95(4):1092–1118.

Busse, M. R., Pope, D. G., Pope, J. C., and Silva-Risso, J. (2015). The psychological effect of weather on car purchase. *The Quarterly Journal of Economics*, 130(1):371.

Carroll, C. D. (2003). Macroeconomic expectations of households and professional forecasters. *the Quarterly Journal of economics*, 118(1):269–298.

Chang, T. Y., Huang, W., and Wang, Y. (2018). Something in the air: Pollution and the demand for health insurance. *The Review of Economic Studies*, 85(3):1609–1634.

Christiano, L. J. and Fitzgerald, T. J. (2003). The band pass filter. *international economic review*, 44(2):435–465.

Cochrane, J. H. (2018). Michelson-morley, fisher, and occam: The radical implications of stable quiet inflation at the zero bound. *NBER Macroeconomics Annual*, 32:113–226.

Cohn, A., Engelmann, J., Fehr, E., and Maréchal, M. A. (2015). Evidence for countercyclical risk aversion: An experiment with financial professionals. *American Economic Review*, 105(2):860–85.

Collin-Dufresne, P., Johannes, M., and Lochstoer, L. A. (2017). Asset pricing when 'this time is different'. *The Review of Financial Studies*, 30(2):505–535.

Conlin, M., O'Donoghue, T., and Vogelsang, T. J. (2007). Projection bias in catalog orders. *American Economic Review*, 97(4):1217–1249.

Curcuru, S., Heaton, J., Lucas, D., and Moore, D. (2009). Heterogeneity and portfolio choice: Theory and evidence. In Ait-Sahalia, Y. and Hansen, L. P., editors, *Handbook of Financial Econometrics*.

Deaton, A. (1992). *Understanding consumption.* Oxford: Oxford University Press.

Dow, J. (1991). Search decisions with limited memory. *The Review of Economic Studies*, 58(1):1–14.

Ericson, K. M. and Laibson, D. (2019). *Chapter 1 - Intertemporal Choice*, volume 2. Elsevier.

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

Fama, E. F. and French, K. R. (1988). Dividend yields and expected stock returns. *Journal of financial economics*, 22(1):3–25.

Frederick, S., Loewenstein, G., and O'Donoghue, T. (2002). Time discounting and time preference: A critical review. *Journal of Economic Literature*, 40:351–401.

Friedman, M. and Schwartz, A. J. (1963). *A monetary history of the US, 1867-1960.* Princeton University Press.

Fuster, A., Hebert, B., and Laibson, D. (2012). Natural expectations, macroeconomic dynamics, and asset pricing. *NBER Macroeconomics Annual*, 26(1):1–48.

Fuster, A., Laibson, D., and Mendel, B. (2010). Natural expectations and macroeconomic fluctuations. *Journal of Economic Perspectives*, 24(4):67–84.

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

Glaeser, E. L. and Nathanson, C. G. (2017). An extrapolative model of house price dynamics. *Journal of Financial Economics*, 126(1):147–170.

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.

Goldstein, W. M. and Weber, E. U. (1995). Content and discontent: Indications and implications of domain specificity in preferential decision making. In *Psychology of Learning and Motivation*, volume 32, pages 83–136. Elsevier.

Greene, R. L. (1986). Sources of recency effects in free recall. *Psychological Bulletin*, 99(2):221.

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

Hall, R. E. and Sargent, T. J. (2018). Short-run and long-run effects of milton friedman's presidential address. *The Journal of Economic Perspectives*, 32(1):121–134. Copyright - Copyright American Economic Association Winter 2018; People - Friedman, Milton; Last updated - 2018-01-31; SubjectsTermNotLitGenreText - Friedman, Milton.

Hogarth, R. M. and Einhorn, H. J. (1992). Order effects in belief updating: The belief-adjustment model. *Cognitive psychology*, 24(1):1–55.

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

Hume, D. (1748). *An Enquiry Concerning Human Understanding.* Oxford: Clarendon Press, 2nd, 1966 edition.

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

Kahana, M. J. and Wagner, A. D., editors (2019). *This is a placeholder for Mack and Palmeri's handbook chapter*, chapter 5.1. Oxford University Press.

Kern, H. L. and Hainmueller, J. (2009). Opium for the masses: How foreign media can stabilize authoritarian regimes. *Political Analysis*, 17:377–399.

Laudenbach, C., Malmendier, U., and Niessen-Ruenzi, A. (2019a). Emotional tagging and belief formation: The long-lasting effects of experiencing communism. In *AEA Papers and Proceedings*, volume 109, pages 567–71.

Laudenbach, C., Malmendier, U., and Niessen-Ruenzi, A. (2019b). The long-lasting effects of experiencing communism on attitudes towards financial markets.

LeRoy, S. F. (2005). Excess volatility. *working paper*.

Li, G. (2014). Information sharing and stock market participation: Evidence from extended families. *The Review of Economics and Statistics*, 96(1):151–160.

Low, H., Meghir, C., and Pistaferri, L. (2010). Wage risk and employment risk over the life cycle. *American Economic Review*, 100(4):1432–67.

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

Luchins, A. S. (1958). Definitiveness of impression and primacy-recency in communications. *The Journal of Social Psychology*, 48(2):275–290.

MacCrimmon, K. R. and Wehrung, D. A. (1990). Characteristics of risk taking executives. *Management science*, 36(4):422–435.

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. (2021). The making of hawks and doves: Inflation experiences on the fomc. Technical report.

Malmendier, U., Pouzo, D., and Vanasco, V. (2020). Investor experiences and financial market dynamics.

Malmendier, U. and Shen, L. S. (2018). Scarred consumption. Technical report, National Bureau of Economic Research.

Malmendier, U. and Steiny, A. (2017). Rent or buy? the role of lifetime experiences of macroeconomic shocks within and across countries.

Mankiw, N. G. and Reis, R. (2002). Sticky information versus sticky prices: a proposal to replace the new keynesian phillips curve. *The Quarterly Journal of Economics*, 117(4):1295–1328.

Mankiw, N. G. and Romer, D., editors (1991). *New Keynesian Economics*, volume 1. MIT Press.

Matter, M. G., Talmi, D., and Daw, N. D. (2019). Memory mechanisms predict sampling biases in sequential decision tasks. Working paper, Princeton Neuroscience Institute.

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.

Milani, F. (2007). Expectations, learning and macroeconomic persistence. *Journal of monetary Economics*, 54(7):2065–2082.

Modigliani, F. (1977). The monetarist controversy or, should we forsake stabilization policies? *The American Economic Review*, 67(2):1–19.

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

Müller, O., Pollack, D., and Pickel, G. (2013). *Religiös-konfessionelle Kultur und individuelle Religiosität: Ein Vergleich zwischen West- und Ostdeutschland*. Springer Verlag.

Murdock, B. B. (1967). Theoretical note: A fixed-point model for short-term memory. *Journal of Mathematical Psychology*, 4(3):501–506.

Nagel, S. and Xu, Z. (2019). Asset pricing with fading memory. Working Paper 26255, National Bureau of Economic Research.

Nakamura, E. and Steinsson, J. (2008). Five facts about prices: a reevaluation of menu cost models. *The Quarterly Journal of Economics*, 123(4):1415–1464.

O'Donoghue, T. and Sprenger, C. (2018). *Chapter 1 - Reference-Dependent Preferences*, volume 1. Elsevier.

Orphanides, A. and Williams, J. C. (2005). The decline of activist stabilization policy: Natural rate misperceptions, learning, and expectations. *Journal of Economic dynamics and control*, 29(11):1927–1950.

Petschow, U., Meyerhoff, J., and Thomasberger, C. (1990). *Umweltreport DDR. Bilanz der Zerstörung. Kosten der Sanierung. Strategien für den Ökologischen Umbau.* Frankfurt am Main.

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.

Restle, F. (1965). Significance of all-or-none learning. *Psychological Bulletin*, 64:313–325.

Richter-Levin, G. and Akirav, I. (2003). Emotional tagging of memory formation - in the search for neural mechanisms. *Brain Research Reviews*, 43:247–256.

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

Schraeder, S. (2015). Information processing and non-bayesian learning in financial markets. *Review of Finance*, pages 1–31.

Shiller, R. J. (1981). Do stock prices move too much to be justified by subsequent changes in dividends? *American Economic Review*, 71:421–436.

Simonsohn, U. and Loewenstein, G. (2006). Mistake #37: The effect of previously encountered prices on current housing demand. *The Economic Journal*, 116(508):175–199.

Simonson, I. and Tversky, A. (1992). Choice in context: Tradeoff contrast and extremeness aversion. *Journal of Marketing Research*, 29(3):281–295.

Sims, C. A., Friedman, B. M., and Woodford, M. (2010). *Chapter 4 - Rational Inattention and Monetary Economics*, volume 3, pages 155–181. Elsevier.

Slovic, P., Fischhoff, B., and Lichtenstein, S. (1986). The psychometric study of risk perception. In *Risk evaluation and management*, pages 3–24. Springer.

Tversky, A. and Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. *Cognitive psychology*, 5(2):207–232.

Tversky, A. and Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 4157:1124–1131.

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

Vissing-Jorgensen, A. (1998). *Limited stock market participation.* PhD thesis, Massachusetts Institute of Technology.

Wachter, J. A. and Kahana, M. J. (2021). A retrieved-context theory of financial decisions. Working Paper 26200, National Bureau of Economic Research.

Weber, E. U., Blais, A.-R., and Betz, N. (2002). A domain-specific risk-attitude scale: Measuring risk perceptions and risk behaviors. *Journal of Behavioral Decision Making*, 15:263–290.

Woodford, M. (2013). Macroeconomic analysis without the rational expectations hypothesis. *Annual Review of Economics*, 5(1):303–346.

Woodford, M. (2019). Modeling imprecision in perception, valuation and choice. NBER Working Paper No. 26258.