

Experience Effects in Finance: Foundations, Applications, and Future Directions*

Ulrike Malmendier

University of California Berkeley

Abstract

This article establishes four key findings of the growing literature on experience effects in finance: (i) the long-lasting imprint of past experiences on beliefs and risk taking; (ii) recency effects; (iii) the domain-specificity of experience effects; and (iv) imperviousness to information that is not experience-based. I first discuss the neuroscientific foundations of experience-based learning and sketch a simple model of its role in the stock market based on Malmendier et al. (2020a, b). I then distill the empirical findings on experience effects in stock-market investment, trade dynamics, and international capital flows, highlighting these four key features. Finally, I contrast models of belief formation that rely on “learned information” with models accounting for the neuroscience evidence on synaptic tagging and memory formation, and provide directions for future research.

JEL classification: D14, D81, D83, D87, D91, F30, G11, G12, G41, G50

Keywords: Experience Effects, Beliefs, Recency, Domain specificity, Information, Stock-market participation, Trade dynamics, International capital flows

1. Introduction

Economic and political crises often shape entire generations. For example, World War I brought about the “lost generation,” whose beliefs and values were shaken by the disorienting war experience. The Great Depression bore a generation of risk-averse Depression Babies, who shied away from the stock market and hoarded food supplies. Positive experiences can also be formative, such as the post-World War II prosperity and optimism giving rise to the generation of “Baby Boomers.”

In recent years, financial economists have started to better understand the long-lasting imprint resulting from the exposure to different economic and political outcomes, and

* This article is based on the 2020 European Finance Association keynote address titled “Exposure, Experience, and Expertise: Why Personal Histories Matter in Finance and Economics.” I thank the editor Alex Edmans, Chris Parsons, and the EFA audience for their comments and suggestions, and Clint Hamilton, Karin Li, and Junru Lyu for excellent research assistance.

especially macro-financial shocks. Personal lifetime experiences appear to deeply shape belief formation and risk taking in finance, whether we consider stock-market participation, mortgage choices, consumption expenditures, or interest-rate expectations. While an earlier literature in corporate finance had identified the role of formative experiences for managerial decision-making (cf. [Malmendier, 2018](#)), there is now a growing understanding of the much broader and more general applicability of these concepts.

In this article, I shed light on the role of past lifetime experiences (experience effects) from three angles. First, I discuss the neuroscience foundations of experience effects and point out how they do or do not square with existing models of belief formation. While some of these links to neuroscience research are still tentative, they help conceptualize the notion of experience effects. Second, I summarize a simple theoretical Overlapping-Generations (OLG) framework, which captures the notion of experience-based learning (EBL) and indicates possibilities of further theoretical development. And third, I present some of the main empirical findings on beliefs and investment in the stock market, with an eye toward synthesizing the stylized features and identifying areas for further research.¹

To guide the discussion, let me preview here four “key features” that have emerged from the growing empirical literature and which, as I will argue, should guide the development of modern models of belief formation in macro-finance.

1. Long-lasting effects: Past experiences of macro-finance realizations shape individual beliefs and choices for years and decades to come. For example, individuals who have lived through stock-market crises are less likely to participate in the stock market for the rest of their lives, both on the extensive and the intensive margin ([Malmendier and Nagel, 2011](#)). Individuals who have lived through the financial and housing crisis of 2008 have remained significantly less likely to purchase a home, and are also significantly curbing their expenditures ([Malmendier and Shen, 2018](#)).
2. Recency bias: More recent experiences have a stronger impact on individual expectations and risk-taking than experiences made earlier in life ([Tversky and Kahneman, 1974](#)). When predicting future returns in the stock, housing, or other asset markets, investors overly rely on the price realizations in those markets over recent months or the past year, though it is also the case that big enough shocks have a detectable impact on individual investors and consumers decades later.
3. Domain specificity: The lingering influence of past experiences on risk taking is specific to the markets or arenas of life in which individuals have had personal exposure to past realizations. For example, while exposure to negative returns in the stock market reduces stock-market participation, and negative exposure in the bond market reduces bond-market participation, there is no evidence of cross-fertilization ([Malmendier and Nagel, 2011](#)). In other words, rather than altering risk attitudes in general, these adjustments concern the concrete “domain,” in which a good or bad outcome has been experienced.²

1 There is a wide range of applications of experience effects outside the realm of the stock market and, in fact, outside finance, which I will briefly point to in Section 6; cf. also [Malmendier \(2021\)](#).

2 In the context of experience-based learning, “domain specificity” indicates that beliefs about future realizations of similar assets or correlated risks are not necessarily following this correlation structure. The notion of “domains” is similar to that of “(content) domains” in [Weber et al. \(1993\)](#) and [Weber, Blais, and Betz \(2002\)](#), though their definition of is broader. I will discuss the ongoing debate about domains and domain-specificity in theoretical neuroscience in Section 2.

4. Robustness to expert knowledge: While much of the existing research in behavioral finance and economics emphasizes biases detected in consumers and individual investors, but not among professional decision-makers, experience effects are observed among highly educated and specialized individuals as well, even in their area of expertise. For example, central bankers at the FOMC tilt their inflation predictions into the direction of their personal experiences (Malmendier, Nagel, and Yan, 2021), and professional forecasters place too much weight on their own beliefs and too little on publicly available information (Bianchi, Ludvigson, and Ma, 2020).

In Section 2, I link these four features to the conceptual foundation of experience effects in the neuroscience and psychology literature. The observation that individuals assign excess weights to outcomes they have personally experienced mirrors the underlying neurological process of synapse formation and reflects our modern understanding of neuroplasticity, that is, of the brain's lifelong ability to change and adapt as a result of experience. The brain forms stronger connections between neurons that are used more frequently, while those that have not been used in a while eventually die, providing the foundations of the first two empirical features: (i) the statistically detectable influences of past experiences and (ii) recency bias. The neuroscience evidence also microfounds (iii) domain specificity since the weights are specific to the outcomes (type of stimuli) the individual has experienced. Finally, the process of "re-wiring" naturally applies to the neural pathways of experts and non-experts alike, as suggested under (iv) above.

The experience-based updating process can be modeled as a form of generalized Bayesian learning (Bissiri, Holmes, and Walker, 2016) with a loss function such that individuals assign more weight to outcomes that they have personally experienced than to other known realizations. In Section 3, I briefly sketch the simple theoretical OLG framework from Malmendier, Pouzo, and Vanasco (2020a, 2020b) to guide the discussion of the empirical evidence. After presenting some of the empirical evidence, with a focus on stock-market experiences and beliefs about stock returns, in Section 4, I point to the similarities and differences of other modern approaches to belief formation in Section 5. The key difference of EBL is the recasting of learning from information-based to experience-based updating. Information can be learned, but learned knowledge will not undo the "re-wiring" that occurs in our brains as we are accumulating experiences. Section 6 provides additional directions for future research.

2. Neuroscience Foundation

Traditional models of economic decision-making do not allow for personally experienced information to affect individual beliefs differently than otherwise acquired information about the same realizations, at least once we account for all economic implications of those realizations. For example, the effect of "living through a depression" on individuals' economic outlook, financial risk-taking, and financial investment is predicted to be no different than reading about it as long as we control for income, wealth, and other economic effects of the depression. Similarly, past experiences of unemployment should have the same influence on future job choices, savings rates, or consumption as full information about the economic implications of such realizations (e.g., about job prospects or potential industry shocks), given controls. And, living through the COVID-19 pandemic would have

the same impact on educational, social, or economic choices as perfect information about the likelihood and implications of a pandemic would.

In practice, it is of course often hard to perfectly control for all direct and indirect implications of economic shocks that might have changed the economic reality of those who have experienced it. But the thought experiment challenges the view that the long-term effects of crisis experiences can *only* be attributed to changed circumstances. Intuitively, even if we could perfectly control for all such channels, would not we expect marked differences in behavior after such an experience?

Modern neuroscience suggests exactly that. The human brain alters its structure and adapts to new experiences throughout life. At every new experience, it forms a connection (synapse) between two neurons through which the neurons communicate how to react in response to the experience. The pre-synaptic cell sends neurochemical messengers (neurotransmitters) to the post-synaptic neuron via the synapse, and the post-synaptic neuron reacts, for example, by sending its own neurotransmitter or by reducing or even shutting down its own signaling (Bear, Connors, and Paradiso, 2020).

Importantly, for our dynamic learning context, the brain keeps reorganizing synaptic pathways throughout life as we accumulate new experiences. While researchers used to believe that the brain became fixed after a certain age, modern neuroscience has established that the brain never stops changing in response to learning. Neurons that are used frequently develop stronger connections, while those that are used infrequently will eventually die. This process of synaptic pruning also explains why, after an explosive growth in synapses since birth (from an estimated 2,500 synapses per neuron in the cerebral cortex at birth to maybe 15,000 synapses per neuron at the age of three), the number of synapses is about halved when we reach adulthood (cf. Chapter 2 in Doidge, 2007). In other words, the brain alters its actual physical structure as the result of experiences (structural brain plasticity).³ EBL directly captures this type of rewiring: past lifetime experiences are predicted to have a lasting impact on people's perceptions and decision-making.

The empirical evidence on experience effects also suggests that it matters how strongly and how often we make an experience, as we will discuss in detail in Section 4. The neuroscientific evidence mirrors this insight, too. The strength and shape of synapse formation depends on how and how often we make an experience. Bliss and Lomo (1973) first discovered that the structure of a synapse changes when it is repeatedly stimulated. As illustrated in Figure 1, the post-synaptic neuron ends up with more available receptors, which heightens its sensibility. A prolonged increase in synaptic strength is known as long-term potentiation (LTP), and has become the primary cellular model of memory in the mammalian brain (Bear, Connors, and Paradiso, 2020). Frey and Morris (1997) later proposed that "synaptic tags" at the potentiated synapses allow to establish particularly longlasting late LTP. Their synaptic tagging hypothesis states that synapses get "tagged" by some previous synaptic activity. That is, prior activity of a neuron is relevant for the longer-lasting persistence of LTP—which is the phenomenon we are interested in to better understand long-term scarring effects.

We can now see how the neuroscientific evidence on structural brain plasticity and the formation of synapses provides a foundation for our four key features of EBL: Both (i) persistent influence of past experiences and (ii) recency bias reflect the process of forming,

3 In addition to structural plasticity, the brain also possesses the functional plasticity to move functions from a damaged area of the brain to other undamaged areas.

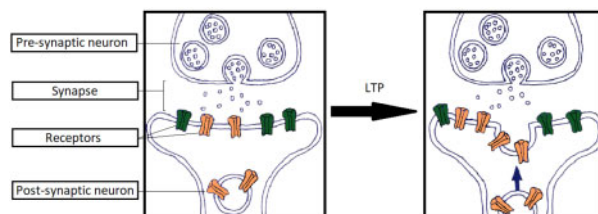


Figure 1. Expression of LTP. Notes: Postsynaptic changes are visible in the increased number of receptors in post-synaptic (lower) neuron. Rendition replicates [Whitlock et al. \(2006\)](#).

strengthening, and removing neural connections. As these processes are specific to an experience, and do not alter the generalized concept of “risky decisions” or “risky assets,” they are naturally (iii) domain-specific. And, as the “re-wiring” of neural pathways and “tagging” of synapses occur in the mammal brains of experts and non-experts alike, we have a foundation of the (iv) robustness to learned knowledge.

At this point, it is worth elaborating on the notion “domain” and “domain-specific.” As used here, domain-specific learning means that experiences in one type of market or with one type of asset risk do not necessarily affect beliefs about another market or asset risk, even if the underlying payoff structures are correlated and the individual knows this. The implied lack of cross-fertilization is similar to that across “(content) domains” in [Weber et al. \(1993\)](#); [Weber, Blais, and Betz \(2002\)](#), though their definition is broader, as they distinguish financial decisions (separately for investing versus gambling), health and safety, recreational, ethical, and social decisions. The theoretical neuroscience literature, instead, utilizes the term “(cognitive) domain” to describe separate brain modules that are specialized for particular types of stimuli. Initially, the observation of the specific neuropsychological deficits in brain-damaged patients lead to the conclusion that different cognitive domains are processed within a specialized region of the brain ([Karmiloff-Smith, 2015](#)), and lead to the distinction of common sense psychological categories such as, say, number, face, and grammatical (semantics and syntax) processing, spatial cognition, and more abstract modules such as “knowledge of the constraints governing the physical world.” Modern brain data, however, suggests that domains need to be defined more abstractly than the common sense psychological categories listed above ([Spunt and Adolphs, 2017](#)).

The same applies to the notion in the financial decision-making context. It is an open question which stimuli are to be grouped together in “domains,” and an even bigger question how to translate these groupings into typical financial variables. The evolving insights from neuroscience will help economists shape their theoretical modeling toward neuropsychological realism as well as improved predictions, both in the time series and in the cross-section. The theory and evidence presented in this article is the first attempt to reflect the features we have gathered so far from the neuroscientific evidence.

I conclude this section with another word of caution. The type of neuroscientific evidence presented in here describes processes on a very disaggregate (molecular) level. Tying such evidence to human beliefs and choice behavior is necessarily daring. After all, current technology does not enable us to pinpoint, at a molecular level, the exact impact of a given experience on the brain or link specific changes in the brain to specific actions. So why not stay at the level of behavioral psychological evidence, for example, refer to the evidence on “availability bias” ([Tversky and Kahneman, 1974](#)) and its successors? The latter is indeed

fully consistent with the underpinnings proposed here: people tend to overweight events that come to mind easily, that is, are “available” to them, and personal experiences are a catalyst for this availability. Hence, the discussion of experience effects from a cognitive neuroscience and a psychology perspective lead to similar insights.

The reason to “dig deeper” is again the open question of which circumstances and features of an experience that helps to anchor it more strongly in memory and makes it more easily retrievable. The hope is that the developing neuroscience research will help us understand better what comes to mind easily, why, and under what circumstances.

3. Theoretical Model

In this section, I sketch the basic ingredients of the simple OLG model of EBL from [Malmendier, Pouzo, and Vanasco \(2020a, 2020b\)](#) to illustrate how the (over-)emphasis on personally experienced realization generates the four features highlighted above, (i) long-lasting scars, (ii) recency bias, (iii) domain specificity, and (iv) robustness to information (applicability to experts). The brief sketch also indicates how such a framework gives rise to additional predictions that can be tested empirically.

Beliefs versus Preferences

Before we can get started, we need to consider the oft-asked “beliefs-versus-preferences” question. It is *ex ante* unclear whether we would want to capture the re-wiring of our brains as the formation (and alteration) of individual preferences or as individuals assigning altered probability weights to possible future outcomes. Or, we might consider another route altogether, akin to cognitive limitations discussed in the behavioral economics literature ([Rabin, 1998](#)). What is the right approach?

There are several possible responses to this question. The first answer is that we can go either way and that the distinction between a preference and a beliefs channel matters less than our standard economic modeling approaches may lead us to think. As [Savage \(1954\)](#) has shown, we can map between subjective beliefs and preferences, and even derive beliefs from preferences. In fact, attempts to distinguish between beliefs and preferences in determining equilibrium prices can be futile.⁴

Second, though, almost every piece of empirical research on experience effects includes direct evidence on past experiences affecting beliefs—whether it is the beliefs of investors about future stock-market performance, the beliefs of consumers and FOMC members about future inflation rates, or the beliefs of mortgage borrowers about future interest rates.⁵ This evidence does not rule out that personal experiences also alter individual preferences and cognitive processes, but it invites belief-based modeling as a (well-documented) starting point, from which we can move to incorporate other channels as evidence emerges.

Finally, regarding the idea of a third, cognitive channel, note that a basic belief-based approach of experience effects could be re-interpreted as a model of cognitive limitation. In

4 As [Kraus and Sick \(1980\)](#) put it: “Since individual agent optimality conditions involve the product of probability and marginal utility, it may be that any set of equilibrium prices that are consistent with some combination of beliefs and preferences could also have resulted from different beliefs combined with different preferences.”

5 [Malmendier and Nagel \(2011, 2016\)](#); [Malmendier, Nagel, and Yan \(2021\)](#); [Botsch and Malmendier \(2020\)](#).

a model of generalized Bayesian learning, where agents overweigh or restrict their data to realizations that they have personally experienced, this limited consideration of data can be re-interpreted as a cognitive restriction. The simple theoretical model sketched below illustrates this point.

Model Set-Up

We start from briefly sketching the infinite-horizon OLG economy from [Malmendier, Pouzo, and Vanasco \(2020a, 2020b\)](#).⁶ There is a continuum of agents with CARA preferences. At each point in time $t \in \mathbb{Z}$, a new generation is born and lives for $q \in \{1, 2, 3, \dots\}$ periods. Each generation has a mass of q^{-1} identical agents, and there are $q + 1$ generations alive at any time t . [Figure 2](#), adapted from [Malmendier, Pouzo, and Vanasco \(2020b\)](#), illustrates the OLG setting for the case of $q = 2$, meaning that each agent born t lives for two periods, until $t + 2$, and trades at three points in time, $t, t + 1$, and $t + 2$. At each point in time (on the horizontal timeline), there are three cohorts alive, with their names indicated on the vertical axis. As time progresses, new cohorts enter, and cohorts that have already lived for two periods exit. For example, at time $t + 1$, the $(t + 1)$ -cohort is born and lives for the next two periods, while the $(t - 1)$ -cohort consumes everything they own at time $t + 1$ and then exits.

Agents transfer resources across time by investing in financial markets. There are two types of assets: the risk-free asset (in perfectly elastic supply) generates a gross return of $R > 1$ at all times, and the risky asset (in unit net supply) generates a dividend $d_t \sim N(\theta, \sigma^2)$ at time t . Trading takes place at the beginning of each period. At the end of the last period of their lives, agents consume the wealth they have accumulated.

The key question is how agents learn about future returns when they do not know the true mean of dividends θ and use past observations to estimate it. (For tractability, agents are assumed to know the variance of dividends.) For the purpose of this article, I will not consider the general model from [Malmendier, Pouzo, and Vanasco \(2020b\)](#) but only the more tractable version with myopic agents, who maximize per-period utility. I also focus on the case of only one young and one older generation actively trading at any point in time ($q = 2$).

A two-period lived generation born at $t = n$ maximizes $E_t^n[-\exp(-\gamma W_{t+1}^n)]$ at each time $t \in \{n, n + 1\}$, where W_{t+1}^n is n 's wealth at time $t + 1$. Denoting with x_t^n the investment in the risky asset, with a_t^n the amount invested in the riskless asset, and with p_t the price of one unit of the risky asset, generation n 's budget constraint is $W_t^n = x_t^n p_t + a_t^n$ at $t \in \{n, n + 1\}$. Wealth next period becomes $W_{t+1}^n = x_t^n (p_{t+1} + d_{t+1}) + a_t^n R = x_t^n (p_{t+1} + d_{t+1} - p_t R) + W_t^n R$, or, if we express the excess payoff from investing in one unit of the risky asset over the riskless asset as $s_{t+1} \equiv p_{t+1} + d_{t+1} - p_t R$ (analogous to the equity premium), $W_{t+1}^n = x_t^n s_{t+1} + W_t^n R$. Hence, the maximization problem can be rewritten as

$$x_t^n \in \arg \max_{x \in \mathbb{R}} E_t^n[-\exp(-\gamma x s_{t+1})], \tag{1}$$

where $E_t^n[\cdot]$ is the (subjective) expectation with respect to a Gaussian distribution with variance σ^2 and a subjective mean of dividends, denoted by θ_t^n .

6 Other approaches to modeling experience-based learning include [Ehling, Graniero, and Heyerdahl-Larsen \(2018\)](#) and [Collin-Dufresne, Johannes, and Lochstoer \(2017\)](#), who also discuss portfolio choice and asset pricing implications of experience-based learning. Reassuringly, all papers reach similar conclusions coming from different modeling angles.

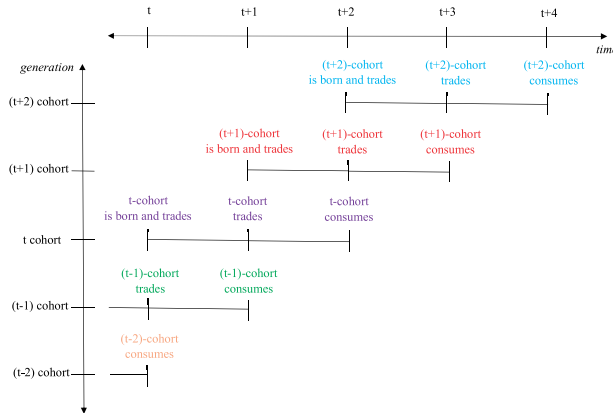


Figure 2. Timeline for an economy with two-period lived generations, $q = 2$.

Experience-Based Learning (EBL)

Under EBL agents (i) overweigh realizations observed during their lifetimes and (ii) tilt the excess weights toward the most recent observations when forecasting dividends. It is important to emphasize that EBL agents choose to discount earlier observations even though they observe the entire history of dividends. Learned knowledge does not alter the way the brain processes experienced realizations. For simplicity, agents are also modeled as fully understanding the model and knowing all the primitives except the mean of the dividend process, differently from, say, agents in reinforcement learning-type models.

Let us consider the limiting case where agents only use observations realized during their lifetimes and apply zero weight to any prior realizations. In this case, the subjective mean of dividends of generation n at time t is

$$\theta_t^n \equiv \sum_{k=0}^{\text{age}} w(k, \lambda, \text{age}) d_{t-k}, \tag{2}$$

where $\text{age} = t - n$. Prior literature parameterizes the weight $w(k, \lambda, \text{age})$ that an agent aged age assigns to dividends k periods earlier as

$$w(k, \lambda, \text{age}) = \frac{(\text{age} + 1 - k)^\lambda}{\sum_{k'=0}^{\text{age}} (\text{age} + 1 - k')^\lambda} \tag{3}$$

for all $k \leq \text{age}$, and $w(k, \lambda, \text{age}) \equiv 0$ for all $k > \text{age}$.⁷ The denominator in (3) is a normalizing constant that depends only on age and on the parameter that regulates the recency bias, λ . The empirical evidence suggests a choice of $\lambda > 0$ so that more recent observations receive relatively more weight.

7 Note that the +1 term in both the numerator and denominator serves to avoid zero weight for individuals aged 0, which are included in this simple theoretical set-up. Empirical studies typically include only older individuals, for example, aged 25–75 in Malmendier and Nagel (2011, 2016), and drop the +1 term from (3). The differences in empirical estimates are virtually undetectable.

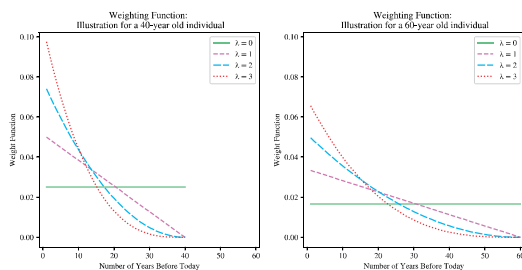


Figure 3. Weighting function: illustration for 40- and 60-year-old individuals.

Figure 3 illustrates the shapes of the weighting function for positive λ values 1, 2, and 3, in line with the range of empirical estimates from Malmendier and Nagel (2011, 2016), and for comparison the case of equal-weighting ($\lambda = 0$). The weights on the left are for a 40-year old investor, and the weights on the right are for a 60-year-old investor. The comparison shows that differently aged investors apply experience weights to different ranges of past years, namely, their personal prior lifetimes. The juxtaposition also highlights that younger generations generally put more weight on recent experiences: Toward the left of the two plots (recent past), the weights are shifted up more for the 40-year-old than the 60-year-old investor. In other words, individuals with a shorter lifespan so far are automatically more “impressed,” or “imprinted,” by recent macro-finance realizations.

Let’s consider a concrete numerical example to see how the mechanics work, using $\lambda = 1$ and the three-generation ($q = 2$) setting from Figure 2. At any time t , there are three generations alive: the oldest cohort ($n = t - 2$), about to exit the market; the middle-aged cohort ($n = t - 1$), and the youngest cohort ($n = t$) that just enters the market. Assume that the dividend has been steady with $d_{t-2} = d_{t-1} = 6$, but there is a spike in dividend for the current period, $d_t = 12$. From (2) and (3), the subjective mean of dividends for the youngest generation ($n = t$) is completely dependent on the most recent dividend with $\theta_t^t = \frac{(0+1-0)^1}{(0+1-0)^1} \cdot 12 = 12$. Meanwhile, d_{t-1} plays a critical role in the subjective mean of dividends for the middle-age generation ($n = t - 1$), with $\theta_t^{t-1} = \sum_{k=0}^1 w(k, 1, 1)$ $d_{t-k} = \frac{(1+1-0)^1}{(1+1-0)^1 + (1+1-1)^1} \cdot 12 + \frac{(1+1-1)^1}{(1+1-0)^1 + (1+1-1)^1} \cdot 6 = 10$. Due to recency effects, this cohort places heavier weights on the more recent dividend d_t , but they still place some, albeit lower weight on the earlier dividend d_{t-1} as well, making $\theta_t^{t-1} < \theta_t^t$. The oldest cohort ($n = t - 2$), instead, exits the market by selling everything they own without considering the dividend. In order to demonstrate the mechanics of the weighting functions, we can still calculate what θ_t^{t-2} would base their trading on if this cohort were to actively trade, namely, $\theta_t^{t-2} = \sum_{k=0}^2 w(k, 1, 2) d_{t-k} = \frac{(2+1-0)^1}{(2+1-0)^1 + (2+1-1)^1 + (2+1-2)^1} \cdot 12 + \frac{(2+1-1)^1}{(2+1-0)^1 + (2+1-1)^1 + (2+1-2)^1} \cdot 6 + \frac{(2+1-2)^1}{(2+1-0)^1 + (2+1-1)^1 + (2+1-2)^1} \cdot 6 = 9$. We see that, although both cohort $n = t - 1$ and cohort $n = t - 2$ have only experienced a dividend of six prior to t , the older cohort places heavier emphasis on prior dividends, which causes $\theta_{t-2}^t < \theta_{t-1}^t$.

An important implication of this theoretical set-up is that θ_t^t does not necessarily converge to the truth as $t \rightarrow \infty$. By construction, $\theta_t^t \sim N(\theta, \sigma^2 \sum_{k=0}^{\text{age}} (w(k, \lambda, \text{age}))^2)$, and

hence, convergence depends on whether $\sum_{k=0}^{\text{age}} (w(k, \lambda, \text{age}))^2 \rightarrow 0$. This in turn depends on how fast the weights on “old” observations decay to zero (i.e., how large λ is). When agents have finite lives, convergence will not occur.

To put it differently, EBL generates persistent belief heterogeneity in investors, which other literatures often have to assume in an ad-hoc manner, for example, to generate trading volume. Experience effects provide an empirical reason why two people may come to hold different beliefs, despite no differences in information.

Equilibrium Outcomes

Malmendier, Pouzo, and Vanasco (2020b) consider the case of linear equilibria where, given a price schedule $\{p_t\}$ that is affine in dividends,

1. the demand profiles $\{(a_t^n, x_t^n) : t \in \{n, n+1\}\}$ for the riskless and the risky asset solve the generation- n problem, and
2. the market clears in all periods: $\frac{1}{2}(x_t^t + x_t^{t-1}) = 1$ for all $t \in \mathbb{Z}$.

In this setting, the risky-asset demand of generation $n \in \{t, t-1\}$ is $x_t^n = \frac{E_t^n[s_{t+1}]}{\gamma V_t^n[s_{t+1}]}$. Most interestingly, one can show that the price loadings depend on past dividends, but only those observed by the oldest generation trading in the market—in our case ($q = 2$):

$$p_t = \alpha + \beta_0 d_t + \beta_1 d_{t-1}. \quad (4)$$

Here, the younger generation of market participants has only experienced the dividend d_t and expects dividends to be identical in the next period. The older (“middle-aged”) generation of market participants has more experience and incorporates the previous dividend d_{t-1} in their weighting scheme. As a result, d_t and d_{t-1} determine p_t , with the price being more sensitive to the more recent dividend ($\beta_0 > \beta_1$) for $\lambda > 0$. The specific value of the coefficients depends on the value of λ .

The dependence of prices on past dividends is a key feature of the model. While models of extrapolation and learning also generate history dependence, experience effects allow us to link it intrinsically to the demographic structure of market participants.

One can develop the model set-up further and allow the mass of a generation to deviate from $1/q$ to better tease out the role of demographic composition. For example, if the fraction of young people in the market increases, current dividends will matter more relative to past dividends for the determination of prices. Vice versa, as a large generation such as the “baby-boom generation” becomes old, prices depend less on contemporaneous dividends and more on past dividends. The model also predicts that an increase in the overall population of market participants generates a level increase in prices.

Further Implications

The simple set-up presented so far captures the imprint of past experiences on beliefs and risk-taking in the stock market, and it naturally incorporates history-dependent fluctuation in the level of market valuation. The model also generates further cross-sectional predictions. It implies (i) persistent generational differences in beliefs, based on personal experiences and (ii) generational differences changing over time as personal experiences of different cohorts diverge or converge. Moreover, it predicts (iii) age differences in response to a recent shock: Intuitively, a given crisis experience exerts stronger influence on younger cohorts, for whom the crisis experience constitutes a larger portion of their lifetime

histories so far. As a result, EBL predicts that younger generations react more optimistically than older generations to positive changes in recent dividends and more pessimistically to negative changes. The bigger the shock (change in dividends) is, the larger is the predicted belief heterogeneity $|d_t - d_{t-1}|$. These differences also imply predictions for history-dependent fluctuations in market participation.

Finally, the model allows to show that cross-sectional differences in lifetime experiences and resulting cross-sectional differences in beliefs affect trade volume. Consider the following measure of total trade volume for our two-period lived agents:

$$TV_t \equiv \left(\frac{1}{2} \sum_{n=t-2}^t (x_t^n - x_{t-1}^n)^2 \right)^{1/2} \tag{5}$$

with $x_{t-1}^t = 0$. That is, trade volume is the square root of the weighted sum (squared) of the change in positions of all agents in the economy. If we rewrite (5) as

$$TV_t = \left[\frac{\chi^2}{2} \left(\theta_t^{t-1} - \theta_{t-1}^{t-1} - \frac{1}{2} \sum_{\tilde{n}=t-2}^t (\theta_t^{\tilde{n}} - \theta_{t-1}^{\tilde{n}}) \right)^2 + \frac{1}{2} (x_t^t)^2 + \frac{1}{2} (x_{t-1}^{t-2})^2 \right]^{\frac{1}{2}} \tag{6}$$

with $\chi = \frac{1}{\gamma\sigma^2(1+\beta_0)}$ and $\theta_{t-1}^t = \theta_t^{t-2} = 0$, we can see that changes in the beliefs θ_t^n of the generations trading at time t , relative to their beliefs at time $t-1$ induce trade, and changes in beliefs are in turn driven by shocks to dividends. Specifically, 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.

Before turning to the empirical evidence, let us reconsider one more time the different possible mechanisms for experience effects—beliefs, preferences, and cognitive limitations. Ostensibly, the model chooses the route of belief formation, restricting the updating behavior to respond to personally experienced realizations. As we emphasized, such behavior does not reflect limited information or failure to understand the model. Instead, agents focus on personally experienced outcomes even though they are fully informed about other historical information. Reflecting the motivating neuroscientific evidence on synaptic strength, the model aims to reflect that past experiences will be accessed more frequently and more directly. Whether we understand this approach to be a type of nonstandard (Gaussian) belief formation or a cognitive limitation on the standard (Gaussian) belief formation process is ultimately a question of labeling. We might even consider the emphasis on one type of information (personally experienced) over other types (learned information) a “preference,” though this latter approach seems less natural. The agent is unlikely to “prefer” or “want” to access personally experienced information more, especially if it is less predictive of future outcomes.

4. Empirical Evidence from Capital Markets

The starting point, and foundation of the literature on EBL, is the evidence on the long-lasting effects of past stock-market crashes (and peaks) on stock-market participation and on the resulting inter-generational differences, as modeled in the prior section. In this section, I synthesize some of the empirical evidence on experience effects in capital markets, with an eye toward highlighting the four key features emphasized throughout the article. In addition, the model from Section 3 revealed several dynamic implications for market

composition and trading behavior, some of which have been tested. I will also discuss cross-sectional sources of variation arising from region- or country-specific exposure to different macro-finance histories, which can be added to the model. The EBL perspective helps to explain several well-known international finance puzzles, including home bias in equity investment, fickleness, and retrenchment.

4.1 Precursors in Earlier Finance Research

Before we dive into the empirical evidence on experience effects in capital markets, I would like to briefly point to precursors in earlier literature. Even before [Malmendier and Nagel \(2011\)](#) coined the term “experience effects,” research in financial economics had documented empirical regularities consistent with this notion. On the individual investor side, [Vissing-Jorgensen \(2003\)](#) shows that young retail investors had the highest return expectations during the stock-market boom in the late 1990s, consistent with the implication of EBL that younger cohorts are most sensitive to recent realizations. In a similar vein, [Greenwood and Nagel \(2009\)](#) document that young mutual fund managers chose higher exposure to technology stocks in the late 1990s than older managers, extending the insight of a stronger reaction of younger cohorts to recent stock-market returns to the allocation to stocks. Another key example is [Kaustia and Knüpfer \(2008\)](#) who show that the returns investors earn from their personal investments in initial public stock offerings (IPOs) are positively related to their future IPO subscriptions. [Choi et al. \(2009\)](#) report that employees who have experienced relatively high returns in their 401(k) account subsequently increase their 401(k) savings rates. The last two papers highlight the importance of personal experiences, in particular.

There is also a parallel literature in Corporate Finance. [Malmendier and Tate \(2005\)](#) and [Malmendier, Tate, and Yan \(2011\)](#), show that corporate managers who grew up during the Great Depression (“Depression babies”) shy away from external financing, in particular stock issuances. [Graham and Narasimhan \(2004\)](#) and [Schoar and Zuo \(2017\)](#) argue that corporate managers who were at the helm of their companies during the Great Depression (or other times of economic downturn) continue to shy away from high debt levels and are more conservative in their capital expenditures. All of this research strongly conveyed the lasting impact of past experiences.

4.2 Experience Effects in the Stock Market—Evidence on the Four Key Features

One famous example of longlasting scars from past economic shocks is the generation of “Depression Babies” in the USA, whose experience of significant financial and emotional trauma during the Great Depression supposedly made them exceedingly risk averse. Given the well-known notion of “Depression Babies,” it is somewhat surprising that no researcher in the twentieth century documented the generation’s aversion to stock-market participation—surprising even more so since already the raw data reveals a striking pattern: the stock-market participation of the generation that experienced the 1930s Great Depression as teenagers or young adults is, at 13%, less than half of later cohorts.

[Malmendier and Nagel \(2011\)](#) are the first to systematically test for the relation between past stock-market experiences and stock-market participation rates, using data from the Survey of Consumer Finances (SCF, available since 1983), combined with data from its precursor, the Inter-university Consortium for Political and Social Research (ICPSR, available since 1947). What is crucial for an empirical analysis of experience effects in field data is to

ensure that, as much as possible, the estimates do not pick up life-cycle effects, such as wealth-dependent, age-dependent, or time-varying risk aversion, nor time trends or any aggregate effects. To achieve this goal the authors focus on the cross-section, and differences in the cross-section over time. They estimate a simple probit model to predict individual (or cohort) i 's stock-market participation y at time t as follows:

$$\Pr(y_{i,t} = 1|x_{i,t}, A_{i,t}(\lambda)) = \Phi(\alpha + \beta A_{i,t}(\lambda) + \gamma'x_{i,t}), \tag{7}$$

where $x_{i,t}$ includes a battery of income, wealth, and demographic controls as well as time effects.

The key explanatory variable is $A_{i,t}(\lambda)$, which is the weighted sum of past lifetime experiences, $A_{i,t}(\lambda) = \sum_{k=1}^{age_{i,t}-1} w_{it}(k, \lambda, age_{i,t})R_{t-k}$. Returns R_{t-k} are annual real S&P 500 returns, running from the end of the previous year ($t-1$) back to the end of the birth year ($t - (age_{i,t} - 1)$) of investor i . Multiplying weight $w_{it}(k, \lambda)$ with β yields the partial effect of a return experienced k years ago on the dependent variable (stock-market participation) for a household of i 's age.

A key ingredient, that has been used in multiple follow-up papers, is the weighting function used in the estimation, $w_{it}(k, \lambda, age_{i,t})$, which was introduced in (3).⁸ It allows weights to decline, increase, or remain constant over time,⁹ while introducing only one additional parameter (λ). Malmendier and Nagel (2011) use maximum likelihood to simultaneously estimate λ and the coefficient of interest β . This approach lets the data determine which weighting scheme best explains households' stock-market participation.

The results in Malmendier and Nagel (2011) suggest a strong and persistent effect of prior lifetime experiences of stock-market realizations on stock-market participation. Comparing an investor with fairly poor stock-market experiences in her life so far (at the tenth percentile of the sample) with an investor with very positive stock-market experiences (at the ninetieth percentile), the estimates imply an increase in the probability of stock-market participation of 10.2 pp. The effect is very large both in absolute and relative to the sample average of stock-market participation of 34.2%.

The estimate of the weighting parameter λ is also of interest. With $\lambda = 1.3$, the weighting function lies between the purple (dashed) line and the blue (dash-dotted) line, in Figure 3. It implies that households' stock-market participation decisions are most strongly

- 8 In the theoretical model and its illustration with two-period lived agents, return experiences include realizations in the current year ($k=0$) and go back to the beginning of the birth year ($k = age$), as shown in (2) and (3). In the empirical analysis, it is preferable to exclude concurrent realizations from the proxy for past experiences in order to distinguish alternative interpretations. Moreover, given the annual availability of much of the data, the end of the birth year is the natural starting point. Correspondingly, Malmendier and Nagel (2011) and later empirical papers specify $w_{it}(k, \lambda, age_{i,t}) = \frac{(age_{i,t}-k)^\lambda}{\sum_{k'=1}^{age_{i,t}-1} (age_{i,t}-k')^\lambda}$. In practice, both formulations imply virtually identical relative weighting.
- 9 Malmendier and Nagel (2011) also experimented with quadratic weighting functions and other functional forms. They found that allowing for "humps," U-shaped weights, step function, and other non-monotonocities did not improve fit. One caveat, though, is that this and other studies use repeated cross-sectional data; longer-term panel data might allow for a sharper distinction between alternative weighting function.

influenced by recent returns, but that returns many years in the past still exert a significant influence.

Interestingly, a different set of estimations, which use survey respondents' self-assessed willingness to take risk in the stock market, rather than actual stock holdings, as the outcome variable, generates a very similar λ estimate ($\hat{\lambda} = 1.4$). Hence, for both risk-taking measures, a significant part of the variation can be traced to variation in experienced real stock-market returns, with roughly similar weights on the history of past returns.

These results establish stylized fact (i) in stock-market data, namely, the presence of longlasting experience effects. The robust estimate of a $\hat{\lambda} > 1$ also confirms stylized fact (ii), namely, the embedded recency bias.

The two findings translate into significant cross-cohort differences in stock-market participation, and in differences in these differences over time, as illustrated in [Figure 4](#). The figure, adapted to more recent data from [figure III.a](#) in [Malmendier and Nagel \(2011\)](#), plots differences in the stock-market participation rates of cohorts above 60 years versus those 40 years or younger over time. The black (dotted) line is the raw difference in rates; the blue (dashed) line plots predicted differences based on fitted stock-market participation probabilities from the probit model; and the red (solid) line shows the counterfactual differences in fitted probabilities when the experienced stock return is set to the average annual return since 1871 until the year prior to the survey year. The plot reveals that the raw and predicted differences are quite closely aligned, confirming the high predictive power of experience effects. The counterfactual plot, however, features some large deviations, especially during the 1970s and 1980s, when the old-minus-young differences would have been much lower and in fact negative rather than positive. Conversely, it would have been more positive in the 1960s and mid-1990s to early 2000s. These are the periods when experienced effects are estimated to have their biggest impact.

Experience effects can also be detected in bond-market data. An interdecile-range shift in lifetime experiences of bond returns (long-term government bonds) predicts a change in bond-market participation of 11.4 pp, relative to a sample average of 37.6%. The estimated weighting parameter is also remarkably similar to the estimates from the elicited risk-tolerance and stock-market participation model, $\hat{\lambda} = 1.3$.

I mention the bond-market results here since they help establish stylized fact (iii) on domain specificity. When [Malmendier and Nagel \(2011\)](#) include experienced stock returns in a regression predicting bond-market participation or, vice versa, include experienced bond returns in a regression predicting stock-market participation, neither experience has predictive power. There is no cross-fertilization between the stock and the bond market.

Finally, [Malmendier and Nagel \(2011\)](#) provide some evidence regarding stylized fact (iv), the persistence of experience effects among more informed investors. They test how the strength of the estimated effect varies with financial sophistication using (i) a dummy for having above-median liquid assets and (ii) a dummy for completion of a college degree. They find that the interactions of these proxies with the experienced return variable (cf. their Appendix-Table A.3) are never statistically nor economically significant. That is, there is little difference between households with high and low financial sophistication in the strength of experience effects. While their SCF and ICPSR data do not allow for more direct proxies, related studies show experience effects directly among professional investors, including, for example, mutual-fund managers ([Greenwood and Nagel, 2009](#)).

Before we turn to tests of the additional implications derived in the theoretical framework, let's briefly touch on the beliefs-versus-preferences discussion. The analysis in

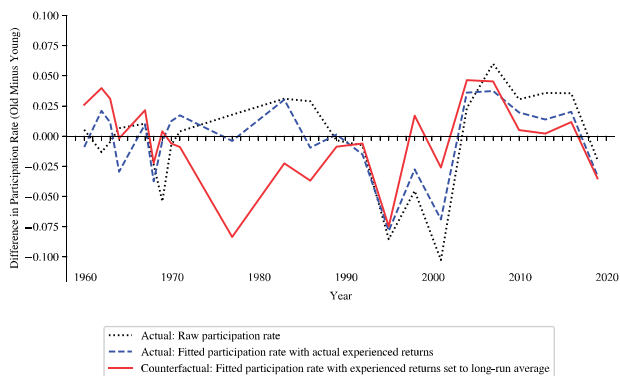


Figure 4. Difference in stock market participation rates: Old (age > 60 years) minus young (age ≤ 40 years): Difference in average fitted stock market participation probabilities between young (age ≤ 40) and old (age > 60 years) (old minus young) from the probit model (dashed line), compared with the counterfactual difference in average fitted probabilities when the experienced stock market return is set to the average annual return since 1871 until the year prior to the survey year (solid line), and the actual participation rate in the raw data (dotted line).

Malmendier and Nagel (2011) provides additional results that directly relate to beliefs about future stock returns. Using 1998–2007 UBS/Gallup survey data from the Roper Center (now at the University of Cornell) on return expectations, they estimate that a 1 pp increase in experienced stock returns is associated with a 0.5–0.6 pp increase in expected 12-month stock-market returns and a 0.6–0.7 pp increase in expected 12-month returns on respondents’ own portfolio. These findings do not rule out that life-time return experiences affect risk preferences as well, but at a minimum, the beliefs channel captures an important part of the experience effects.

4.3 Experience-Based Market Dynamics

The model of EBL has further implications for aggregate valuations and the dynamics of stock-market participation and trading. I present four sets of findings, and point to open questions and future research directions.

Market Valuations

A first implication regards aggregate valuation levels. As experienced returns fluctuate over time, they alter aggregate risk-taking in the stock market. Figure 5, which updates figure IV from Malmendier and Nagel (2011), illustrates this point. The plot relates population-wide averages in lifetime experienced returns, shown as red bars, to the price/earning (P/E) ratio, shown as the blue solid line. The experience-effect model predicts that periods of high experienced returns (and hence high willingness to invest in the stock market) should coincide with periods of high price/earnings ratios, that is, high market valuations. Indeed, the figure indicates a strong positive correlation between aggregated experienced stock-market return of US investors and the P/E ratio. Periods of high equity-market valuations (the 1960s, 1990s, and 2010s) coincide with periods when the average investor has experienced high stock-market returns over their life so far. Periods of low valuation and high subsequent returns (1940s, early 1980s, and to some extent late 2000s with the financial crisis) coincide with the average investor having experienced low stock-market returns over their life

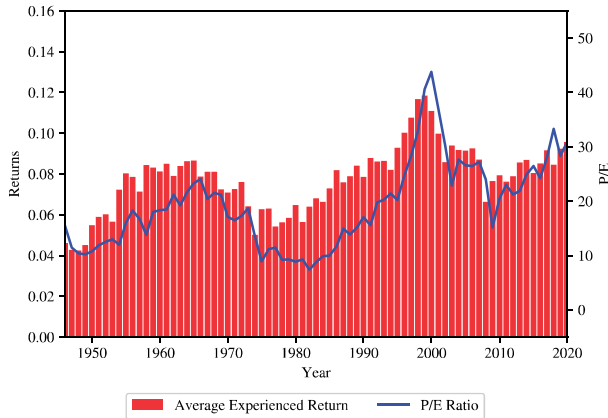


Figure 5. Aggregate perspective: Market Valuation. Average Experienced Returns are calculated as a weighted average of lifetime real S&P 500 returns across all age groups at a given point in time (from age 25 to age 74), where the weights are the percentage share of aggregate liquid assets held by the respective age group (on average across SCF waves) and the weighting over individuals' lifetimes uses the average estimated λ parameter across the different estimations in Malmendier and Nagel (2011), $\lambda = 1.5$. The annual P/E Ratio is an updated series from Shiller (2005), which is calculated with a 10-year moving average of trailing earnings of firms in the S&P 500 index in the denominator. Since the stock return data dates back to 1871 and we consider age up to 75, we can only compute average experience from 1946 onward.

so far. This correlation is not mechanical as it would not arise for other plausible values of the weighting parameter λ that the cross-sectional differences in the SCF microdata could have generated.

The empirical estimates imply that EBL about stock returns is a plausible explanation for time-variation in the aggregate demand for risky assets and the resulting variation in stock-market valuation levels and P/E ratios. A deeper analysis of this back-of-the-envelope illustration, also separately for different asset classes, is a promising avenue for future research.

Market Participation

A second aggregate implication regards the time variation in market participation: differences in cohorts' lifetime experiences of past stock returns predict differences in cohorts' willingness to participate in the stock market. Cohorts that have experienced particularly high stock returns over their lives so far should be more heavily represented in the stock market than those cohorts whose lifetimes so far featured lower average returns. In other words, a theory of experience effects carries predictions for asset-market composition and its fluctuation over time. While much of the prior research has focused on life-cycle variations, experience effects allow for a different perspective: not only age, but also prior life experiences influence market participation.

Figure 6 shows that the dynamic participation predictions hold in the data.¹⁰ For visualization purposes, the figure combines again all cohorts over sixty into the set of "old

¹⁰ The figure updates both Figure 1 from Malmendier and Nagel (2011) and Figure 5 from Malmendier, Pouzo, and Vanasco (2020b).

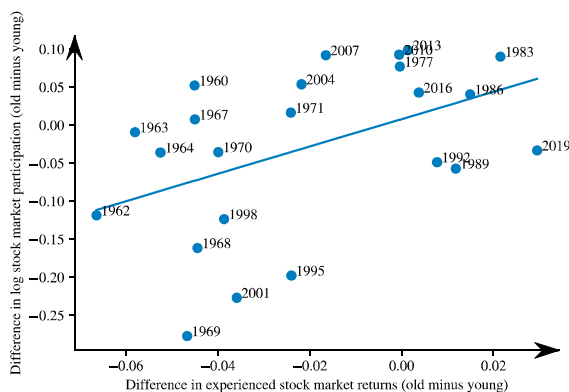


Figure 6. Aggregate perspective: market composition. Differences in stock-market participation rates of old and young individuals plotted against differences in experienced stock returns.

cohorts,” and all cohorts aged 40 years or below into the set of “young cohorts.” It then plots the difference in their average lifetime return experiences, defined as average S&P 500 returns in the past 50 years versus the past 20 years, against the difference in their (log) stock-market participation rates. The figure reveals a clear positive correlation, which has in fact become stronger compared to the original version in [Malmendier and Nagel \(2011\)](#), which ended in 2007. This is also echoed in several variants of the graph that use lifetime experiences in real dividends, earnings, or (log) GDP instead of stock returns (cf. [Malmendier, Pouzo, and Vanasco 2020b](#)).

The pattern confirms the predictive power of experience effects for a better understanding of market participation. It also points to future research possibilities: Can we dig deeper and better capture the differential experiences of those individuals who never participate in the stock market (or drop out)? Do they update their beliefs differently? In fact, might negative beliefs that caused nonparticipants not to enter or to drop out of the stock market be sticky as they do not “experience” current, more positive realizations as closely as participants do?

Another direction for future research is a more asset-specific analysis. Rather than considering stock market investment in general, can we be more specific about different cohorts’ experiences with different types of assets? An example is the analysis of prior investment experiences in specific industries in [Huang \(2019\)](#). She finds that positive return experiences in a given industry increase the likelihood of further investment in the same industry. Zooming in further, we might ask whether this pattern also holds on the stock level, or zooming out, we might turn to assets other than shares.

Finally, it would be interesting to further explore the interaction of cross-cohort differences and recent experiences versus those further in the past. Do younger cohorts react more strongly to a (recent) dividend shock than older cohorts as it makes up a larger part of their lifetimes? More specifically, does a positive shock induce younger cohorts to invest relatively more in risky assets, while a negative shock tilts the composition toward older cohorts? How about earlier shocks, which both younger and older cohorts have experienced? As [Figure 3](#) illustrates, some of the earlier shocks (say, 35–40 years ago) should be weighted more heavily by an older, 60-year-old individual than a younger, 40-year-old individual. Can we find data that allow us to tease out these differential effects?

Trade Volume

The emphasis on cross-sectional differences in experience-based beliefs allows us to turn to yet another aspect of market dynamics, namely, trading decisions. Does EBL help us better understand fluctuations in trade volume among market participants? The theory model in Section 3 revealed that it is not the differences in lifetime experiences (and resulting level of disagreement between cohorts) that predict trading, but changes in those differences. The prediction is that trade volume is higher as disagreement increases.

Malmendier, Pouzo, and Vanasco (2020b) test precisely this implication, using the detrended turnover ratio as their proxy for trade volume.¹¹ Figure 7 plots the corresponding deviations of the turnover ratio from the trend (red solid line) against changes in experience-based disagreement between cohorts, calculated as the standard deviation of experienced returns (dashed blue line).¹² The plot, which updates Figure 9 from Malmendier, Pouzo, and Vanasco (2020b), reveals that increases in experience-based disagreement strongly predict higher abnormal trade volume. Together with the findings on market valuation and market participation, the figure confirms the explanatory power of EBL for market dynamics. The results on trading behavior also give rise to additional avenues for future research. More fine-grained data would allow for asset-specific analyses, comparing trading in those assets with the asset-specific prior experiences.

International Capital Flows

As an example of an analysis that moves away from the US stock market, I briefly report results from a related study that takes the notion of experience effects to the international context and distinguishes between two types of asset classes—domestic and foreign.

Experience effects have direct implications for international macro models: Investors in different countries have different “experiences,” that is, differ in their exposures to domestic versus foreign outcomes. Different countries also feature different demographics, resulting in different cross-cohort differences and market composition. For example, the experience-effects model predicts that the impact of the COVID-19 crisis on risk attitudes and beliefs are different across countries depending on how strongly the pandemic affected the economy and the local stock-market, say, in Spain, where the IBEX 35 was still significantly below pre-pandemic levels a year later, versus neighboring France, where the CAC 40 reached and exceeded pre-pandemic levels around the same time. Moreover, a higher fraction of market participants exposed to the crisis were relatively young in France compared to the (more strongly) aging population in Spain.

Malmendier, Pouzo, and Vanasco (2020a) demonstrate that experience effects help explain classic international macro puzzles regarding capital flows and portfolio investment. Specifically, they focus on the tendency of investors to

1. hold an over-proportional fraction of their equity wealth in domestic stocks (**home bias**, cf. Cooper and Kaplanis, 1986 and French and Poterba 1991);

11 Malmendier, Pouzo, and Vanasco (2020b) calculate the average monthly turnover ratio (shares traded divided by shares outstanding), weighted by market-capitalization across all firms in January and in December of the preceding year. They then log, linearly detrend, and CF-filter the yearly variable to obtain the deviation of turnover ratio from the trend.

12 The variable takes the change in experienced returns as of the current year relative to the prior year, and calculates the current-year age-cohort population-weighted standard deviation of this difference variable.

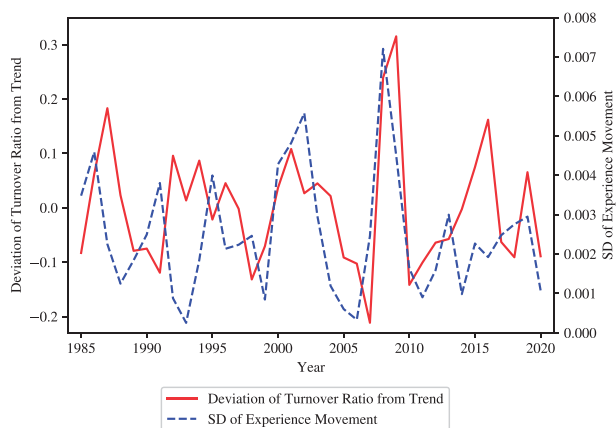


Figure 7. Aggregate perspective: market dynamics.

2. invest in domestic equity markets in periods of domestic or global crises (**retrenchment**, cf. [Forbes and Warnock, 2012](#); [Broner et al. 2013](#)); and
3. withdraw capital from foreign equity markets in periods of foreign and global crises (**fickleness**, cf. [Forbes and Warnock, 2012](#); [Caballero and Simsek, 2018](#)).

The basic intuition relies on investors having had more exposure to domestic risky-asset returns than foreign investors. In that sense, foreign investors behave similarly to younger investors in the baseline experience-effect model. [Malmendier, Pouzo, and Vanasco \(2020a\)](#) show that EBL jointly rationalizes the above puzzles.¹³ First, the setting generates home bias in portfolio holdings as agents perceive domestic output as less risky. Hence, they tend to overweight domestic assets. Second, agents overreact to foreign output realizations because they are generally less confident about what they know about the foreign country. As a result, negative shocks, for example, low output realizations abroad, together with general equilibrium effects imply that capital inflows of domestic agents (retrenchment) and outflows of foreign agents (fickleness) both increase. The reverse holds for booms. That is, crises are followed by an increase in domestic inflows and foreign outflows of capital, and booms are followed by the corresponding decrease, even though all agents perceive the shocks the same way, as negative or positive.

In addition, EBL generates new testable predictions based on the demographic composition of different countries. EBL implies that countries with a larger number of young market participants overreact (more) to both domestic and foreign shocks. Hence, retrenchment and fickleness are both alleviated in young-demographics countries, and exacerbated in old-demographics countries.

All predictions are tested and confirmed in data from the IMF, World Bank, and the World Federation of Exchanges. Thus, the concept of experience effect provides a unifying framework for a range of international macro-puzzles. At the same time, both the theoretical and the empirical analysis are rather “high-level,” often due to data constraints on the

13 To simplify the theoretical analysis, [Malmendier, Pouzo, and Vanasco \(2020a\)](#) endow agents with more precise prior beliefs about domestic than foreign output, which capture that they are more confident about their knowledge of their own country than a foreign country.

precise nature of the assets, details of investment and holding patterns, and the demographics of different countries. The international context is probably the setting that provides for the richest and most promising set of further analyses for researchers with access to more fine-grained country-specific data.

5. Related Approaches

Before concluding with some further suggestions for future research, I would be remiss not to point to other modern formalizations of belief formation, many of which are able to capture several—though typically not all—of the four stylized features.

A large strand of the literature focuses on over-inference from more recent observations, including models of natural expectation formation (Fuster, Hebert, and Laibson, 2011; Fuster, Laibson, and Mendel, 2010) and over-extrapolation (Barberis *et al.*, 2015, 2016). These approaches capture key feature (ii), the significant recency bias in investors' belief formation. There is also evidence of such over-extrapolation occurring in professional experts such as top managers running pension funds (Andonov and Rauh, 2020), akin to feature (iv), the robustness to learned expert knowledge.

Other approaches incorporate both (i) long-lasting scarring effects and (ii) recency bias. An example is the model in Kozlowski, Veldkamp, and Venkateswaran (2020), which assumes that the experience of economic shocks changes how people think about tail risk. Agents in their model do not know the distribution of shocks, which induces updating and a long-lasting economic response after the realization of such a tail event. The latter approach does not naturally capture (iii) the domain specificity of experience-based belief formation and (iv) the robustness to learned (expert) knowledge, though. One could, of course, assume that individual updating about tail events is domain-specific, that is, individuals ignore or are uninformed about the correlation between the returns of related assets, though this seems implausible at least for professional forecasters, bankers, and other professionals.

The most fundamental difference between the concept of experience effects and most other modern models of belief formation, such as the ones cited above, is that the latter models still rely on “information” as the core determinant. There is no room for “experiences” to be stronger than learned information. Instead, information-based theories tend to attribute the effect of “experiences” to information constraints, possibly paired with heterogeneous priors and model uncertainty assumptions. The issue is that “learned” information will not necessarily undo experience effects. The lasting effect of experiences reflects a “re-wiring” that occurs as we accumulate experiences throughout our lives.

Modern theoretical approaches that directly account for such re-wiring, outside the literature on experience effects, tend to fall in the realm of memory research. A first example is Wachter and Kahana (2019), who introduce the concept of retrieved-context memory to financial decision-making. They define “context” as a record of associations that arise from experiences and the environment. “Retrieved context” is the mechanism through which this belief system responds to the current environment and which directly links to EBL. As the authors emphasize: “Unless agents have full access to all decision-relevant information at the moment of choice, they must use their memory of past experiences to guide their decisions.”

The authors also introduce the concepts of similarity, contiguity, and recency as the “three major laws governing the human memory system.” Similarity is a concept that has

been incorporated into some economic theories, notably Gilboa and Schmeidler's (1995) case-based decision theory, though their axioms leave the form of similarity unrestricted. Psychologists, instead, define similarity to capture the brain's tendency to recall past information which is similar to currently active features or experiences. Relatedly, contiguity refers to the idea that people are more likely to recall features from past memory which have previously co-occurred with features or experiences that are currently active. And recency naturally captures the concept from stylized feature (ii). The first two, similarity and contiguity provide for a neuroscience underpinning of stylized feature (iii), domain specificity. As the authors write, they imply that "remembering an item involves a jump-back-in-time to the state of mind that obtained when the item was previously experienced," generating—when triggered—feature (i), the long-lasting imprint.

A related model by Bordalo, Gennaioli, and Shleifer (2020) is also in line with the notion of experience effects and its neuroscience foundations. The authors emphasize that the basic mechanisms of memory make information embedded in past experiences disproportionately accessible to decision makers. As such, their approach naturally captures the stylized features (i), (ii), and (iv). Moreover, much of their model revolves around the role of "context," which is related to (iii) domain specificity, as we just discussed. They emphasize that contextual stimuli trigger recall of past experiences based on contextual similarity.

Overall, a modern understanding of memory-based anchoring and adjustment mechanisms provides a direct underpinning for the evidence on experience effects and EBL. At the same time, this work still lacks clarity about what exactly characterizes a "context" and what exactly qualifies as "similar," mirroring the lack of clarity around "domains" and "domain specificity" in the literature on experience effects. Here, progress likely needs to come from the side of neuroscience to allow for more predictive economic theories and empirical tests.

6. Implications for Future Research

The neuroscientific and theoretical underpinnings of experience effects as well as the empirical findings have allowed us to discern some of the key aspects of how our past experiences influence our beliefs, attitudes, and decision-making. I have emphasized four baseline findings: (i) the notion of a long-lasting imprint, which is slowly altered over time as individuals accumulate new experiences; (ii) the observation of a bias toward more recent experiences; (iii) the domain specificity, or lack of cross-fertilization, in experience-based belief formation; and (iv) the robustness of these findings to "learned knowledge," including that of professionals and experts. Both the theoretical and the empirical findings have then pointed us toward further implications of experience effects and these four features. The notion of a long-lasting imprint gives rise to (i) persistent cross-cohort differences, and recency bias triggers (ii) a stronger (over-)reaction of younger cohorts to recent shocks than older cohorts. Moreover, domain specificity invites (iii) more asset- or market-specific research; and the robustness to "learned knowledge" invites (iv) more analyses of professional, highly trained subjects such as doctors, fund managers, or central bankers.

Going forward, it would be desirable for theoretical work to aim not only for more psychological and empirical realism, as much of the behavioral finance literature has done already, but to also build more closely on the neuroscientific evidence of brain functions and brain re-wiring. Concepts such as "similarity" and "contiguity" and "context" appear to not have been fully sorted out in terms of their meaning for financial and economic decision

making. It remains to be fleshed out what exactly delineates a domain and how “similarity” and “contiguity” help to identify it. We might also ask whether “domains” can be actively (re-)shaped, rather than being set in stone. For example, I discussed that [Malmendier and Nagel \(2011\)](#) show how investors who get burned in the stock market, do not seem to change their beliefs or behaviors regarding the bond market. This domain specificity, though, relies implicitly on a categorization that originally creates those domains. That is, the observed behavior occurs at a time when stock and bond markets are fairly well segregated, with different news stories, analysts, and other resources dedicated to one or the other. But if, at different times or in different countries, stocks and bonds were not categorized as cleanly as (quasi) separate, or there was a larger convertible market, one might observe less specificity in behavior with respect to experiences in those domains.¹⁴ It would be very interesting to know the price and trading implications of widening and narrowing domains, and identify the effects empirically.

Another angle in need of theoretical development is work moving beyond information-based theories—in the sense “information” is currently understood. An important step forward, in terms of theoretical development, has been the move from theories of over-extrapolation and over-inference, that still have “theoretically understood information” at their core, to theories that emphasize “encoding” and “retrieval.” An early predecessor is the case-based decision theory of [Gilboa and Schmeidler \(1995\)](#), whose mapping from “problems” and “actions” to results clearly sees the importance of individually experienced situations. While Gilboa and Schmeidler dispense entirely with the notion of beliefs, later models such as [Bordalo *et al.* \(2019\)](#) and [Wachter and Kahana \(2020\)](#) incorporate a role of “context” while remaining grounded in a (quasi-)Bayesian modeling framework. For a more extensive discussion, see [Malmendier and Wachter \(2021\)](#).

More closely related to the model proposed here, one may ask whether the weights on prior lifetime experiences should truly be decreasing over all age ranges. While capturing recency effects, such parameterization misses out on other aspects of brain function, for example, higher brain plasticity earlier in life. For example, some childhood experiences might be particularly influential as the evidence in [Duchin, Simutin, and Sosyura \(2021\)](#) (on the long-term effects of CEOs growing up in a male-dominated family on how they allocate capital to male versus female divisional managers) suggests.

On the empirical side, the applicability of “experience effects” within finance has already proven to be far-reaching. We have discussed papers on stock-market participation, trading behavior, and international capital flows. There are also numerous applications to consumption and to inflation expectations in macro-finance, especially monetary economics, for example, [Malmendier and Shen \(2018\)](#); [Malmendier and Nagel \(2016\)](#); [Malmendier, Nagel, and Yan \(2021\)](#) (for an overview cf. [Malmendier \(2021\)](#)). However, we still have little evidence on how past experiences shape other dimensions of household financial decision-making. For example, whether individuals accumulate too much debt or too little (debt aversion) given their financial situation, might be affected by experiences they made earlier in life, or they saw their parents make. Similarly, prior exposures could help explain the large heterogeneity in savings rates. Do early experiences with the benefits of saving make kids life-long savers, as sometimes said about the Germans and their tradition of local S&L banks opening saving accounts for first-graders? Other applications include insurance choices—from home to life and fire, flood, or earthquake insurances.

14 I thank Chris Parsons for making this point.

While differences in the choice of coverage have been related to demographic and financial determinants, including proxies for risk aversion, EBL points to the deeper question of where individual-level differences in the willingness to pay for insurance may come from. What role does prior exposure to negative realizations of the underlying risk play?¹⁵

The household finance applications also illustrate that there is potential for more research on personal experiences. While some of the papers discussed earlier, such as [Kaustia and Knüpfer \(2008\)](#) and [Choi et al. \(2009\)](#), leverage truly personal experiences, much of the experience-effect literature uses exposure to macro-finance realizations. The latter approach comes with the advantage of avoiding certain endogeneity concerns; but it misses some of the experiences that are encoded most strongly. Personal experiences also offer a pathway to incorporating parents' experiences, family background, and other channels of intergenerational transfer into the research on experience effects.

For all of these potential applications, researchers will benefit from the increasing availability of within-person "big data." These data sources allow us to distinguish standard information-based explanations from attributing effects to personal experiences, and to make progress on answering some of the open questions.

References

- Andonov, A. and Rauh, J. D. (2020): The return expectations of institutional investors. Working Paper.
- Barberis, N., Greenwood, R., Jin, L., and Shleifer, A. (2015): X-capm: an extrapolative capital asset pricing model, *Journal of Financial Economics* 115, 1–24.
- Barberis, N., Greenwood, R., Jin, L., and Shleifer, A. (2016): Extrapolation and bubbles. Technical report, National Bureau of Economic Research, Cambridge, Massachusetts.
- Bear, M., Connors, B., and Paradiso, M. (2020): *Neuroscience: Exploring the Brain, Enhanced Edition*. Jones & Bartlett Learning, LLC, Burlington, Massachusetts.
- Bianchi, F., Ludvigson, S. C., and Ma, S. (2020): Belief distortions and macroeconomic fluctuations. Technical report, National Bureau of Economic Research, Cambridge, Massachusetts.
- Bissiri, P. G., Holmes, C. C., and Walker, S. G. (2016): A general framework for updating belief distributions, *Journal of the Royal Statistical Society Series B, Statistical Methodology* 78, 1103.
- Bliss, T. V. and Lomo, T. (1973): Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path, *The Journal of Physiology* 232, 331–356.
- Bordalo, P., Gennaioli, N., Porta, R. L., and Shleifer, A. (2019): Diagnostic expectations and stock returns, *The Journal of Finance* 74, 2839–2874.
- Bordalo, P., Gennaioli, N., and Shleifer, A. (2020): Memory, attention, and choice, *Quarterly Journal of Economics* 135, 1399–1442.
- Botsch, M. J. and Malmendier, U. (2020): The long shadows of the great inflation: evidence from residential mortgages. Working Paper.
- Botzen, W. W. and van den Bergh, J. C. (2012): Risk attitudes to low-probability climate change risks: wtp for flood insurance, *Journal of Economic Behavior & Organization* 82, 151–166.
- Broner, F., Didier, T., Erce, A., and Schmukler, S. L. (2013): Gross capital flows: dynamics and crises, *Journal of Monetary Economics* 60, 113–133.

15 Some of the existing insurance literature has related past natural disasters to a higher uptake of the corresponding type of insurance; see for example [Botzen and van den Bergh \(2012\)](#) and [Froot and O'Connell \(2007, 2008\)](#).

- Caballero, R. J. and Simsek, A. (2018): A model of fickle capital flows and retrenchment. Technical report, National Bureau of Economic Research, Cambridge, Massachusetts.
- Choi, J. J., Laibson, D., Madrian, B. C., and Metrick, A. (2009): Reinforcement learning and savings behavior, *The Journal of Finance* 64, 2515–2534.
- Collin-Dufresne, P., Johannes, M., and Lochstoer, L. A. (2017): Asset pricing when ‘this time is different’, *The Review of Financial Studies* 30, 505–535.
- Cooper, I. A. and Kaplanis, E. (1986): Costs to crossborder investment and international equity market equilibrium, *Recent Developments in Corporate Finance*, Cambridge University Press, Cambridge, pp. 209–240.
- Doidge, N. (2007): *The Brain That Changes Itself: Stories of Personal Triumph from the Frontiers of Brain Science*. Penguin, London.
- Duchin, R., Simutin, M., and Sosyura, D. (2021): The origins and real effects of the gender gap: evidence from ceos’ formative years, *The Review of Financial Studies* 34, 700–762.
- Ehling, P., Graniero, A., and Heyerdahl-Larsen, C. (2018): Asset prices and portfolio choice with learning from experience, *The Review of Economic Studies* 85, 1752–1780.
- Forbes, K. J. and Warnock, F. E. (2012): Capital flow waves: surges, stops, flight, and retrenchment, *Journal of International Economics* 88, 235–251.
- French, K. R. and Poterba, J. M. (1991): Investor diversification and international equity markets, *The American Economic Review* 81, 222–226.
- Frey, U. and Morris, R. G. (1997): Synaptic tagging and long-term potentiation, *Nature* 385, 533–536.
- Froot, K. A. and O’Connell, G. (2007): *The Pricing of US Catastrophe Reinsurance*. University of Chicago Press, Chicago, Illinois.
- Froot, K. A. and O’Connell, P. G. (2008): On the pricing of intermediated risks: theory and application to catastrophe reinsurance, *Journal of Banking & Finance* 32, 69–85.
- Fuster, A., Laibson, D., and Mendel, B. (2010): Natural expectations and macroeconomic fluctuations, *Journal of Economic Perspectives* 24, 67–84.
- Fuster, A., Hebert, B., and Laibson, D. (2011): Natural expectations, macroeconomic dynamics and asset pricing, in: D. Acemoglu and M. Woodford (eds), *NBER Macroeconomics Annual*, University of Chicago Press, Chicago, Illinois. pp. 1–48.
- Gilboa, I. and Schmeidler, D. (1995): Case-based decision theory, *The Quarterly Journal of Economics* 110, 605–639.
- Graham, J. R. and Narasimhan, K. (2004): Corporate survival and managerial experiences during the great depression, *Experimental & Empirical Studies eJournal*
- Greenwood, R. and Nagel, S. (2009): Inexperienced investors and bubbles, *Journal of Financial Economics* 93, 239–258.
- Huang, X. (2019): Mark twain’s cat: industry investment experience, categorical thinking and stock selection, *The Journal of Finance Economics* 131, 404–432.
- Karmiloff-Smith, A. (2015): An alternative to domain-general or domain-specific frameworks for theorizing about human evolution and ontogenesis, *AIMS Neuroscience* 2, 91–104.
- Kaustia, M. and Knüpfer, S. (2008): Do investors overweight personal experience? evidence from ipo subscriptions, *The Journal of Finance* 63, 2679–2702.
- Kozlowski, J., Veldkamp, L., and Venkateswaran, V. (2020): The tail that wags the economy: beliefs and persistent stagnation, *Journal of Political Economy* 128, 2839–2879.
- Kraus, A. and Sick, G. A. (1980): Distinguishing beliefs and preferences in equilibrium prices, *The Journal of Finance* 35, 335–344.
- Malmendier, U. (2018): Behavioral corporate finance, in: D. Bernheim, S. DellaVigna, and D. Laibson (eds.), *Handbook of Behavioral Economics*, Vol. 1, Elsevier North-Holland, Amsterdam, Netherlands. pp. 277–379.
- Malmendier, U. (2021): Exposure, experience, and expertise: why personal histories matter in economics. Working Paper.

- Malmendier, U. and Nagel, S. (2011): Depression babies: do macroeconomic experiences affect risk taking, *The Quarterly Journal of Economics* 126, 373–416.
- Malmendier, U. and Nagel, S. (2016): Learning from inflation experiences, *The Quarterly Journal of Economics* 131, 53–87.
- Malmendier U. and Shen L. S. (2018): Scarred consumption. Technical report, National Bureau of Economic Research, Cambridge, Massachusetts.
- Malmendier, U. and Tate, G. (2005): Ceo overconfidence and corporate investment, *The Journal of Finance* 60, 2661–2700.
- Malmendier, U. and Wachter, J. (2021): Memory of past experiences and economic decisions. Working Paper.
- Malmendier, U., Tate, G., and Yan, J. (2011): Overconfidence and early-life experiences: the effect of managerial traits on corporate financial policies, *The Journal of Finance* 66, 1687–1733.
- Malmendier, U., Pouzo, D., and Vanasco, V. (2020a): Investor experiences and international capital flows, *Journal of International Economics* 124, 103302.
- Malmendier, U., Pouzo, D., and Vanasco, V. (2020b): Investor experiences and financial market dynamics, *Journal of Financial Economics* 136, 597–622.
- Malmendier, U., Nagel, S., and Yan, Z. (2021): The making of hawks and doves, *Journal of Monetary Economics* 117, 19–42.
- Rabin, M. (1998): Psychology and economics, *Journal of Economic Literature* 36, 11–46.
- Savage, L. J. (1954): *The Foundations of Statistics*, Wiley, New York, NY.
- Schoar, A. and Zuo, L. (2017): Shaped by booms and busts: how the economy impacts CEO careers and management styles, *The Review of Financial Studies* 30, 1425–1456.
- Shiller, R. J. (2005): *Irrational Exuberance*, 2nd ed., Princeton, NJ: Princeton University Press.
- Spunt, R. P. and Adolphs, R. (2017): A new look at domain specificity: insights from social neuroscience, *Nature Reviews Neuroscience* 18, 559–567.
- Tversky, A. and Kahneman, D. (1974): Judgment under uncertainty: heuristics and biases, *Science* 4157, 1124–1131.
- Vissing-Jørgensen, A. (2003): Perspectives on behavioral finance: does “irrationality” disappear with wealth? Evidence from expectations and actions, *NBER Macroeconomics Annual* 18, 139–194.
- Wachter J. A. and Kahana M. J. (2019): A retrieved-context theory of financial decisions. Technical report, National Bureau of Economic Research, Cambridge, Massachusetts.
- Wachter, J. A. and Kahana, M. J. (2020): Associative learning and representativeness. SSRN 3602327. Working Paper.
- Weber, E. U., Böckenholt, U., Hilton, D. J., and Wallace, B. (1993): Determinants of diagnostic hypothesis generation: effects of information, base rates, and experience, *Journal of Experimental Psychology: Learning, Memory, and Cognition* 19, 1151.
- Weber, E. U., Blais, A. R., and Betz, N. E. (2002): A domain-specific risk-attitude scale: measuring risk perceptions and risk behaviors, *Journal of Behavioral Decision Making* 15, 263–290.
- Whitlock, J. R., Heynen, A. J., Shuler, M. G., and Bear, M. F. (2006): Learning induces long-term potentiation in the hippocampus, *Science* 313, 1093–1097.

