

Neurofinance

Forthcoming in “Handbook of Financial Decision Making”, Edward Elgar Publishing,
Gilles Hillary and David Mclean Editors

Abstract

A nascent field combining theory and methodology from neuroscience and finance, neurofinance takes what we know at the neurobiological level and applies it to investor decision-making. It aims to propose new theories about investor behavior and to test existing theories that are difficult to test with behavioral data alone but can however be tested with a combination of behavioral and neural data. The path taken in recent research has uncovered neurobiological foundations for well-known anomalies in investor behavior, revealing enticing opportunities for improving investor decision-making in the future.

Keywords: Neurofinance, neuroeconomics, behavioral finance, decision-making under uncertainty, risk taking, neuroforecasting

1. Introduction

“Financial behavior that may seem irrational now is behavior that hasn’t had sufficient time to adapt to modern contexts across evolutionary time. An obvious example from nature is the great white shark, a near perfect predator that moves through the water with fearsome grace and efficiency, thanks to 400 million years of adaptation. But take that shark out of the water and drop it onto a sandy beach, and its flailing undulations will look silly and irrational. It’s perfectly adapted to the depths of the ocean, not to dry land.” (Lo (2017), p.9)

1.1 Homer Simpson versus Zelig—what of the investor brain?

The nascent scientific field neurofinance combines theory and methodology from neuroscience and finance to study how the brain perceives and reacts to financial risks. These emerged only recently in evolutionary time, implying little time for the brain to adapt. After all, evolution shaped our brain to maximize survival in the ancestral lifestyle. This begs the question of which aspects of modern financial risk-taking clash with the latter. Take, for example, the challenge of risk appraisal—learning about the risk/reward profiles of potential investment opportunities. We will see later in this chapter that the same investor who repeatedly makes suboptimal investing choices in contexts where the investment set is presented through explicit statistics may well be able to choose the right investment in contexts that are more congenial to the brain, for instance, if the investor becomes acquainted with the investment set through return simulations. This is one example of context-dependent behavior.

The notion of context dependence constitutes an important paradigmatic departure from the common belief in the social sciences that “bounded rationality” resonates with “irrationality”.¹ Thaler (2016) famously represented that view using the figure of Homer Simpson as a metaphor of human limited intelligence to be contrasted with the idealistic (“Spock-like”) Homo Economicus model. Neuroscientists would instead advance the figure of “Zelig”—the human chameleon depicted by Woody Allen—to stress the context dependence of human cognition—the fact that investor behavior is expected to be Spock-like in contexts designed to match the ancestral lifestyle in key regards (I elaborate next) and Homer-like otherwise.

1.2 Practical implications

A key challenge for behavioral finance research, therefore, is to examine how decisions come about in the brain and to pin down the contextual elements of investor decision-making to which the brain is not well adapted. The ultimate purpose is to help improve investor decisions by removing those elements of mismatch between the brain and modern financial risk-taking. This goal involves both reshaping the investor brain and redesigning key aspects of the investing landscape to better match the investor brain (“*neuro-ergonomics*”).

Reshaping the investor brain In Section 2.1.2, I will discuss the topic of changing investors’ mental models related to market instability. Investors facing fast-changing market conditions tend to ignore it by default due to the nature of their [prior knowledge](#) regarding stability. Evidence indeed suggests that by default, investors consider their environment to be stable, which is at odds with the instability observed in many markets. Recent research points to the possibility of reshaping investor beliefs about stability so that investors facing fast-changing conditions become aware that instability is an inescapable element of their world.

Neuro-ergonomics Prospectuses furnished by financial institutions to investors typically show a recent history of asset returns in a numeral format. We will see in Section 2.1.2 that alternative information formats could be used instead to allow investors to grasp the statistics underlying realized asset returns. We will also see in Section 3.1 that making capital gains less salient on trading platforms may help “debias” the celebrated “*disposition effect*”—the robust bias whereby investors have a greater propensity to sell stocks trading at a gain relative to purchase price than stocks trading at a loss.

These are examples of behavioral interventions aimed at reducing the mismatch between the investor brain and the investing landscape. Recent research suggests that “neuro-ergonomics” could further involve leveraging matching elements between the two. Take, for instance, the way visual and auditory components systematically covary in trading floors: turbulent times mean both high asset price volatility displayed on trader screens and high ambient noise (what traders call “*the buzz*”), and low volatility is typically associated with quiet on the floor.² Evidence suggests that such covariation of visual volatility and ambient noise helps traders gauge volatility (Payzan-LeNestour et al., 2021) possibly because using both visual and auditory cues to predict events provided a survival advantage over using either modality alone in the ancestral lifestyle.³

¹ Tversky and Kahneman (1974), the cornerstone of what is called the “*heuristics and biases*” literature, inspired many studies across the social sciences on the role of cognitive biases and restrictions, such as limitations in the number of variables agents can keep track of or pay attention to. Classical examples in behavioral finance include Huberman and Regev (2001), Hirshleifer et al. (2009), and Gao et al. (2011).

² For example, Jones et al. (1994) and Coates (2012).

³ For example, Garcia and Koelling (1966) and Brown and Richerson (2014).

1.3 Structure of this chapter

Improving investor decision-making based on the logic explained above is difficult without pinning down the root causes of investor behavior at the neurobiological level. Research in neurofinance sets out to achieve this goal. By doing so, it contributes to our understanding of investor behavior at each stage of the scientific process. To illustrate this, I have framed this chapter around Charles Peirce’s description of the scientific process as a three-stage inquiry (Peirce, 1877). The first stage consists of conceiving new ideas and selecting those that are worth testing (“*abduction*”).⁴ Selected ideas are then made “*falsifiable*” (Popper, 1959), i.e., testable implications are drawn from them (“*deduction*”). Finally, these implications are tested with empirical data (“*induction*”). At the end of the process, ideas are either rejected or amended through abduction again, whereby the scientific process is akin to a loop (Fig. 1).

Research in neurofinance contributes to our understanding of investor behavior at each stage of this process. At the abduction stage, well-established neuroscientific knowledge can be harnessed to generate new ideas about investor behavior (Section 2). At the deduction and induction stages, theories about investor behavior that are difficult to test with behavioral data alone can however be tested with a combination of behavioral and neural data, with potential implications for practitioners and policy-makers (Section 3).

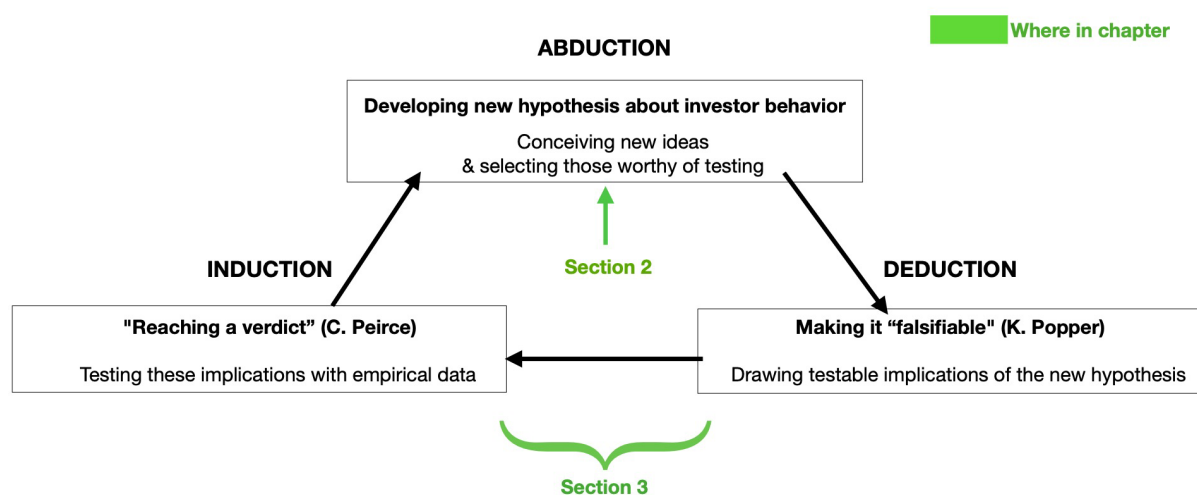


Fig. 1. Neurofinance contributes to behavioral finance at each stage of the scientific process.

2. New Ideas Leading to New Discoveries

Albert Einstein’s experience as a scientist was that new ideas come “suddenly and in a rather intuitive way”.⁵ New ideas can also come from revisiting old ideas, and neuroscience provides finance scholars with new lenses to study investor behavior. In this section, I provide examples of recently proposed hypotheses about investor behavior that emerged from taking what we know at the neurobiological level and apply it to investor decision-making.

⁴ The criterion for deeming an idea “worthy” is that if the idea is true, some otherwise-puzzling fact becomes a “*matter of course*” (Burch, 2018).

⁵ Isaacson, W., 2007. Albert Einstein, Letter to Dr. H. L. Gordon (May 3, 1949).

2.1 The brain as a “Bayesian sampler”

One key neuroscience insight is that the brain is a “*Bayesian sampler*” (Sanborn and Chater, 2016): it did not evolve to interpret explicit probabilities (the exercise was useless or “*ecologically irrelevant*” in primitive environments⁶) but to perform natural sampling tasks that were critical for survival (e.g., foraging for food). As a result, people flounder with elementary probability tasks⁷ but often fare well when the same tasks allow for sampling from probability distributions or, in the language of psychologists, “*learning from experience*” (e.g., Erev et al., 2010). For example, when Gigerenzer et al. (1988) allowed people to directly experience random sampling in Kahneman and Tversky’s (1973) classic “*Engineer-Lawyer problem*”, the celebrated “*base rate neglect*” bias was markedly reduced. In a path-breaking study that pioneered neurofinance as a new scientific field, the notion that through natural sampling, the brain can assess the value of predictors as if it were endowed with a calculator for the celebrated *Black-and-Scholes* formula was advanced (Montague and Berns, 2002).

This representation of the brain allows us to reconcile the heuristics and biases literature, which as stressed above emphasizes cognitive limitations in people, with substantial evidence that human cognition approximates the optimal Bayesian benchmark within important domains of human psychology.⁸ It motivated experimental finance scholars to study how the investor brain, when placed in ecologically relevant contexts, learns about key statistics of return distributions, starting with the first two moments.

2.1.1 Mean-variance analysis in the brain Preuschoff et al. (2006) provided evidence that the brain engages in mean-variance analysis. To establish this, the researchers presented subjects with a deck of ten randomly shuffled cards numbered 1 to 10. Two cards were drawn from the deck without replacement. Task participants were asked to bet whether the second card would be lower than the first. Once the participant posted her bet (for \$1), the two cards were displayed, one after the other. Brain activation was measured through [blood oxygenation level-dependent \(BOLD\) imaging](#) upon revelation of the first card, which represents the point when the participant’s brain should encode the corresponding changes in expected payoff and risk. By design, expected payoff and payoff variance changed each in their own manner as a function of reward probability. Specifically, expected reward increased linearly, and risk displayed a symmetric, inverted U-shaped pattern with a peak at probability 0.5 (Fig. 2A). If expected reward and risk are encoded somewhere in the brain, activation should exhibit the same patterns. The researchers found the ventral striatum—a key node of the [reward system](#)—to encode expected value (Fig. 2B) and [the anterior insula](#) to track variance (Fig. 2C and Preuschoff et al., 2008).

⁶ For example, Brennan and Lo, 2011.

⁷ For example, Tversky and Kahneman (1974) and Grether (1992).

⁸ For example, Körding and Wolpert (2004), Griffiths and Tenenbaum (2006), and Xu and Tenenbaum (2007).

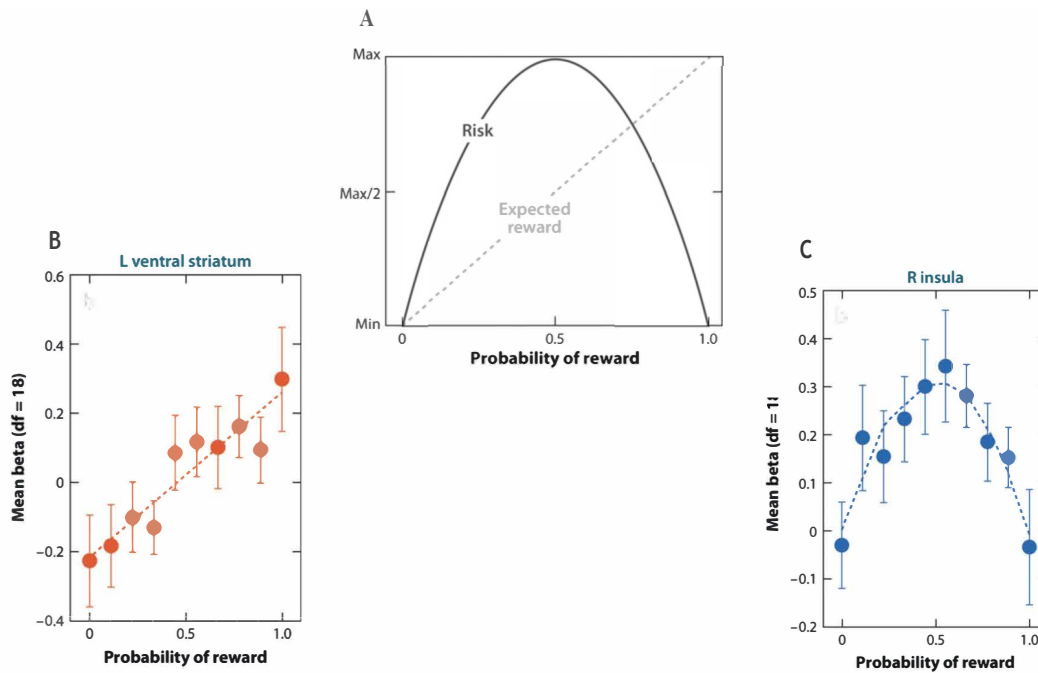


Fig. 2. Mean-variance analysis in the brain. A. Expected payoff and payoff variance as a function of the probability of reward in the card task used in Preuschoff et al. (2006). Brain signals that encode these statistical features should show similar relationships. B. Mean activations across 19 participants (red circles) and 95% confidence intervals (line segments) stratified by level of reward probability in the ventral striatum (a key node of the [reward system](#)). The activation patterns are increasing and linear, which is required if the activations are to be associated with expected reward. C. Mean activation in the right [anterior insula](#) across the participants (blue circles) as a function of reward probability shows the inverted U-shaped pattern needed for it to encode payoff variance. Source: Preuschoff et al. (2006).

Glossary

Amygdala: An almond-shaped cluster of nuclei located deep and medially within the temporal lobes of the brain. As a key node of the [fear system](#), the amygdala plays a primary role in emotional responses related to fear, anxiety, and aggression.

Anterior cingulate cortex: Located along the medial wall of the frontal lobes, immediately posterior to the medial prefrontal cortex, the anterior cingulate cortex is implicated in salience detection and in signaling when action outcomes are misaligned with expectations (i.e., prediction error).

Anterior insula: A region of the cerebral cortex containing so-called “von Economo neurons” (VENs), which play a fundamental role in human emotional awareness and self-consciousness (Allman et al., 2005; Craig, 2009). Its functional roles also include detection of relevant environmental stimuli within the “salience network” (Downar et al., 2000), and signaling risk and potential losses as part of the [fear system](#).

Blood oxygenation level-dependent (BOLD) imaging: The standard technique used to measure brain activity in functional MRI ([fMRI](#)) experiments. Participants in such experiments

are asked to lie down on a narrow table that slides into a metal tube surrounded by a large magnet. A mirror installed inside the tube allows the participant to see a computer screen outside the MRI machine on which the experimental task is displayed. A strong magnetic field is created inside the tube, causing the hemoglobin [the protein that transports oxygen throughout the body] in the participant's blood to emit its own magnetic field that the MRI machine can detect. Hemoglobin without oxygen responds more strongly to a magnetic field than hemoglobin with oxygen. The MRI machine can thus track changes in blood oxygenation in the brain. The neurons in an active region of the brain will require more oxygen than average, and brain activity depletes oxygen levels. Therefore, contrasts in a fMRI scan reflect changes in oxygen levels, which reflect activity in the region.

Dorsomedial prefrontal cortex (dmPFC): A key node of the “*central executive function network*” (Weissman et al., 2006, Mitchell et al., 2009), which is involved in cognitive appraisal and response expression when a task is to be performed. Its key functional roles include detecting response conflict and signaling the need to reallocate attentional resources to other areas. It is also involved in the representation of the mental state of other individuals, which is also known as **Theory of Mind** or ToM (e.g., Frith and Frith, 2003, and Amodio and Frith, 2006).

Fear or “threat detection” system: A collection of serotonergic (serotonin-releasing) neurons that project from the raphe nuclei (a cluster of nuclei found in the brain stem) to several serotonergic targets, including the **anterior insula** and the **amygdala**. The **vmPFC** is also a key node of that pathway, which is activated by the perception of threat (LeDoux, 1996).

Inferior frontal gyrus: A key node of the central executive function network; it is implicated in the adaptation to abrupt contextual changes when previous responses are suboptimal. Its key functional roles include reorienting attention, inhibiting responses to previously rewarded stimuli, and overcoming avoidance of previously punished stimuli.

Locus coeruleus (LC): Subcortical nucleus located in the brain stem and the primary source of the neurotransmitter norepinephrine in the brain. It is believed to be implicated in the signaling of abrupt changes in one's environment (Payzan-LeNestour et al., 2013).

Nucleus accumbens (NAc): Subcortical nucleus located in the ventral striatum and a primary source of the neurotransmitter dopamine in the brain. It is a key node of the **reward system**.

Prior knowledge: The accumulation of learned associations, patterns, generalizations, and abstractions acquired through intentional or incidental learning events and lived experiences.

Model free reinforcement learning: Applied to financial decision-making, this term refers to the use of adaptive expectations to forecast action payoff, which in many contexts results in behaviors well described by the “*win stay lose move*” heuristic. Evidence suggests that it is hardwired in the human brain (Doll et al., 2012; Schultz, 2016).

Reward system: This term refers to the mesolimbic dopaminergic pathway, a collection of dopaminergic (i.e., dopamine-releasing) neurons that project from the ventral tegmental area (VTA) to the ventral striatum and **vmPFC**. Dopaminergic activation of the pathway is accompanied by the perception of reward. As such, the reward system is the neural substrate of desire or “wanting” (Panksepp, 1998; Wyvell and Berridge, 2000).

Nudge: A strategy to shape behavior in a predetermined direction by designing the choice context in a manner that is known to produce a behavioral pattern. Nudges operate either on automatic processing (e.g., making opt-in the default in organ donation systems increases transplantation rates) or, in the case of so-called “**information nudges**” and “booster” (Hertwig, 2017), on people’s deliberation (e.g., warning signs for drivers prime their vigilance system in dangerous situations; providing nutritional charts to consumers shows them the calories hidden in some foods).

Ventromedial prefrontal cortex (vmPFC): A key node of the brain valuation system (Rolls, 2006; Fellows and Farah, 2007). The orbitofrontal cortex (OFC), an area within the vmPFC, is involved in the computation of the value of available options at the time of a decision (Berridge and O’Doherty, 2013). The vmPFC is also implicated in the processing of fear and is critical in the regulation of amygdala activity in humans, thereby playing an important role in the inhibition of emotional responses and the process of self-control.

2.1.2 Learning about unstable payoffs

These findings suggest that the investor’s brain is well equipped to learn about the risk/reward profiles of investment opportunities. However, one could argue that the experimental task used in Preuschoff et al. (2006) is less challenging than what real-world investors face in one key aspect, namely, instability. Specifically, the risk/reward profiles of investment opportunities occasionally change in a sudden manner, for diverse reasons, including shifts in the demands of the goods or services produced, jumps in asset returns, etc. When observing a streak of low payoffs, the investor should determine whether the streak signifies a break with the past, in which case learning about the value of the option at hand should resume from scratch.

Payoff instability complicates investor learning substantially (it involves simultaneously assessing payoff and break probability through “hierarchical Bayesian learning” or similar methods⁹). Payoff instability also substantially increases the incentives of optimally learning about payoffs. Indeed, under stability, i.e., when payoff value does not change or changes slowly such that the past is predictive of the present, the most plausible suboptimal (“boundedly rational”) learning approach, **model free reinforcement learning**, effectively works like Bayesian learning (Masanao, 1987). In contrast, under conditions of instability, Bayesian learning greatly outperforms **model free reinforcement learning** (Payzan-LeNestour & Bossaerts, 2015). Given that instability is a common feature of modern markets, the question of whether investors can learn optimally about unstable payoffs is thus important for finance practitioners.

This motivated an experimental finance study in which task participants were asked to perform a six-armed “restless” bandit task. Each arm represented an investment opportunity for which the expected value jumped regularly over time due to regime shifts affecting the payoff probabilities (Payzan-LeNestour & Bossaerts, 2015). The instructions for the task informed participants that such shifts would occur during the task without telling them *when*, leaving to the participants the task of detecting the shifts to adjust their investment choices accordingly throughout the task. In this context, participants exhibited Bayesian learning despite the complexity. Similar evidence of sophisticated learning in unstable settings was documented in an experimental study where task participants were similarly provided with detailed information about the statistics underlying the task (Meyniel et al., 2015).

⁹ For example, Payzan-LeNestour and Bossaerts (2015).

Context dependence This is just half of the story though. In a follow-up experimental treatment replicating the original task used in Payzan-LeNestour & Bossaerts (2015)’s study, except for that this time, the participants were not informed about the presence of shifts in the outcome probabilities, the participants failed to detect the shifts and ended up resorting to [model free reinforcement](#) to forecast outcomes. One explanation for this reversal is that people consider their environment to be stable by default, which may reflect both the ancestral lifestyle (natural sources of reward shift at low frequency relative to financial returns) and the modus operandi of human perception (generating stability would be among its fundamental properties; see, e.g., Manassi et al., 2017).

The bottom line, therefore, is that the emergence of sophisticated learning in task participants seems to hinge on providing them sufficiently detailed information about the stochastic structure of the task. Similarly, informing investors about the occurrence of shifts in asset returns—by using simulations or natural frequencies, to match the way in which humans have experienced statistical information over most of their history¹⁰—could [nudge](#) them into detecting these shifts.

Too beautiful to be true? However, some may question the effectiveness of such nudging in practice. First, why would the average investor understand that it is worth implementing Bayesian learning to learn about financial returns given that it is commonly *not* worth it? (Recall indeed that the [model free reinforcement learning](#) alternative works equally well in stable environments, and it is “frugal”, i.e., recruiting fewer cognitive resources.¹¹) Second, even if investors understand the importance of learning optimally, understanding it is one thing, implementing it is another. Does the laboratory evidence of sophisticated learning scale to the average investor in the field?

Although a decisive answer to these two questions will require empirical testing, recent work suggests that the answer may be positive for both. First, evidence suggests that people are good intuitive meta-learners; they can intuit what to do in a given context even though they cannot always implement it.¹² Second, evidence suggests that learning about unstable payoffs is ingrained in human neurobiology, inasmuch as the human brain appears to separately signal different aspects of uncertainty, namely, jump likelihood (referred to by neuroscientists as “*unexpected uncertainty*”), the degree of ambiguity/ignorance associated with a given payoff distribution (“*estimation uncertainty*”), and the variance or entropy of the payoff distribution (“*risk*”). For example, in a [fMRI](#) study, the brain of task participants in the foregoing restless bandit task tracked each aspect of uncertainty throughout the task (Fig. 3).¹³ The successful implementation of Bayesian learning hinges on such an exquisite representation of uncertainty.¹⁴

¹⁰ For example, Sedlmeier and Gigerenzer (2001).

¹¹ For example, Gigerenzer et al. (1999).

¹² For example, Boureau et al. (2015), Lieder & Griffiths (2017), and Griffiths et al. (2019).

¹³ For more evidence of neural encoding of these different aspects of uncertainty, see Nassar et al. (2012), McGuire et al. (2014), and Meyniel et al. (2015), among others.

¹⁴ For example, Yu & Dayan (2005) and Payzan-LeNestour & Bossaerts (2011).

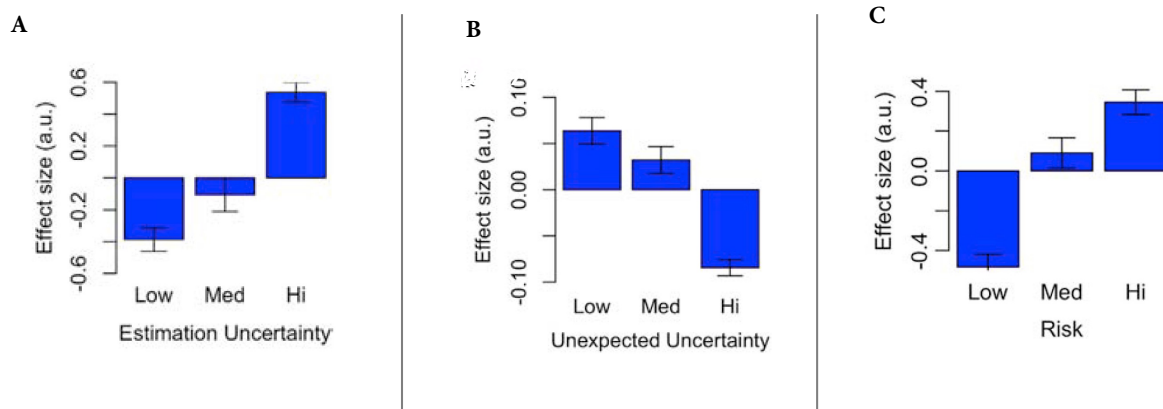


Fig. 3. Separate encoding of different aspects of uncertainty while learning about unstable payoffs. A. Bar plot shows the average effect of low, medium, and high estimation uncertainty in the [anterior cingulate](#) of participants in the decision-making task used in Payzan-LeNestour et al. (2013). Activation in this region increased linearly in estimation uncertainty. B. Bar plot shows the average effect of low, medium, and high unexpected uncertainty in the left middle temporal gyrus (an area located on the temporal lobe in the brain). Activation in this region increased linearly in unexpected uncertainty. A significant response to unexpected uncertainty was also observed in the left [locus coeruleus](#), which is consistent with the proposal that the neurotransmitter norepinephrine plays a key role in human adaptation to abrupt changes (e.g., Yu & Dayan, 2005 and Payzan-LeNestour & Bossaerts, 2011). C. Bar plot shows the average effect of low, medium, and high risk in the [inferior frontal gyrus](#). Activation in this region increased linearly in risk. To generate each plot, trials were sorted according to their level of uncertainty (A: estimation uncertainty; B: unexpected uncertainty; C: risk) into one of three equal-sized bins, which were then fitted to the [BOLD](#) signal. Error bars represent the standard error of the mean (SEM). Source: Payzan-LeNestour et al. (2013).

2.1.3 Hedging in the brain In addition to learning expected returns and payoff variance, investors ought to learn the correlation strengths between different assets given that this is essential for optimal hedging—the process of combining multiple positions in different assets to reduce total risk in a portfolio, as specified by modern portfolio theory (Markowitz, 1952). An influential model in behavioral finance advanced that learning about asset correlations is out of most people’s reach; hence, the optimal hedging strategy shall be replaced in practice by the “*1/N heuristic*”, which disregards environmental structure and uses fixed weights across assets (Benartzi and Thaler, 2001).

This idea is probably true in contexts where information about asset correlation is provided in the form of explicit statistics. But what if the same information is presented in a format that is congenial to the brain? In a [fMRI](#) study, Wunderlich et al. (2011) found that participants in a task requiring minimization of outcome fluctuations were able to learn from experience the correlations between two individual resources and to use this information to make risk-optimal choices. By design, the correlation strength between the two resources changed probabilistically over time, requiring task participants to continuously update their estimate of the current correlation structure. Through repeated observations, the participants learned a strategy that outperformed the *1/N heuristic*. This learning was made possible by the way their [anterior insula](#) tracked the correlation between the two resources throughout the task.

2.1.4 Revisiting the “rhetoric of irrationality” To summarize, research in neurofinance conducted since the early 2000s suggests that our ancestors were endowed with the equivalent of mean-variance and option-pricing calculators in their head and that we are evolutionarily designed to implement computations critical for sound investing, such as detecting abrupt changes in investment opportunities and tracking correlations between assets.

Therefore, altogether, this set of discoveries supports the point made by Gigerenzer and Hoffrage (1995) that “complexity” (the amount of computations needed to perform a given task) is not a reliable predictor of whether humans find a task difficult relative to ecological relevance (the extent to which our ancestors had to cope with similar tasks). Relatedly, it supports the idea proposed by Lopes (1991) among others that the “rhetoric of irrationality” emanating from the heuristics and biases literature became too prominent in the behavioral sciences at some point, leading to overemphasizing human limitations and underaappreciating the key notions of ecological relevance and context dependence. After all, it is perhaps not surprising that when asked to perform ecologically irrelevant tasks, people fare poorly; they are like the white shark put out of the water in the opening quote.¹⁵ The evidence suggests that the same people can perform well in contexts of repeated sampling that are directly relevant to financial investing.

However, this does not imply that investors will always behave optimally in ecologically relevant contexts. As we saw above, the quality of their learning about asset returns is highly context dependent in the sense that seemingly minor changes in their environment, such as whether they are given some piece of information about the statistics underlying asset returns, can have dramatic effects on the quality of their learning. The following discusses other key aspects of this context dependence, and how modifying specific contextual elements of the investing landscape can help improve investor decision-making.

2.2 Efficient Coding

Researchers recently studied risk-taking through the lens of “*efficient coding theory*”, the neuroscientific framework commonly used to study how the brain records information about the external world. The basic idea is that the brain is designed to communicate information in a way that economizes on its limited resources (the fact that we have a finite number of neurons, and each neuron has a finite number of spike outputs).¹⁶

2.2.1 Adaptive normalization and the “after-effect” in investor perception of volatility

Two aspects of efficient coding are particularly relevant for financial decision-making. The first aspect is that the perception of a stimulus depends not on its absolute intensity but rather on its intensity relative to the recent past and neighboring stimuli. As a result, perception is sensitive to contrasts or deviations from the mean (Glimcher, 2014), giving rise to a number of effects, including the celebrated “*after-effect*”, in which prolonged exposure to a given stimulus level systematically creates the illusion of an opposite stimulus. For example, after prolonged viewing of the downward flow of a waterfall, static rocks to the side appear to ooze upward (Barlow and Hill, 1963). This would reflect the “normalized” response of the neurons coding for downward motion (their activity is below baseline, as if they were “fatigued” following the prolonged stimulation) relative to the neurons coding for upward motion (which feature baseline activity). Similarly, after prolonged viewing of a feminine face, a gender-neutral

¹⁵ An extreme example of ecologically irrelevant task is a “thought experiment” in which participants are asked for hypothetical choices in imagined situations (e.g., Tversky and Kahneman, 1974).

¹⁶ For example, Barlow (1972), Laughlin (1981), and Glimcher (2014).

(androgynous) face looks masculine for the same reason (diminished response of neurons coding for the feminine property relative to the neurons coding for the masculine property).

After-effects have been observed for all types of stimuli. Given their ubiquity, and given that financial markets feature volatility “regimes” (prolonged periods of extreme—very high or very low—volatility), it is natural to apply the concept of after-effect to investor perception of volatility. To be more specific, when experiencing neutral volatility levels (e.g., 1% appears to be medium for equity volatility) after prolonged exposure to very high volatility (e.g., 4%), investors may perceive neutral volatility levels as lower than actual levels (Payzan-LeNestour et al., 2016; Payzan-LeNestour et al., 2022b). Evidence suggests that a volatility after-effect of this type exists and distorts the volatility perception of both the average individual in the laboratory and the marginal trader himself, causing distortions in the VIX that constitute an exploitable arbitrage opportunity (Payzan-LeNestour et al., 2022a; Payzan-LeNestour et al., 2022b).

2.2.2 “Outlier blindness” and tail risk neglect Allocating most of the brain’s resources to represent the outcomes frequently encountered at the expense of the unlikely outcomes is an effective method to economize on these resources.¹⁷ A side effect of this strategy, however, is that the brain tends to perceive tail events—the so-called “*black swans*” (Taleb, 2004)—as less extreme than they actually are (Payzan-LeNestour & Woodford, 2022). Applied to financial decision-making, this idea implies that we should expect to see underestimation of tail risk in financial investing, not because extreme events are not included in the sample on the basis of which risks are estimated but because they are initially perceived to be less extreme than they actually are. For example, traders could initially underappreciate the size of volatility “jumps” (consistent with recent findings, e.g., Lochstoer and Muir, 2020). Similarly, the importance of macroeconomic shocks, such as the *Long-Term Capital Management* and *Global Financial Crises*, could have initially been underappreciated.

2.2.3 Context dependence in risk attitude Frydman and Jin (2022) provided experimental evidence that when the frequency of large payoffs is increased, people perceive the upside of a risky lottery more accurately and take greater risk, as a result of neural efficient coding. This suggests that risk-taking in a given individual depends on the payoff distribution to which the individual has recently been exposed.

2.3 Craving for financial returns and the “cheap call selling anomaly”

Craving is the term used by neuroscientists to refer to an intense desire for a given reward irrespective of its expected value.¹⁸ The phenomenon is believed to reflect so-called “Pavlovian influences”: following the repeated exposure to the reward associated with a given reward cue, the [reward system](#) responds more powerfully to the prospect of getting the reward every time the cue is presented in the form of dopamine boosts. Neurochemical “reward signals”, these are automatically triggered by the reward system, i.e., they are not under cognitive control. Hence, the decision-maker can find himself “wanting” a reward cue despite knowing that its expected value is negative.¹⁹ For example, when craving for a food item while in a satiated state, we find ourselves to be lured into eating the food, although we do not expect to particularly enjoy the experience (“I’m truly full!”). Kent Berridge called this phenomenon “wanting without liking” or “irrational wanting”.

¹⁷ For example, Tobler et al. (2005) and Wei and Stocker (2017).

¹⁸ See Berridge & O’Doherty (2013) for a survey of the literature.

¹⁹ For example, Wyvell & Berridge (2000) and Dickinson & Balleine (2002).

Ample neuroscientific evidence indicates that craving occurs for different types of rewards (food, drug, etc.), and the reward system processes the different types similarly (using a “*common currency*”²⁰). This motivated researchers to apply the concept of craving to financial investing. The basic idea is that repeatedly obtaining good returns with a given investment choice could lead the investor’s **reward system** to bias the investor toward choosing the investment, even if its expected value is negative. Take for example the repeated play of a gamble delivering \$2 with a probability of 0.9 and a loss of \$100 otherwise, and assume the agent knows these underlying statistics (there is no learning involved). The gamble’s expected value ($.9 \times 2 - .1 \times 100$) is clearly negative. But given how the brain works, through repeated exposure to the gamble, the agent’s brain comes to automatically associate the gamble with its good outcome. Hence, the agent may find himself taking it despite knowing the expected value is negative—a gambling bias (Payzan-LeNestour & Doran, 2022).

New market anomaly Empirical tests of this idea led to the discovery that selling call options priced below \$1 (so called “cheap calls”) has consistently delivered negative long-term returns and negative skew (Payzan-LeNestour & Doran, 2022). An anomaly viewed from the prevailing body of knowledge in finance, this finding is, however, a “*matter of course*”²¹ when this body of knowledge is augmented with the idea of investor craving for financial returns. Indeed, relative to trading other options, cheap call selling yields a good outcome with very high probability. As such, it is inherently tempting (“craving provoking”), like the gamble that the agent comes to want despite the negative expected value in the foregoing example.

2.4 Emotion as information

2.4.1 “Somatic Marker Hypothesis”

Behavioral scientists have long embraced the view that emotions are not only unnecessary but disruptive (Shiraev, 2010). Behavioral finance research is a case in point. Two landmark articles painted a dualist picture of the mind whereby emotions such as fear and greed (René Descartes’ “*animal spirits*”) trump rational decision-making (Kahneman, 2003; Camerer et al., 2005). This narrative is commonly used to explain departures from the “Efficient Markets Hypothesis” (Samuelson, 1965; Fama, 1970). At odds with such a Cartesian view of the mind, the current consensus in neuroscience is that emotion is *central* to rational decision-making (Damasio, 1994), as encapsulated in the celebrated “*Somatic Marker Hypothesis*” (Bechara et al., 1997) and congenial “*affect as information*” idea (Loewenstein et al., 2001). One key insight is that emotion’s primary role is to motivate us for decisions linked to survival (Rolls, 1992; Panksepp, 1998) and to help us focus: being scared by a potential threat and being positively aroused by a potential reward both help focus (Hansen et al., 2008; Berke, 2018).

Techniques to measure brain activity, such as **BOLD imaging**, have allowed researchers to show the central involvement of emotion not only in choices linked to immediate survival but also in more abstract choices, including those involved in financial risk-taking. For example, financial options involving uncertain large gains have consistently been found to elicit “anticipatory reward signals” triggered by the **NAC** at the time of a decision, while potential large losses elicit “anticipatory loss signals” generated by the **amygdala** and **anterior insula** (Knutson & Bossaerts, 2007; Bossaerts, 2009; Wu, 2012; d’Acromont & Bossaerts, 2016). Rather than being epiphenomenal, these anticipatory signals *guide* choice. Specifically, anticipatory **NAC** activity at the time of a decision appears to be a reliable predictor of financial

²⁰ For example, Levy & Glimcher (2012).

²¹ The term should be understood in the epistemological sense of Footnote 4.

risk-taking—e.g., lower activity predicting lower buying during bubbles in Smith et al., 2014, higher activity predicting higher risk-taking in the investment task of Kuhnen & Knutson, 2005. Likewise, increased anticipatory [anterior insula](#) activity at the time of a decision predicts reduced financial risk-taking both in lab settings (e.g., Kuhnen & Knutson, 2005; Knutson et al., 2007) and real-life stock trading (Häusler et al., 2018).

2.4.2 Optimizing emotions In addition to their crucial role in guiding decisions, emotions appear to be essential for learning in markets. For example, the neural mean-variance signals mentioned in Section 2.1.1 are believed to be emotional rather than cognitive. Evidence also suggests that professional traders with higher [anterior insula](#) activity upon critical market events are more successful at appraising these events as a result of being more attuned to their gut feelings in those moments (Lo & Repin, 2002; Steenbarger, 2002; Critchley et al., 2004; Fenton-O’Creedy et al., 2011; d’Acremont & Bossaerts, 2016; Kandasamy et al., 2016). The bottom line is that harnessing emotion is central to sound financial investing (Bossaerts, 2019). This however does not imply that ‘more emotion’ is always optimal. Evidence suggests that chronic arousal hampers trading performance (Lo and Repin, 2002), and inhibiting prepotent ‘flight’ responses to facilitate sticking to one’s ongoing strategy is essential during aberrational events such as “flash crash” (d’Acremont & Bossaerts, 2016). One key practical implication for investors, therefore, is that rather than systematically minimizing—in the language of psychologists, “downregulating”—their level of emotional arousal, investors ought to *optimize* it, which opens fascinating avenues for applied research in behavioral finance (Section 4).

3. Using Neural Data to Test Theories of Investor Behavior

3.1 “Realization utility” theory of trading

Neural data can be helpful in testing theories of investor behavior that are difficult to test with behavioral data alone. “*Realization utility theory*” (Shefrin and Statman, 1985; Barberis and Xiong, 2012; Ingersoll and Jin, 2013) is one example. It posits that, in addition to deriving utility from consumption, investors derive utility directly from realizing gains and losses on the sale of risky assets that they own. For example, if an investor realizes a gain (e.g., by buying a stock at \$20 and selling it at \$40), he receives a positive burst of utility proportional to the capital gain. In contrast, if he realizes a loss, he receives a negative burst of utility proportional to the realized loss. Under plausible conditions, realization utility will lead investors to exhibit the disposition effect (Barberis and Xiong, 2012).

Realization utility is difficult to test with behavioral data alone because its behavioral predictions are similar, on many dimensions, to those of other theories (Kaustia, 2010). However, the distinctive implications of the theory *at the neural level* can be tested. For example, realization utility theory implies that realizing a capital gain (/ a loss) increases (/ decreases) activity in a part of the [reward system](#) called the ventral striatum, which is known to signal utility changes. Frydman et al. (2014) provided evidence for the neural predictions of the theory in a laboratory setting.

3.2 Neurofoundations for market “bubbles”

A bubble is identified when an asset is traded at prices well above its intrinsic fundamental value. For example, housing prices increased and crashed in many countries from 2000–2008, causing long-lasting macroeconomic disruptions. The classical explanation for why bubbles occur, “irrational exuberance”, encompasses two potential (not mutually exclusive) mechanisms: faulty beliefs on one hand (investors being unaware that the market is

overpriced),²² and the overstimulation of the [reward system](#) described in Section 2.3 on the other hand, which underlies what has been referred to as market “greed.”²³ From a macroeconomic perspective, it is important to distinguish between the faulty cognition and greed factors as they require different interventions. Neural data can be useful in this regard. For example, Smith et al.’s (2014) experimental finding that larger anticipatory reward signals are associated with increased market participation during bubbles points to the greed factor being among the root causes of bubbles.

Neural data allowed researchers to uncover other potentially relevant mechanisms that are difficult to investigate with behavioral data alone. For example, evidence points to the desire to avert the disappointment caused by forgoing potentially large gains as being a powerful motivator behind market participation during a bubble (Lohrenz et al., 2007). Evidence also suggests that market bubbles are more likely to form when investors try to infer the intentions of other market participants (De Martino et al., 2013). To assess this hypothesis, investigators scanned the brain of participants in experimental markets during their trading. They found neural signals in the participants’ [dmPFC](#), a region well known to be involved in [Theory of Mind](#) computations, to modulate the computation of trading value in the participants’ [vmPFC](#), stimulating the formation of financial bubbles in their market experiment.

4. Future Directions

Nudging After providing evidence that the disposition effect originates from the realization utility of selling winner stocks (Section 3.1), researchers turned toward deriving practical implications for investors. Evidence suggests that decreasing the salience of capital gains on investors’ online trading screens may help reduce the disposition effect in investors by nudging them into focusing less on the capital gain aspect of their investments and by decreasing the realization utility of selling winner stocks (e.g., Frydman & Rangel, 2014, and Frydman & Wang, 2020). This is an example of behavioral interventions—the so-called “[nudges](#)”—aimed to debias investor behavior. The strategy consisting of nudging investors into paying attention to the instability underlying asset returns so they can learn about unstable payoffs (cf. Section 2.1.2) is an example of a different category of nudges, namely, [information nudges](#).

Neuroforecasting We saw in Section 2.4.1 that neural activity can be used to predict individual choice. Recently, researchers have begun to explore whether it might also inform forecasts of [aggregate choice](#). Although some relegate such an idea to the realm of science fiction, recent findings suggest that it can work. For example, by averaging brain activity in teenagers exposed to music clips culled from an internet site while undergoing [fMRI](#), researchers found that the sample’s averaged brain activity in the [NAc](#) in response to these songs could forecast aggregate song downloads 2 years later (Berns & Moore, 2012). Strikingly, averaged explicit ratings of liking collected from the laboratory sample did not forecast aggregate song downloads. These findings along with others in the same vein led researchers to argue for a “*brain-as-predictor*” approach, in which brain activity might offer unique information capable of improving forecasts of aggregate choice. Evidence suggests that in some cases, the contribution of neural information may supersede even that afforded by individual choice itself or behavioral measures, such as self-report ratings (Knutson & Genevsky, 2018).

²² For example, Shiller (2000), Shefrin (2009), and Ubel (2009).

²³ For example, Kindleberger and Alibert (2005) and Lo (2017).

Optimizing emotion in real time through biofeedback After establishing that investors' level of emotional arousal ought to be optimized rather than minimized (Section 2.4.2), researchers could turn toward developing machine learning algorithms that would allow us to predict the trading performance of a trader from the way he emotionally reacts to market events. This would require simultaneously collecting in real-time psychophysiological markers of emotion and trading performance data and then using machine learning techniques to learn associations between the two. Ultimately, this learning could be concretized into a live decision support system based on biofeedback that investors could use in their workplace to improve their decision-making in real time.

Thus, although future research in neurofinance may move in diverse directions, the path taken in recent research has opened up enticing opportunities for improving investor decision-making in the field.

References

- Allman, J.M., Watson, K.K., Tetreault, N.A., Hakeem, A.Y., 2005. Intuition and autism: A possible role for von Economo neurons. *Trends Cogn. Sci.* 9, 367–373.
- Amodio, D.M., Frith, C.D., 2006. Meeting of minds: The medial frontal cortex and social cognition. *Nat. Rev. Neurosci.* 7, 268–277.
- Ang, A., Timmermann, A.G., 2011. Regime changes and financial markets. *Netspar. Disc. Papers.* 6, 2011-2068.
- Barberis, N., Xiong, W., 2012. Realization utility. *J. Financ. Econ.* 104, 251–271.
- Bargh, J.A., Chartrand, T.L., 1999. The unbearable automaticity of being. *Am. Psychol.* 54, 462-479.
- Barlow, H.B., 1972. Dark and light adaptation: Psychophysics. in Jameson, D., Hurvich, L.M., (eds.) *Handbook of Sensory Physiology*. Springer-Verlag, New York.
- Barlow, H.B., Hill, R.M., 1963. Evidence for a physiological explanation of the waterfall phenomenon and figural after-effects, *Nature.* 200, 1345–1347.
- Bechara, A., Damasio, H., Tranel, D., Damasio, A.R., 1997. Deciding advantageously before knowing the advantageous strategy. *Science.* 275, 1293-1295.
- Benartzi, S., Thaler, R.H., 2001. Naive diversification strategies in defined contribution saving plans. *Am. Econ. Rev.* 91, 79–98.
- Berke, J.D., 2018. What does dopamine mean? *Nat. Neurosci.* 21, 787-93.
- Berns, G. S., & Moore, S. E., 2012. A neural predictor of cultural popularity. *J. Consum. Psychol.* 22, 154–160.
- Berridge, K.C., O'Doherty, J.P., 2013. *Neuroeconomics: Decision making and the brain*, chapter from experienced utility to decision utility. Academic Press, San Diego. 335–351.
- Bossaerts, P., 2009. What decision neuroscience teaches us about financial decision making. *Annu. Rev. Financ. Econ.* 1, 383-404.
- Boureau, Y., Sokol-Hessner, P., Daw, N., 2015. Deciding how to decide: Self-control and meta-decision making. *Trends Cogn. Sci.* 19,700-710.
- Brennan, T.J., Lo, A., 2011. The origin of behavior. *Q. J. Finance* 1, 55–108.
- Brown, G.R., Richerson, P.J., 2014. Applying evolutionary theory to human behaviour: Past differences and current debates. *J. Bioecon.* 16, 105–128.
- Burch, R., 2018. Charles Sanders Peirce. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University.

- Camerer, C., Loewenstein, G., Prelec, D., 2005. Neuroeconomics: how neuroscience can inform economics. *J. Econ. Lit.* 43, 9-64.
- Caplan, B., 2007. *The myth of the rational voter: Why democracies choose bad policies.* Princeton University Press, New Jersey.
- Coates, J., 2012. *The hour between dog and wolf: How risk taking transforms us, body and mind.* Penguin Books, London.
- Craig, A.D., 2009. How do you feel — now? The anterior insula and human awareness. *Nat. Rev. Neurosci.* 10, 59–70.
- Critchley, H.D., Wiens, S., Rotshtein, P., Ohman, A., Dolan, R.J., 2004. Neural systems supporting interoceptive awareness. *Nat. Neurosci.* 7, 189-195.
- d'Acremont, M., Bossaerts, P., 2016. Human reaction to leptokurtic noise and its neurobiological foundations. *Cereb. Cortex.* 26, 1818-30.
- Damásio, A.R., 1994. *Descartes' Error: Emotion, reason, and the human brain.* Avon Books, New York.
- De Martino, B., O'Doherty, J.P., Ray, D., Bossaerts, P., Camerer, C., 2013. In the mind of the market: Theory of mind biases value computation during financial bubbles. *Neuron.* 79, 1222-1231.
- Dickinson, A., Balleine, B., 2002. *Stevens' Handbook of Experimental Psychology, Chapter: The role of learning in the operation of motivational systems.* John Wiley and Sons, New York. 497–533.
- Doll, B.B., Simon, D.A., Daw, N.D., 2012. The ubiquity of model-based reinforcement learning. *Curr. Opin. Neurobiol.* 22, 1075–81.
- Downar, J., Crawley, A., Mikulis, D., Davis, K.D., 2000. A multimodal cortical network for the detection of changes in the sensory environment. *Nat. Neurosci.* 3, 277–283.
- Epstein W.R., 1960. Perceptual set as an artifact of recency. *Am. J. Psychol.* 73, 214–228.
- Erev, I., Ert, E., Roth, A.E., Haruvy, E., Herzog, S.M., Hau, R., Hertwig, R., Stewart, T., West, R., Lebiere, C., 2010. A choice prediction competition: choices from experience and from description. *J. Behav. Decis. Mak.* 23, 15–47.
- Fama, E.F., 1970. Efficient capital markets: A review of theory and empirical work. *J. Finance.* 25, 383-417.
- Fellows, L.K., Farah, M.J., 2007. The role of ventromedial prefrontal cortex in decision making: Judgment under uncertainty or judgment per se? *Cereb. Cortex.* 17, 2669–2674.
- Fenton-O'Creevy, M., Soane, E., Nicholson, N., Willman, P., 2011. Thinking, feeling and deciding: The influence of emotions on the decision making and performance of traders. *J. Organ. Behav.* 32, 1044-1061.
- Frith, U., Frith, C.D., 2003. Development and neurophysiology of mentalizing. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 358, 459–473.
- Frydman, C., Barberis, N., Camerer, C., Bossaerts, P., Rangel, A., 2014. Using neural data to test a theory of investor behavior: An application to realization utility, *J. Finance.* 69, 907-946.
- Frydman, C., Jin, L.J., 2022. Efficient coding and risky choice. *Q. J. Econ.* 1, 161-213.
- Frydman, C., Rangel, A., 2014. Debiasing the disposition effect by reducing the saliency of information about a stock's purchase price. *J Econ Behav Organ.* 107, 541-552.
- Frydman, C., Wang, B., 2020. The Impact of salience on investor behavior: Evidence from a natural experiment, *J. Finance.* 75, 229-276.
- Gao, P., Da, Z., Engelberg, J., 2011. In search of attention. *J. Finance* 66,1466–1491.
- Garcia, J., Koelling, R.A., 1966. Relation of cue to consequence in avoidance learning. *Psychon. Sci.* 4, 123–124.

- Gigerenzer, G., Hell, W., Blank, H., 1988. Presentation and content: The use of base rates as a continuous variable. *J Exp Psychol Hum Percept Perform.* 14, 513–525.
- Gigerenzer, G., Hoffrage, U., 1995. How to improve bayesian reasoning without instructions: Frequency formats. *Psychol. Rev.* 102, 684–704.
- Gigerenzer, G., Todd, P., The ABC Research Group, 1999. Simple heuristics that make us smart. Oxford University Press, Oxford.
- Glimcher, P., 2014. Understanding the hows and whys of decision-making: From expected utility to divisive normalization. *Cold Spring Harb. Symp. Quant. Biol.* 79, 169–176.
- Grether, D.M., 1992. Testing Bayes rule as a descriptive model: The representativeness heuristic. *Q. J. Econ.* 95, 537–557.
- Griffiths, T.L., et al. 2019. Doing more with less: Meta-reasoning and meta-learning in humans and machines. *Curr. Opin. Behav. Sci.* 29, 24-30.
- Griffiths, T.L., Tenenbaum J.B., 2006. Optimal predictions in everyday cognition. *Psychol. Sci.* 17, 767–773.
- Hansen, A.L., Johnsen, B., Thayer, J., 2008. Relationship between heart rate variability and cognitive function during threat of shock. *Anxiety stress coping.* 22, 77-89.
- Häusler, A.N., Kuhnen, C.M., Rudorf, S., Weber, B., 2018. Preferences and beliefs about financial risk taking mediate the association between anterior insula activation and self-reported real-life stock trading. *Sci Rep* 8, 11207.
- Hertwig, R., 2017. When to consider boosting: Some rules for policy-makers. *Behav. Public Policy.* 1, 143-161.
- Hirshleifer, D., Lim, S.S., Teoh, S.H., 2009. Driven to distraction: Extraneous events and underreaction to earnings news. *J. Finance.* 64, 2289–2325.
- Hoffrage, U., Lindsey, S., Hertwig, R., Gigerenzer, G., 2000. Communicating statistical information. *Science.* 290, 2261–2262.
- Huberman, G., Regev, T., 2001. Contagious speculation and a cure for cancer: A nonevent that made stock prices soar. *J. Finance.* 56, 387–96.
- Ingersoll, J., Jin, L.J., 2013. Realization utility with reference-dependent preferences. *Rev. Financ. Stud.* 26, 723–767.
- Jones, C., Kaul, G., Lipson, M. 1994. Transactions, volume, and volatility. *Rev. Financ. Stud.* 7, 631-651.
- Kahneman, D., 2003. A psychological perspective on economics. *Am. Econ. Rev.* 93, 162-168.
- Kahneman, D. and Tversky, A., 1973. On the psychology of prediction. *Psychol. Rev.*, 80:237–251.
- Kandasamy, N., et al. 2016. Interoceptive ability predicts survival on a london trading floor. *Sci. Rep.* 6, 32986.
- Kaustia, M., 2010. Prospect theory and the disposition effect. *J. Financ. Quant. Anal.* 45, 1–36.
- Kindleberger, C.P., Alibert, R.Z., 2005. Manias, panics, and crashes. John Wiley and Sons, New York.
- Knutson, B., Bossaerts, P., 2007. Neural antecedents of financial decisions. *J. Neurosci.* 27, 8174-8177.
- Knutson, B., Rick, S., Wimmer, G.E., Prelec, D., Loewenstein, G., 2007. Neural predictors of purchases. *Neuron.* 53,147-156.
- Knutson, B., Genevsky, A., 2018. Neuroforecasting Aggregate Choice. *Curr Dir Psychol Sci.* 27(2), 110-115.
- Körding, K.P., Wolpert., D.M., 2004. Bayesian integration in sensorimotor learning. *Nature.* 427, 244–247.

- Kuhnen, C.M., Knutson, B., 2005. The neural basis of financial risk taking. *Neuron*. 47, 763-770.
- Laughlin, S.B., 1981. A simple coding procedure enhances a neuron's information capacity. *Z Naturforsch C Biosci*. 36, 910–912.
- LeDoux, J.E., 1996. *The emotional brain*. Simon and Schuster, New York.
- Levy, D.J., Glimcher, P.W., 2012. The root of all value: A neural common currency for choice. *Curr. Opin. Neurobiol*. 22, 1027–1038.
- Lieder, F., Griffiths, T.L., 2017. Strategy selection as rational metareasoning. *Psychol. Rev.* 124, 762-794.
- Lo, A., 2017. *Adaptive markets: Financial evolution at the speed of thought*. Princeton University Press, New Jersey.
- Lo, A., Repin, D.V., 2002. The psychophysiology of real-time financial risk processing. *J. Cogn. Neurosci*. 14, 323-39.
- Lochstoer, L.A., Muir, T., 2022. Volatility expectations and returns. *J. Finance*. <https://doi.org/10.1111/jofi.13120>
- Loewenstein, G.F., Weber, E.U., Hsee, C.K., Welch, N., 2001. Risk as feelings. *Psychol. Bull.* 127, 267–286.
- Lohrenz, T., McCabe, K., Camerer, C.F., Montague, P.R., 2007. Neural signature of fictive learning signals in a sequential investment task. *Proc. Natl. Acad. Sci. USA*. 104, 9493–9498.
- Lopes, L.L., 1991. The Rhetoric of Irrationality. *Theory Psychol.* 1, 65-82.
- Manassi, M., Liberman, A., Chaney, W., Whitney, D., 2017. The perceived stability of scenes: serial dependence in ensemble representations. *Sci. Rep.* 7, 1971.
- Mandelbrot, B.B., 1957. *Fractales, Hasard et Finance*. Flammarion, Paris.
- Markowitz, H., 1952. Portfolio selection. *J. Finance* 7, 77–91.
- Masanao, A., 1987. *State space modeling of time series*. Springer-Verlag, Berlin.
- McGuire, J.T., Nassar, M.R., Gold, J.I., Kable, J.W., 2014. Functionally dissociable influences on learning rate in a dynamic environment. *Neuron*. 84, 870–881.
- Meyniel, F., Schlunegger, D., Dehaene, S., 2015. The sense of confidence during probabilistic learning: A normative account. *PLoS Comput. Biol.* 11, e1004305.
- Mitchell, D.G., et al. 2009. Adapting to dynamic stimulus-response values: Differential contributions of inferior frontal, dorsomedial, and dorsolateral regions of prefrontal cortex to decision making. *J. Neurosci.* 29, 10827-10834.
- Montague, P.R., Gregory, S.B., 2002. Neural economics and the biological substrates of valuation. *Neuron*. 36, 265-84.
- Nassar, M.R., et al. 2012. Rational regulation of learning dynamics by pupil-linked arousal systems. *Nat. Neurosci.* 15, 1040–1046.
- Panksepp, J., 1998. *Affective neuroscience: The foundations of human and animal emotions*. Oxford University Press, USA.
- Payzan-LeNestour, E., Balleine, B.W., Berrada, T., Pearson, J., 2016. Variance after-effects distort risk perception in humans. *Curr. Biol.* 26, 1500-1504.
- Payzan-LeNestour, E., Bossaerts, P., 2011. Risk, unexpected uncertainty, and estimation uncertainty: Bayesian learning in unstable settings. *PLoS Comput. Biol.* 7, 1001048.
- Payzan-LeNestour, E., Bossaerts, P., 2015. Learning about unstable publicly unobservable payoffs. *Rev. Financ. Stud.* 28, 1874–1913.
- Payzan-LeNestour, E., Doran, J., 2022. Craving for money? Empirical evidence from the laboratory and the field. SSRN Working Paper.
- Payzan-LeNestour, E., Doran, J., Pradier, L., Putnins, T., 2022a. Harnessing neuroscientific insights to generate alpha. *Financ. Anal. J.* 78, 79-95.

- Payzan-LeNestour, E., Dunne, S., Bossaerts, P., O'Doherty J.P., 2013. The neural representation of unexpected uncertainty during value-based decision making. *Neuron*. 79, 191-201.
- Payzan-LeNestour, E., Pradier, L., Doran, J., Nave, G., Balleine, B.W., 2021. Impact of ambient sound on risk perception in humans: Neuroeconomic investigations. *Sci. Rep.* 11, 5392.
- Payzan-LeNestour, E., Pradier, L., Putnins, T., 2022b. The “waterfall illusion” in the perception of risk. Working Paper.
- Payzan-LeNestour, E., Woodford, M., 2022. Outlier blindness: A neurobiological foundation for neglect of financial risk. *J. Financ. Econ.* 143, 1316-1343.
- Peirce, C.S., de Waal, C., (eds.) 2014. *Illustrations of the Logic of Science 1877*. Open Court, Chicago.
- Popper, K.R., 1959. *The Logic of Scientific Discovery*. 2005 Routledge eds., London.
- Preuschoff, K., Bossaerts, P., Quartz, S.R., 2006. Neural differentiation of expected reward and risk in human subcortical structures. *Neuron*. 51,381–90.
- Preuschoff, K., Quartz, S.R., Bossaerts, P., 2008. Human insula activation reflects risk prediction errors as well as risk. *J. Neurosci.* 28, 2745–2752.
- Rolls, E.T., 1992. Neurophysiology and functions of the primate amygdala. in Aggleton, J.P., *The Amygdala* (eds). Wiley-Liss, New York. 143-165.
- Rolls, E.T., 2006. The neurophysiology and functions of the orbitofrontal cortex. in Zald, D.H., Rauch, S., *The orbitofrontal cortex* (eds). Oxford University Press, New York. 95–125.
- Samuelson, P., 1965. Proof that properly anticipated prices fluctuate randomly. *Ind. Manag. Rev.* 6, 41-49.
- Sanborn, A.N., Chater, N., 2016. Bayesian brains without probabilities. *Trends Cogn. Sci.*, 20, 883-893
- Schultz, W., 2016. Dopamine reward prediction-error signalling: A two-component response. *Nat. Rev. Neurosci.* 17, 183–195.
- Sedlmeier, P., Gigerenzer, G., 2001. Teaching bayesian reasoning in less than two hours. *J. Exp. Psychol. Gen.* 130, 380–400.
- Shefrin, H., 2009. *Insights into the Global Financial Crisis*. Number 5. Research Foundation of CFA Institute.
- Shefrin, H., Statman, Meir., 1985, The disposition to sell winners too early and ride losers too long: Theory and evidence, *J. Finance.* 40, 777–790.
- Shirayev, E., 2010. *A History of Psychology: A Global Perspective*. Sage, Los Angeles.
- Shiller R.J., 2000. *Irrational Exuberance*. Princeton University Press, New Jersey.
- Smith, A., Lohrenz, T., King, J., Montague, P.R., Camerer, C.F., 2014. Irrational exuberance and neural crash warning signals during endogenous experimental market bubbles. *Proc. Natl Acad. Sci.* 111, 10503-10508.
- Steenbarger, B.N., 2002. *The psychology of trading: Tools and techniques for minding the markets*. John Wiley and Sons. New Jersey.
- Taleb, N.N., 2004. *Foiled by randomness: The hidden role of chance in life and in the markets*. Penguin Books, London.
- Thaler R.H., 2016. *Misbehaving: The making of behavioral economics*. W.W. Norton, New York.
- Tobler, P.N., Fiorillo, C.D., Schultz, W., 2005. Adaptive coding of reward value by dopamine neurons. *Science.* 307, 1642–1645
- Tversky, A., Kahneman, D., 1974. Judgment under uncertainty: Heuristics and biases: Biases in judgments reveal some heuristics of thinking under uncertainty. *Science.* 185, 1124-1131.

- Ubel, P., 2009. Human nature and the financial crisis. *Forbes*. Printed at https://www.forbes.com/2009/02/20/behavioral-economics-mortgage-opinions-contributors_financial_crisis.html?sh=67022557281d
- Wei, X.X., Stocker, A.A., 2017. Lawful relation between perceptual bias and discriminability. *Proc. Natl Acad. Sci.* 114, 10244–10249
- Weissman, D.H., Roberts, K.C., Visscher, K.M., Woldorff, M.G., 2006. The neural bases of momentary lapses in attention. *Nat. Neurosci.* 9, 971-978.
- Wunderlich, K., Symmonds, M., Bossaerts, P., Dolan, R.J., 2011. Hedging your bets by learning reward correlations in the human brain. *Neuron.* 71, 1141–1152.
- Wyvell C.L., Berridge, K.C., 2000. Intra-accumbens amphetamine increases the conditioned incentive salience of sucrose reward: Enhancement of reward wanting without enhanced liking or response reinforcement. *J. Neurosci.* 20, 8122–8130.
- Wu, C.C., Sacchet, M.D., Knutson, B., 2012. Toward an affective neuroscience account of financial risk taking. *Front. Neurosci.* 6, 159.
- Xu, F., Tenenbaum, J.B., 2007. Word learning as bayesian inference. *Psychol. Rev.* 114, 245–272.
- Yu, A.J., Dayan, P., 2005. Uncertainty, neuromodulation, and attention, *Neuron.* 46, 681-692.