

Outlier Blindness: A Neurobiological Foundation for Neglect of Financial Risk*

Elise Payzan-LeNestour and Michael Woodford[†]

April 20, 2021

Abstract

How do people record information about the outcomes they observe in their environment? Building on a well-established neuroscientific framework, we propose a model in which people are hampered in their perception of outcomes that they expect to seldom encounter. We provide experimental evidence for such ‘outlier blindness’ and discuss how it provides a microfoundation for neglected tail risk by investors in financial markets. (JEL: C91, D87, G41)

Keywords: Tail risk, Imprecise perception, Efficient coding, Adaptation, Decision making under uncertainty, Behavioral finance, Neuroeconomics, Experiments

*We are indebted to Patrick Bolton, Peter Bossaerts, Brice Corgnet, Peter Dayan, Samuel Gershman, Philippe Gouillou, Lawrence Jin, Ronald Masulis, Carsten Murawski, Amos Nadler, Antonio Rangel, Christian Ruff, Philippe Tobler, seminar participants at Caltech, Melbourne University, the Behavioral and Experimental Economics and Finance workshop at the University of Sydney, the Department of Economics at the University of Zurich, University College London, and the 2019 American Finance Association, for very helpful comments and suggestions. Special thanks to Andrew Lo for discussing an early version of the paper at the 2019 American Finance Association, and to an anonymous reviewer for providing invaluable comments and suggestions. The analyses and conclusions set forth are those of the authors. Daniela Grinblatt, Philippe Gouillou, and Yeji Sung provided outstanding research assistance. Funding from Columbia University and UNSW Business School is gratefully acknowledged.

[†]Payzan-LeNestour: UNSW Business School, University of New South Wales, Sydney, elise@unsw.edu.au. Woodford: Department of Economics, Columbia University, New York, michael.woodford@columbia.edu. Send correspondence to elise@unsw.edu.au and michael.woodford@columbia.edu.

Introduction

Most business activities involve making decisions with uncertain outcomes, and even though predicting the future is impossible, keeping track of how often different things happen can help one to make a better decision. As Dixit and Nalebuff (1991) suggest, “Even though you can’t guess right all the time, you can at least recognize the odds.” However, behavioral scientists have long established that “recognizing the odds”—i.e., learning outcome probabilities from observing the realized outcomes in one’s environment— is itself no simple task.

Many discussions of this issue emphasize the inherent difficulty in telling apart different possible statistical models based on finite data samples, or many people’s limited understanding of the principles of statistical inference. However, these discussions typically take for granted that people are aware of exactly what has happened in the specific events that they have observed; it is the regularity that underlies the individual observations that is assumed to not always be learned. But what if even perception, or the accuracy of recognizing and later recalling these observations, cannot be taken as given? In this study, we propose that people are hampered in their perception of outcomes that they expect to seldom encounter, and view the magnitude of such outcomes as less extreme than they really are. Metaphorically, it is as if people are “blind to outliers.”

“Occam’s Razor” might suggest that we should stick with the default assumption that perception is unbiased (on the ground that theories that proceed on this basis are simpler). But a recent neuroscientific paradigm called “efficient coding” suggests otherwise, presenting evidence that the brain is designed to communicate information in a way that economizes on its limited resources.¹ Since an equal focus on all possible values of a given stimulus would imply poor discrimination among different values, it seems plausible that the brain learns which outcomes are more likely to occur and allocates most of its resources to representing the outcomes frequently encountered in the environment at the expense of the unlikely outcomes. This may result in a diminished capacity to discriminate between outcomes far from the most likely ones (“outlier blindness”).

While the principle of efficient coding has been extensively discussed in the neuroscience literature, the applicability of these ideas to financial decision-making may not be obvious. Here we offer both a novel theoretical exposition and new empirical evidence, in order to clarify the relevance of the efficient coding principle for finance. We propose a simpler mathematical formulation of the resource constraint, that cap-

¹For example, Laughlin (1981), Tobler et al. (2005), Wei and Stocker (2015), and Ganguli and Simoncelli (2016).

tures the essential structure of important models in the neuroscience literature, with the goal of providing a model that should be “portable” across domains in the sense advocated by Rabin (2013). We also analyze efficient coding under novel assumptions about the objective to be maximized and the statistics of the variables that are to be encoded, that we believe are more appropriate for applications to financial decision-making.

In our exposition, we emphasize in particular the result of outlier blindness, and an important implication of this general principle for financial decision-making, namely the neglect of tail risk. In our model, outlier blindness 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 may initially underappreciate the size of volatility “jumps.” Likewise, the importance of macroeconomic shocks such as the LTCM and Global Financial Crises may have initially been underappreciated.

How much of a problem such “outlier blindness” causes should depend on the statistics of the decision maker’s environment. In our model, the distorted perception of outliers occurs even if the investing environment is frequently shifting (we elaborate in Section 3.1.2). This distorted perception should be particularly extreme in the case of distributions with especially long tails, which unfortunately is a common feature in financial markets.² Our model further predicts quick investor adaptation to a shift in the return distributions, such that outlier blindness occurs even if the investing environment is frequently changing. This contrasts with recent theories which imply that investor adaptation should be quite slow (e.g., Robson and Whitehead (2019)).

We provide experimental evidence for our model in a perceptual task specifically designed to measure outlier blindness in the laboratory. Essentially, we compare how well task participants discriminate between two given values when these values are within the range of values to which they have been exposed (“adapted”) in the previous trials, versus when it is an outlier (outside the range to which they have been adapted). We find that participant perception of outliers is significantly hampered, consistent with our model. The experimental data further support the model prediction that investor adaptation takes only a small number of trials and hence outlier blindness can occur even when financial returns are affected by changing economic conditions.

The phenomenon of outlier blindness could be an important factor in explaining the vulnerability of capitalist economies to financial instability. For example, the crisis triggered by the failure of Lehman Brothers in the fall of 2008 had much more

²For evidence on the fat-tailedness of asset returns, see, e.g., Mandelbrot (1957), Fama (1965), Gabaix et al. (2003), Gabaix et al. (2006), Kelly and Jiang (2014).

severe consequences owing to many institutions' having taken positions that left them exposed to significant risk in such an eventuality. Many observers have been puzzled by the lack of reaction of investors to the news of the housing bubble bursting about two years earlier. Despite evidence for a regime shift, investors continued to neglect tail risk (even though, from the summer of 2007, daily realized volatility began to rise above 20%, with peaks above 40%).³ One explanation for this apparent paradox argues that investors chose to ignore the 2006 regime shift news for reasons of greed (Lo, 2017), or for other psychological reasons related to “motivated beliefs” (e.g., Cheng et al., 2014, and Benabou, 2015). The present study offers an additional, complementary explanation, which is the idea that investors underestimated the importance of the news as a consequence of outlier blindness. For a more recent example, Giglio et al. (2020) report that investor long-term market expectations did not adjust to the implications of the Covid-19 pandemic, which could be related to outlier blindness as well.

Outlier blindness theory also provides a neurobiological foundation for the neglect of financial risk's becoming particularly acute after two decades of a “Great Moderation” or the like. Minsky (1986) has argued that investors become progressively more inclined to underestimate tail risk the longer a period of relative macroeconomic and financial stability continues, so that periods of stability predictably give rise to the fragility that allows a crisis to occur. The current study provides a cognitive mechanism through which this can occur: in our model, after a prolonged period of good returns and low volatility, when investors imagine—simulate in their mind—the occurrence of extreme negative returns, they under-estimate their size. As a result, they give such outcomes insufficient weight in decision making.⁴

Relation to the literature This paper adds to a recent and growing literature on risk perception. The literature on “neglected risks” shows how agents can underestimate the probability of extreme events until they occur, as a result of a number of cognitive biases such as the availability and representativeness heuristics,⁵ undersampling extreme events (Hertwig et al., 2004), and wrongly assuming a Gaussian functional form when learning about fat-tailed payoffs under model uncertainty.⁶ Here we propose an additional source of tail risk neglect, which is not learning-related. Rather, it pertains to the way that outliers are perceived as less extreme than in reality in the first instance (even before the process of inferring patterns from individual

³See Lo (2017) and Gennaioli and Shleifer (2018) for further discussion.

⁴We thank Patrick Bolton for bringing to our attention this connection between Minsky's hypothesis and our theory of outlier blindness.

⁵For example, e.g., Gennaioli et al., 2012, Gennaioli et al., 2015, and Jin, 2015.

⁶For example, Taleb (2004), Donnelly and Embrechts (2010), and Payzan-LeNestour (2018).

instances), leading to underweighting rare events.

One may wonder how such underweighting of rare events can be reconciled with the *overweighting* of low-probability events predicted by prospect theory, for which laboratory experiments have provided ample evidence. In addition, a wide range of empirical findings in finance (e.g., Barberis and Huang (2008), Kumar (2009), and Kumar et al. (2011)), insurance (e.g., Barseghyan et al. (2013)), and gambling (e.g., Barberis (2011)), suggest that people overweight rare events in their decision-making. While at first glance, this may seem at odds with our theory, it needs not be, for two reasons.

First, and as noted by Barberis (2013), many of the foregoing applications of probability weighting studied in prior work can be thought of as “*decisions from description*”, that is, settings in which the decision maker knows both the possible outcomes of a gamble and the probabilities with which each of them should occur — either because they are informed of these probabilities (as in many laboratory experiments), or because the probabilities can be deduced through simple reasoning, or because historical data have been tabulated in a way that makes the (empirical) frequency distribution evident. As we discuss further below, our theory of outlier blindness applies primarily to a different kind of settings, namely, “*decisions from experience*”, in which people have to learn the distribution of the gambles by sampling from them, as is often the case in the business world. For instance, investment valuations—both in publicly traded securities and other asset classes such as private equity, and of course trading decisions—often rely on experience. Laboratory evidence suggests that rather than being over-weighted, low-probability outcomes are under-weighted in the case of decisions from experience (e.g., Hertwig et al. (2004), Ungemach et al. (2009), and Erev et al. (2010)). However the mechanisms underlying such under-weighting are not well-understood (e.g., Ungemach et al. (2009) and Barberis (2013)). Our results here suggest that outlier blindness may be among these mechanisms.

Furthermore, and importantly, the same principles of efficient coding that lead to our prediction of tail risk neglect in decisions from experience can also predict the well-established pattern of overweighting of rare events in decisions from description, when these principles are applied to the perception of outcome probability, a crucial aspect of decisions from description (see Section 3.3).

The paper also adds to a growing literature that seeks to understand the implications of “imprecise perception” (the idea that decisions are based on internal representations that are imprecise; see Woodford (2020)). For example, Gabaix and Laibson (2017) show how imprecise representation of future payoffs can generate apparent temporal discounting in perfectly patient agents. Payzan-LeNestour et al. (2020) provide evidence for a perceptual bias (“after-effect”) in investor risk percep-

tion. Frydman and Jin (2020) study how a change in the distribution of outcomes from one environment to another affects risk taking; they emphasize the effects of a change in the variance without changing the mean outcome. We instead consider the effects of a change in the mean of the distribution and focus on the phenomenon of outlier blindness. Robson and Whitehead (2019) also consider the effects of a change in the mean but their focus differs from ours in that they emphasize the effects of a change in the coding rule on which stimuli should be considered “large” or “small”, rather than on the fineness of the discriminations that are made between “outliers”. Moreover, their study is purely theoretical.

The rest of the paper is organised as follows. Section 1 and Section 2 respectively describe the core prediction of our model and the experiments used to test it. Section 3 describes key extensions of the model and explains their implications for financial decision-making. Section 4 concludes by discussing several possible extensions of the current study for future behavioral finance research.

1 Benchmark Model

1.1 From Objective Return Magnitudes to their Internal Representations

To model investor perception of financial returns, we propose a new model that incorporates the basic structure of a number of models of efficient coding from the neuroscience literature, while abstracting from any specific biophysical model of how neurons represent environmental features. Moreover, and importantly, our model defines efficiency of a coding scheme in terms of the reliability of the decisions that are made on the basis of the imprecise internal representations, rather than assuming that some measure of the information content of the internal representation is itself the criterion for efficiency, as is often assumed in the neuroscience literature.

We provide a detailed exposition of our model, and the ways in which it differs from previous expositions of efficient coding theory, in the Appendix. For brevity, this section and Section 3 only present the gist of the model.

We assume that when a given investor observes a return of magnitude x , the investor’s brain produces an internal representation, the information in which can be summarized by a sufficient statistic r (which is assumed to be a single number for simplicity). The representation r is drawn from a probability distribution conditional on the true return magnitude, which we denote by $p(r|x)$. If two return values x_1 and x_2 are presented and the investor must decide which one is greater, the investor’s judgment must be based on the internal representations r_1 and r_2 evoked by the

respective return values.⁷

In our baseline model in this section, we assume that mental processing is efficient, in the sense of maximizing the probability of the investor's correctly judging which of two magnitudes x_1 and x_2 is larger, taking into account inherent limits on the precision of internal representations (which we further specify below). This objective corresponds to the reward structure in the experiment discussed in the next section. In Section 3, we show that the qualitative conclusions of our model still obtain under a variety of alternative choices for the efficiency criterion.

In the case of a binary comparison of the kind considered here, under fairly weak conditions, the investor should respond that the second return value is greater if and only if $r_2 > r_1$.⁸ The probability of an erroneous judgment then depends on the degree of overlap of the two conditional distributions $p(r|x_1)$ and $p(r|x_2)$. When there is little overlap, as in the case shown in the upper panel of Figure 1, in which the means of the two conditional distributions are far apart, judgments should almost always be correct. (Here we assume that the true values satisfy $x_2 > x_1$; thus an error occurs if $r_2 < r_1$.) When instead there is substantial overlap, as in the case shown in the lower panel (the means of the two conditional distributions are close to each other), errors occur with substantial frequency (though still less than half the time).

Following a long-standing tradition in psychophysics (dating at least to Thurstone (1927)), we take the conditional distribution of values for r to be a normal distribution $N(m(x), \sigma^2)$, where $m(x)$ is a continuously increasing function of x , and σ is the same for all returns. The function $m(x)$ thus maps the objective return magnitude x to a so-called "Thurstone scale." The distance between two return magnitudes on the scale indicates the degree of overlap between the two conditional return distributions, that is, the extent to which the two returns are distinguishable (in the sense illustrated in Figure 1). So the steeper the function $m(x)$ over a given range of values of x , the sharper the investor's ability to discriminate between realized returns within this range.

To reflect the fact that the degree of precision with which different values can be distinguished is limited by the available resources of the brain (we have a finite number of neurons and each neuron can vary its activity over only a finite range), we assume that the function $m(x)$ takes values within some bounded interval $[\underline{m}, \bar{m}]$. The ratio $K \equiv |\bar{m} - \underline{m}|/\sigma$ determines the number of distinct return values that can

⁷See Woodford (2020) for a discussion of this approach to modeling human perception.

⁸This is the optimal decision rule to maximize the frequency of correct responses if the conditional probabilities satisfy the monotone likelihood ratio property (i.e., for any true magnitudes x, x' such that $x' > x$, the likelihood ratio $p(r|x')/p(r|x)$ is an increasing function of r), and if any pair of returns are equally likely to be presented in either order (i.e., under the prior probability distribution for the true values (x_1, x_2) , (x, x') and (x', x) are equally likely, for any magnitudes x, x').

be effectively distinguished from one another.




Figure 1: **Two examples of the distributions of possible internal representations of two return magnitudes (x_1 and x_2 , with $x_2 > x_1$), with different implications for how distinguishable the return values are.** Lower panel: the two return magnitudes are difficult to tell apart (there is considerable overlap between the two distributions). Upper panel: they are easily distinguished (there is little overlap).

1.2 Outlier Blindness

The basic idea at the core of efficient coding theory is that given the limited available resources of the brain, increasing the number of states that can be effectively discriminated over one interval of the state space requires a corresponding reduction in the number that can be effectively discriminated elsewhere. To put it differently, there is an inherent trade-off between the degree of accuracy of perception in different parts of the state space. The optimal solution to this trade-off (i.e., the “efficient” way to allocate resources) will depend on one’s objective; here we assume that the function $m(x)$ is chosen to maximize the probability of a correct response in our binary comparison task, where the computation of this probability depends on a prior distribution over the comparisons that one might face.⁹

We show that to maximize the probability of a correct response in the foregoing binary comparison of return values, it is optimal to allocate fewer possible distinctions to regions of the return space that are infrequently encountered, given that the capacity to make fine distinctions in such cases will seldom matter, and finer distinctions

⁹We discuss alternatives to this assumption in Section 3.1 below.

to those parts of the return space where values are most frequently observed. This is accomplished computationally by having the $m(x)$ function increase most steeply in the part (or parts) of the return space with the greater prior density. Here we sketch the basic intuition, leaving the details of the analysis for the Appendix.

Let $F(x)$ be the cumulative distribution function for the prior distribution from which returns are drawn in some environment; thus $q = F(x)$ is the quantile of the return x . Any monotonic encoding rule $m(x)$ can then alternatively be specified as $m(x) = n(F(x))$, where $n(q)$ is a monotonic function that maps quantiles into the Thurstone scale. In our model, the probability of judging a quantity x_1 to be greater than another quantity x_2 depends only on the distance $m(x_1) - m(x_2)$ between the locations to which they are mapped on the Thurstone scale. Thus it depends only on $n(q_1) - n(q_2)$, where q_1 and q_2 are the quantiles of these two returns. At the same time, whether such a judgment is correct or incorrect depends only on the relative order of the two quantiles; thus we can define the probability of a correct decision on the basis of the function $n(q)$ that maps quantiles to the Thurstone scale, in a way that is independent of what the prior $F(x)$ may be.

It then follows that the quantile coding rule $n(q)$ that maximizes the probability of a correct decision (and hence that maximizes expected reward in our experiment) will be the same function $n^*(q)$, regardless of the prior distribution from which returns are drawn. The optimal encoding rule $m(x)$ will be different in the case of a different prior, but *only* because the assignment of quantile ranks q to given returns x has changed. Thus if the prior $\tilde{F}(x)$ associated with a new environment is simply a mean shift of the prior $F(x)$ (so that $\tilde{F}(x) = F(x - \Delta)$ for all x , where Δ measures the size of the shift), and $m(x)$ is the optimal encoding rule for the original prior, then the optimal encoding rule for the new prior will be

$$\tilde{m}(x) = m(x - \Delta).$$

Hence the degree of accuracy with which any two returns x_1 and x_2 can be distinguished is predicted to depend on the quantiles of these returns under the prior distribution. In the case of a prior under which both returns lie in a low-probability-density part of the distribution, their quantiles will not be very different, and they will be difficult to discriminate (the lower panel of Figure 1).

We obtain an even stronger result for the limiting case in which the precision of internal representations is relatively high; in this case, we can show that the optimal quantile encoding rule $n^*(q)$ is close to linear. This means that the optimal encoding rule $m(x)$ satisfies

$$m'(x) \sim f(x), \tag{1}$$

where $f(x) = F'(x)$ is the probability density function implied by the prior. The steepness of the optimal encoding rule depends only on the prior density at that point.¹⁰ In the case of a unimodal prior distribution, this implies that $m(x)$ should increase most steeply over an intermediate range near the mode of the frequency distribution of return magnitude x , and less steeply for more extreme values of x .

As an illustration of this key feature of the model, consider Figure 2, in which the $m(x)$ function (bottom right panel) represents the efficient coding scheme in the case of an environment in which the return distribution is given by the normal distribution $f(x)$ (top right panel). Note that $m(x)$ increases most steeply at the value of x where $f(x)$ is highest, and progressively less steeply for more extreme values of x both above and below this value. Instead, the $\tilde{m}(x)$ function on the bottom right panel represents an efficient coding scheme when the return distribution is given by $\tilde{f}(x)$ (top right panel): $\tilde{m}(x)$ increases most steeply at the value of x where $\tilde{f}(x)$ is highest, and progressively less steeply for more extreme values.

Functions $m(x)$ and $\tilde{m}(x)$ have the same range of possible values (the horizontal axis of the lower right panel), and the same bounded range of variation ($[\underline{m}, \bar{m}]$ on the vertical axis, which represents the Thurstone scale mentioned above). Nonetheless, the degree to which the two magnitudes x_1 and x_2 are discriminated is different in the two cases, as shown by the differing degrees of overlap of the conditional distribution of values for the internal representation r (shown in the lower left panel of the figure). With the function $m(x)$, magnitudes x_1 and x_2 are well discriminated, whereas with $\tilde{m}(x)$, they should be frequently confused. The two alternative coding rules represent alternative uses of a fixed range of possible internal representations; which one is better will depend on the environment in which it is used. Return values like x_1 and x_2 are encountered more often in the case of distribution $f(x)$; hence it is more important to be able to distinguish them well in that case, than if the environment is instead described by frequency distribution $\tilde{f}(x)$. In the latter case, the ability to tell apart x_1 and x_2 is sacrificed, for the sake of sharper discriminations in a different part of the return space (the part where the prior density $\tilde{f}(x)$ is larger).

It is clear from Figure 2 that the $m(x)$ (resp. $\tilde{m}(x)$) efficient coding function is essentially flat in the tails of the frequency distribution $f(x)$ (resp. $\tilde{f}(x)$), implying that outliers (whether values in the far left tail or the far right tail) are poorly distinguished from one another. This is the phenomenon of outlier blindness mentioned in the Introduction, which is our main focus in this study. In Section 3.2 we show that it leads people to underestimate the magnitude of tail events.

¹⁰A similar result is obtained in a number of models of efficient neural coding, though under assumptions different from the ones that we make here, and less obviously appropriate for financial applications, as discussed in Appendix A.4.

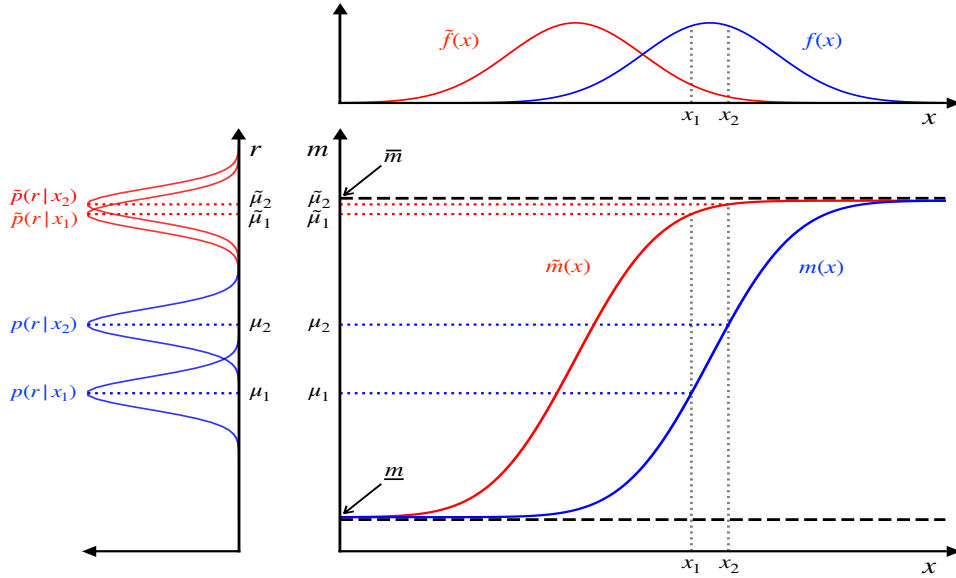


Figure 2: **Two possible distributions for the return magnitude x ($f(x)$ and $\tilde{f}(x)$, top right panel), and their corresponding efficient coding schemes (respectively $m(x)$ and $\tilde{m}(x)$, bottom right panel). The degree to which the two magnitudes x_1 and x_2 are discriminated is different in the two cases, as shown by the differing degrees of overlap of the conditional distributions of values for their internal representations (lower left panel).**

Of note, the phenomenon of outlier blindness should occur with any frequency distribution, as long as there is a long tail on at least one side of the distribution, with increasingly more extreme values occurring with lower and lower probability. The specific Gaussian functional form for the frequency distribution of the return magnitude x used in Figure 2 (and our experimental task uses a Gaussian distribution as well, see Section 2) is not essential for outlier blindness to occur.

Figure 2 further illustrates the important observation that what counts as an outlier depends on what one has reason to expect. The same two return values x_1 and x_2 are not outliers in the environment in which the frequency distribution of returns is described by $f(x)$, and hence they are fairly accurately discriminated; yet the same realized returns are outliers in the other environment, in which the frequency distribution is instead described by $\tilde{f}(x)$, and hence they are frequently confused with one another.

This has important consequences for applications of the principle of outlier blindness to economic and financial decision-making (Section 3). One might well ask how one can say that there is lower discriminability of nearby states in one part of the state space than in another, if it isn't clear how "near" two states should be considered to be to one another, in a way that makes it possible to say that pairs of states in different parts of the state space are "equally distant" from one another. The problem

is solved if we consider how the discriminability of *the same two states* changes when we change the frequency distribution of states for the context in which they occur. This provides the basis for a particularly compelling test of the existence of outlier blindness, which we report next.

2 Experimental Test of the Model

2.1 Experimental Design

It would be difficult to conclusively demonstrate the existence of outlier blindness in investor perception of financial returns by observing investor choices in the field, owing to the inevitable presence of many confounding factors. Therefore, to test our model, we turn to controlled experimentation in a laboratory setting. Even in the case of a laboratory experiment, an experimental task involving economic payoffs would make a clear demonstration of our predicted effect difficult, as participant choices would reflect both their beliefs and their risk preferences, without our being able to disentangle these two determinants of their decisions. Hence, to be able to test the prediction of our model depicted in Figure 2 above, we design a purely perceptual task in which on each trial, the participants have to discriminate between shades of grey in two rectangles (Figure 3). This enables us to examine participants' perception *per se*, since any dependence on personal risk preferences is absent from the task by design.

One may wonder, of course, whether perception works in the same way for a low-level stimulus such as shades of grey, and for the quantitative information on which economic decisions are based (financial returns, monetary payoffs, etc.). A long tradition in behavioral economics argues that there are important analogies between distortions in sensory perception (such as visual illusions) and distortions in the assessment of economic or financial opportunities, starting from Kahneman (2003). There is also a growing consensus in neuroscience that such analogies are pervasive.¹¹ In Section 4, we further discuss this point along with other questions also related to the applicability of the current findings to financial decision-making.

To test the prediction of our model depicted in Figure 2 above, we compute participant accuracy (the fraction of correct replies) in “test trials” in which the shades that are presented to the participant are outliers, in the sense that they are outside

¹¹For example, Carandini and Heeger (2012), Glimcher (2014), Padoa-Schioppa and Rustichini (2014), Summerfield and Tsetsos (2015), and Bhui and Gershman (2018). Empirical evidence for this idea [that the principles underlying the perception of quantitative information and low-level stimuli are the same] can be found in Tremblay and Schultz (1999), Padoa-Schioppa and Assad (2008), and Khaw et al. (2017), among others.

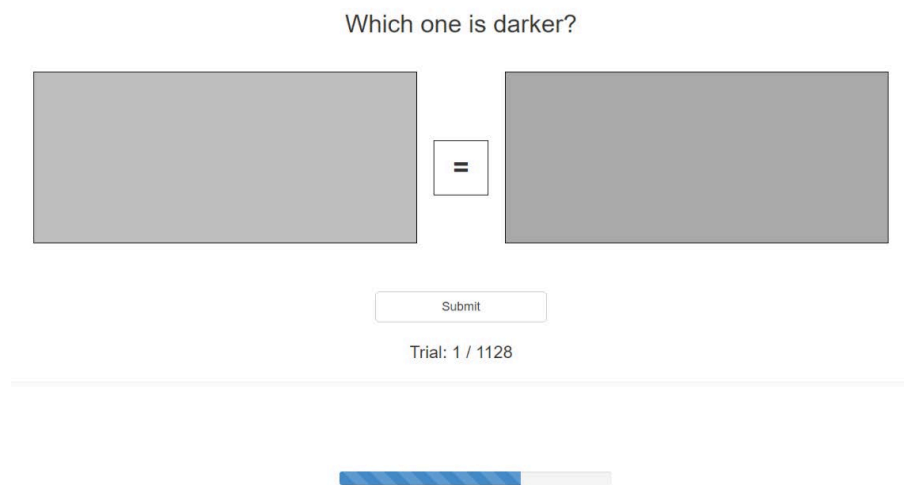


Figure 3: **User interface of the experimental task.** On each of 1,1228 trials, the participant is to click on the rectangle that looks darker, or on an “=” icon displayed in the middle of the screen if the rectangles look of the same color. The time allowed to make a reply is 2 seconds (remaining time is indicated through a timer at the bottom of the screen). One run of the task lasts for about 25 minutes.

the range of shade values to which the participant has been exposed (“adapted”) in the previous 40 trials. We then compare it to participant accuracy in “control trials” in which the exact same shades are presented to the participant, but this time they are not outliers as they are within the range of values to which the participant has been exposed in the previous 40 trials. See Figure 4.

For example, a participant may be asked to discriminate between shade values in the range 2-5¹² for 40 trials, and then be asked to discriminate between shades 8 and 9 (test trial). At some other time during the task, the participant is going to be asked to discriminate between shades 8 and 9 again, this time after exposure to shades around 8 and 9 in the previous 40 trials (control trial). For each participant there are 24 test/control trial pairs of this kind. The main statistic of interest is the difference between participant accuracy in the control vs. test trials: the model predicts a decreased accuracy for outlier perception, i.e., in the test trials, vis-à-vis the control trials.

All the information needed to replicate the following experimental findings is provided in the Online Appendix. We provide a detailed description of the experimental task (its stochastic structure in particular), the instructions for the task, and a detailed justification for each feature of the experimental design. The code to generate the experimental task used in each experiment, as well as the code to run the analyses

¹²The scale of grey used in our task goes from 1 (very light) to 12 (very dark). See the Online Appendix for more details.

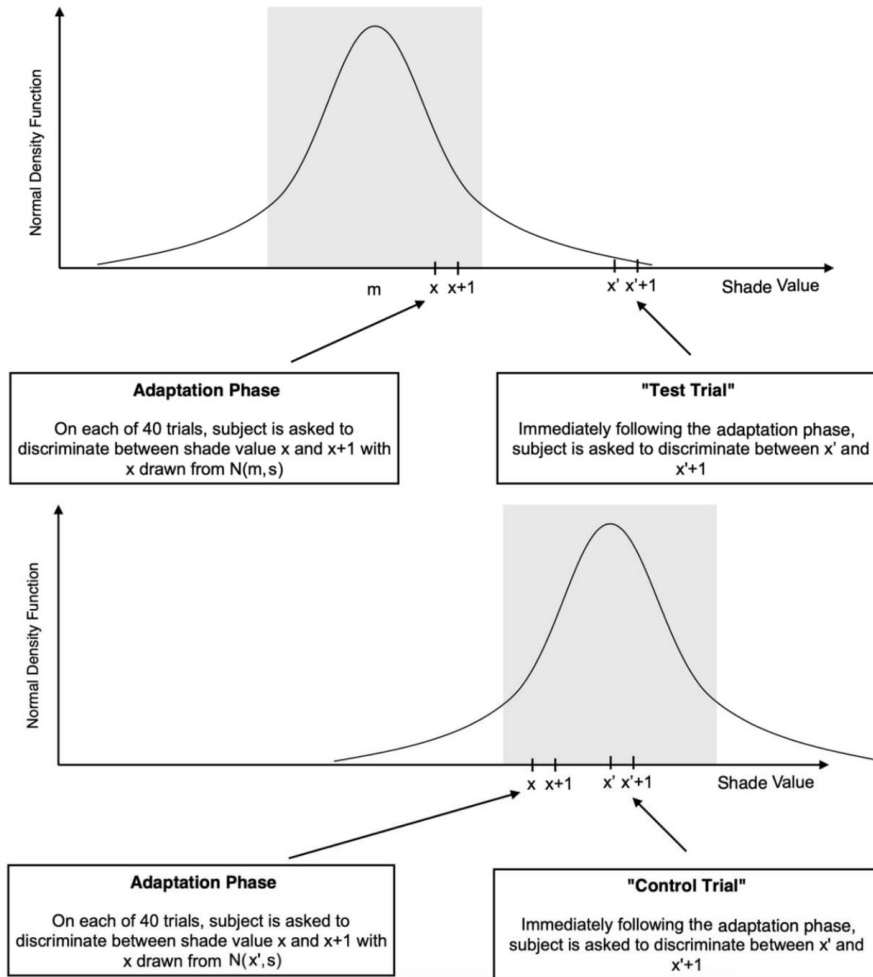


Figure 4: **Diagram explaining the core of the experimental strategy.** Top graph: The participant goes through an adaptation phase in which on each of 40 trials, shade value x is drawn from a normal distribution with mean m and standard deviation s , immediately followed by a “test trial” with shade x' randomly drawn in the range of values located at least three standard deviations from m . Bottom graph: Each value x' in turn defines a 40-trial adaptation phase in which on each trial, the shade value presented to the participant is normally distributed around x' (standard deviation is either 1 or 2; choice is random), followed by a “control trial” in which the shade values presented to the participant are exactly like in the test trial (in this example, x' and $x'+1$). For each participant there are 24 pairs of test and control trials of this kind. Values for m and s are randomly drawn in the sets $\{3, 4, \dots, 10\}$ and $\{1, 2\}$ respectively.

reported in this section, can be downloaded at <https://supplementarymat.weebly.com>. The experimental data for each experiment are also provided there.

2.2 Results

2.2.1 Main Findings

Sixty three undergraduate students from the University of New South Wales participate in the experiment. A-priori power calculations conducted using G*power (Faul et al., 2007) show that this sample size gives us 98% power to detect a medium outlier blindness effect. To maximize data quality by ensuring that the participants pay attention on each trial, we provide them with high monetary incentives. One key aspect of this strategy consists of rewarding high performance in the task through high payoffs (for the other aspects, see the Online Appendix). Eight participants earn more than \$70, 32 earn more than \$50 (mean: 43.8; median: 50.10; mode: 45.65; std: 36.7). All aspects of our experimental protocol are approved by the *UNSW Human Research Advisory Panel (HREAP)* as part of the requirements of the *National Statement on Ethical Conduct in Human Research*, as well as by *UNSW Business School Experimental Research Laboratory* for compliance with the standard rules imposed on all lab users. These rules include providing participants with significant monetary incentives, and the “no-deception rule”.

The average response time in the experiment is 0.9 sec (min: 0.66; max: 1.41; std: 0.14). There is a negative correlation between accuracy and response time (one-way Anova test: $f = 1454$; $p < .00001$, two-sided). Missed trials (trials in which the participant fails to reply within the required time) seldom occur (average percentage of missed trials across participants: 0.005%). Supplementary statistics (regarding response time, earnings, missed trials, and accuracy) are provided in the Online Appendix.

The main finding is that across participants, the perceptual accuracy for a given value is significantly lower when the value is an outlier (in the test trials) than when it is not (in the control trials), as predicted by our model. Our statistic of interest (i.e., the accuracy difference between control and test trials) averaged across participants is significantly positive (paired t-test: $t = 11.7$; $p \sim 0$; see Figure 5 and Table 1). Strikingly, the statistic is positive for almost all the participants (60 out of 63; see Figure 6).

2.2.2 Placebo Test

To check that the foregoing findings truly reflect outlier blindness rather than other causes, we run a follow-up “placebo experiment” which replicates the design of the original experiment in all respects except for that the participants [N=35, same cohort as in the original experiment] do not have time to form expectations about shade

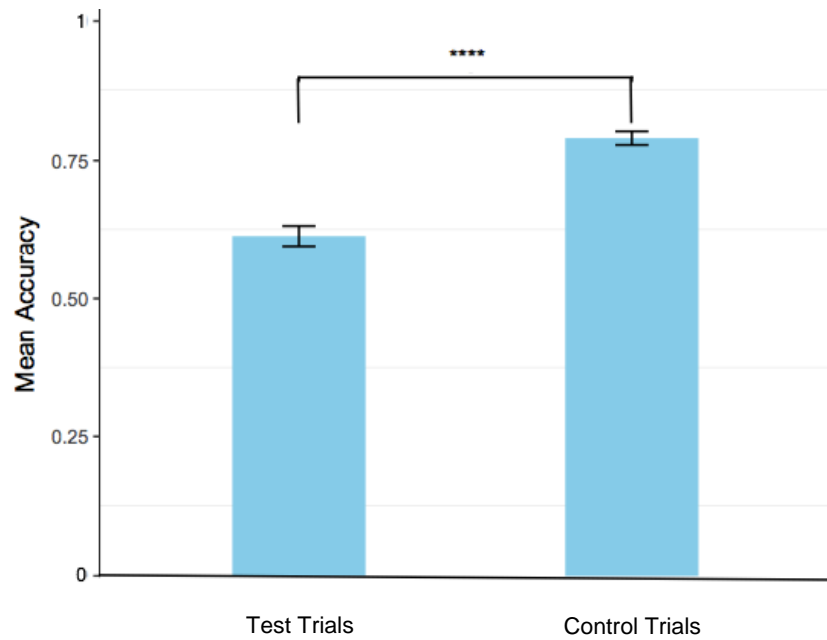


Figure 5: **Mean accuracy in the test trials (left) and control trials (right).** Heights of the bars indicate the mean accuracy averaged across the 63 participants. Accuracy is defined as the fraction of correct replies. Line segments indicate standard error of the mean (sem). **** $p < .00001$. By design, the shade values presented in the test and control trials are the same; they are outliers when presented in the test trials and within the expected range when presented in the control trials.




Figure 6: **Comparative accuracy levels in the test versus control trials for each task participant.** Each data point corresponds to one participant (N=63; 24 observations per participant). x axis: accuracy in the test trials. y axis: accuracy in the control trials. Data points above the 45 degree line correspond to participants for whom accuracy is decreased in the test trials, as predicted by our model.

values as the adaptation lasts for only 3 trials (the minimal achievable in the current design; see the Online Appendix for details). Thus, the values presented in the test

trials do not count as “outliers” and hence there is no reason to expect a decreased accuracy in the test trials in this case. Accordingly, we cannot reject the null hypothesis that accuracy is the same in the test and control trials of the placebo test; see Table 1. The absence of effect in the placebo test is also apparent in Figure 7 which displays the distribution of accuracy in the test and control trials across participants, for both the original experiment (blue curves) and the placebo experiment (red curves). For the latter, the distributions of accuracy in the test and control trials overlap. This contrasts with the original experiment for which the outlier blindness effect is manifest (blue curves: the mean accuracy level at the test trials is shifted to the left relative to the mean accuracy at the control trials).

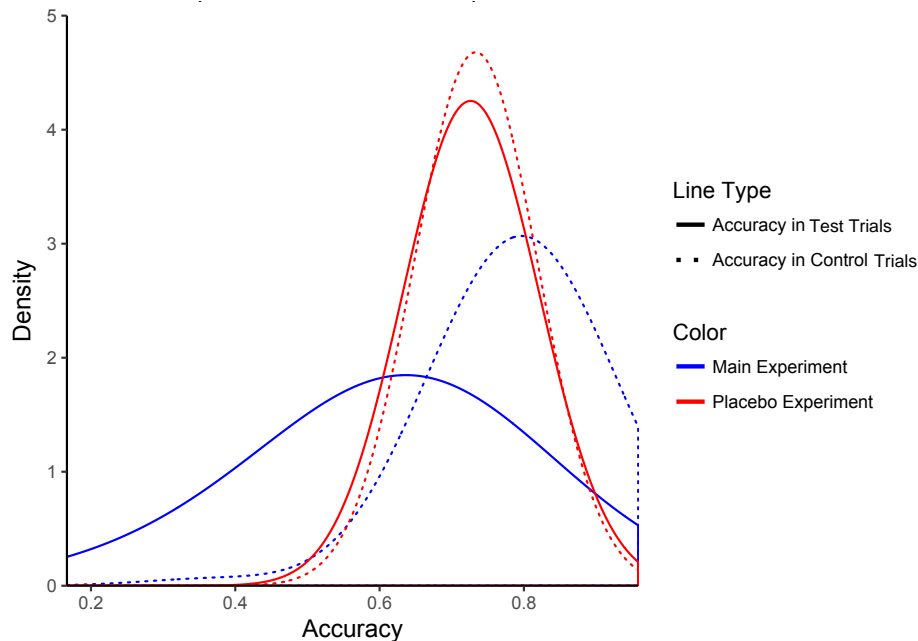


Figure 7: **Density plot of accuracy in the test and control trials, for both the original experiment (blue) and the placebo test (red).** The density is derived across participants, based on 63 participants, 24 observations per participant (original experiment), and 35 participants, 320 observations per participant (placebo test).

2.2.3 Evidence for High-Speed Adaptation

An important question for applications of our theory is how quick the brain is to adapt to changes in the frequency distribution of the magnitudes to be encoded. To test whether rapid adaptation is possible, we run the “5-trial adaptation experiment” [N=31; same cohort as in the original experiment] in which adaptation lasts for only 5 trials. We find a significant outlier blindness effect in that experiment (paired t-test: $t = 4.2$, $p < .001$, two-tailed; see Table 1). Thus, an adaptation length of five trials

appears to be sufficient for the agent to form expectations and hence for the outlier blindness effect to reappear.

Note that the effect found in our baseline experiment is only partially recovered: the outlier blindness effect is significantly larger in the original experiment than in the 5-trial adaptation experiment (two sample t-test to compare the statistic of interest in the two experiments: $t = 8.8$, $p < .001$, see Table 2), as expected (a longer period of experience of a distribution of magnitudes allows more precise learning and hence more complete adaptation to the distribution).

2.2.4 Robustness Checks

We run an extensive set of robustness checks which include examining if the outlier blindness effect is robust to an increasing response time in participants, by running a final experiment in which the time allowed to provide a reply on each trial is doubled relative to that in the original experiment. We find a significant outlier blindness effect in that experiment (Table 1). We cannot reject the null hypothesis that outlier blindness is similar in magnitude in the original experiment and the experiment with double response time (Table 2).

Table 1: **Tests of the outlier blindness effect, for each experiment.** The outlier blindness effect is measured by the differential accuracy in the control vs. test trials as explained in the main text. Numbers in parenthesis: standard error of the mean (sem).

	Mean Accuracy in		Two-Tailed	Wilcoxon Signed
	Test Trials	Control Trials	Paired t-test	Rank Test
Main Experiment (N = 63)	0.61 (0.02)	0.79 (0.01)	t = 11.76 p ~ 0	V = 1858.5 p ~ 0
5-Trial Adaptation Experiment (N = 31)	0.72 (0.01)	0.75 (0.01)	t = 4.25 p < 0.001	V = 330 p = 0.006
Placebo Experiment (N = 35)	0.73 (0.01)	0.73 (0.01)	t = 1.12 p = 0.272	V = 396 p = 0.131
Experiment with Double Response Time (N = 33)	0.72 (0.02)	0.84 (0.01)	t = 6.27 p < 0.0001	V = 390 p < 0.0001

Table 2: **Comparison of the magnitude of the outlier blindness effect across experiments.** In each experiment, the magnitude of outlier blindness is measured by the difference in participant accuracy in the control vs. test trials, averaged across participants. Numbers in parenthesis: standard error of the mean (sem).

	(Mean Accuracy in Control Trials) -(Mean Accuracy in Test Trials)	Two-sample t-test
Main Experiment (N = 63)	0.18 (0.02)	t = 1.94, p = 0.057
Experiment with Double Response Time (N = 33)	0.13 (0.02)	
Main Experiment (N = 63)	0.18 (0.02)	t = 8.89, p ~ 0
5-Trial Adaptation Experiment (N = 31)	0.03 (0.01)	

3 Extensions of the Model and Implications for Financial Decision Making

Having demonstrated the existence of outlier blindness in a setting where the nature of the effect can be established unambiguously, we now study the generality of the outlier blindness result and key implications for financial decision-making.

3.1 Generality of the Outlier Blindness Result

3.1.1 Consequences of Assuming Alternative Objective Functions

Some may argue that typical financial decisions involve judgments about magnitudes that are not even formally analogous to the problem considered in Sections 1 & 2, in which the objective is to correctly rank the relative magnitudes of two stimuli as frequently as possible. This is not obviously the right criterion in the case of financial decisions.

First, while it obviously is relevant to an investor to correctly recognize that the mean return on one asset is greater than that on another, it is not true that it should *only* matter to the investor how often they correctly rank two expected returns; it is more valuable to correctly rank two returns when the difference is large than when it is small. An important question, therefore, is whether the result of outlier blindness still prevails if, instead of maximizing the fraction of correct rankings, the coding scheme is optimized to maximize the average magnitude of the return identified as larger by the decision maker. This alternative objective function is the one that would be relevant in the case of an investor who must choose which of two assets to invest in, when the investor’s reward is not a prize for making the “correct” decision (the reward scheme in a typical perceptual experiment), but rather the return on the investment

that they choose.

In Appendix A.5, we show that in the case of this alternative criterion for efficiency, it remains the case that the optimal encoding function should increase more steeply over those parts of the return space with greater prior density. Specifically, we show that because it is now particularly valuable to make correct judgments when larger quantities are at stake, it is no longer efficient to reduce discriminability in the tails of the distribution to quite the same extent as with the benchmark accuracy-of-ranking objective; but this effect is not strong enough to overturn the result of outlier blindness. Recall that in our baseline case, the optimal encoding rule satisfies (1). In the case of the alternative expected-value objective, this takes the more general form

$$m'(x) \sim f(x)^\alpha, \quad (2)$$

where now $\alpha = 2/3$, rather than having $\alpha = 1$ as with the accuracy-of-ranking objective. So the slope of the encoding rule is still an increasing function of the prior density.

Another alternative criterion of relevance for some economic and financial decisions is that of minimizing the mean squared error of the estimated value \hat{x} of the state, under the assumption that the estimate must be based on the internal representation r . To do well on this criterion, it is again necessary not simply to be able to recognize the correct ordinal ranking of states, but to be able to judge *how much* larger some are than others. Yet again we show that the optimal encoding rule under this alternative objective satisfies (2), except now with an exponent of $\alpha = 1/3$. So the prediction of outlier blindness holds in this case as well.

Moreover, and importantly, our model accounts for the possibility that investors' losses from poor recognition of the magnitude of a monetary payoff may not be well described by any translation-invariant loss function of the kinds considered above.¹³ Instead, errors in recognizing the exact magnitude of the payoff might have more serious consequences for magnitudes in a certain range. For example, monetary losses larger than a particular size might become especially disastrous for an investor, so that the utility loss is a nonlinear function of the monetary loss. If such nonlinearities become particularly severe in the case of especially large (though unlikely) losses, the increasing significance for utility of further increases in the magnitude of the loss can

¹³Let $L(x^{chosen}, x^{not})$ denote the loss function, where x^{chosen} is the monetary amount that is judged to be larger and x^{not} is the monetary amount judged to be smaller. The functions considered above are decreasing in the true difference $x^{chosen} - x^{not}$, but independent of the two absolute magnitudes. A second category of translation-invariant functions considered above is $L(\hat{x}, x)$, where x is the true magnitude and \hat{x} the decision maker's estimate, that is increasing in the difference $|\hat{x} - x|$ but independent of the two absolute magnitudes.

be a reason for it to be valuable to make sharper discriminations between different large quantities, despite the infrequency with which they will be encountered.

We find that the model's main prediction of outlier blindness holds in cases of this kind as well. Suppose that an investor must choose between investments that promise two different possible returns x_1, x_2 , but that the reward for choosing a particular investment is not the monetary return x but a utility $u(x)$, where $u(x)$ is an increasing but nonlinear function — perhaps with $u'(x)$ becoming particularly large for large negative values of x . The problem of efficient coding in this case is mathematically identical to the expected-value maximization problem discussed above, except that the “ x ” in the previous problem (the state variable, the expected value of which is to be maximized) is now the utility associated with a particular return, not the monetary value of the return. Thus condition (2) still applies, if the function m described there is understood to encode the level of utility rather than the financial return. Since the utility is itself a monotonic function of the financial return, we can again write the mean of the distribution from which the internal representation r is drawn as an increasing function $m(x)$ of the financial return. In this case, the optimal rule for encoding the financial return will satisfy the more general relationship

$$m'(x) \sim u'(x)^{1-\alpha} \cdot f(x)^\alpha, \quad (3)$$

where again $\alpha = 2/3$.¹⁴

If the prior distribution for x has a long left tail, with larger negative values occurring with progressively smaller prior densities, the factor $f(x)^\alpha$ will become small for large negative outliers. But if on the other hand $u'(x)$ also becomes larger the more negative is x , this would be a reason for it to be efficient to discriminate *more* sharply, rather than less sharply, between larger negative values of x . These two considerations cut in opposite directions, and the marginal-utility effect would tend to undercut the logic of outlier blindness. Nonetheless, the rate at which marginal utility increases will have to be quite severe in order for the prediction that efficiency requires outlier blindness to be overturned. Suppose, for example, that the prior $f(x)$ is a Gaussian distribution, while for negative x , $u(x)$ is proportional to $-(-x)^\gamma$, for some $\gamma > 1$. Then no matter how large γ is, it will still be the case that for large negative values of x , $m'(x)$ will decrease to zero as x becomes more negative. Thus our model still implies outlier blindness, even if the rate at which $m'(x)$ shrinks for large negative returns is somewhat reduced by the consideration that (for ranges of values with equal prior densities) it is desirable to be able to make sharper discriminations

¹⁴The same formula holds, but with $\alpha = 1/3$, if instead the objective is to minimize the mean squared error of an estimate of the utility associated with monetary return x .

over a range of values where marginal utility is greater.

3.1.2 Perceptual Adaptation, Fast and Slow

Tests of the implications of efficient coding in sensory domains often treat the distribution of sensory features in the environment as something that can be regarded as fixed. But in financial economics applications, it is less obvious that the distributions from which random variables are drawn do not change over time; and to the extent that they do, this raises the following question: does adaptation occur fast enough for the brain to be able to implement efficient coding even with little experience with the frequency distribution associated with a transient context?

Discussions of processes through which adaptation could occur on the basis of successive observations in a new environment are only recent, and they often imply that a very large number of new samples should be needed in order for the new prior distribution to be learned with much accuracy.¹⁵ If that is true, then one should not expect outlier blindness to prevail in changing environments.

Our model however suggests the opposite, namely, that outlier blindness may prevail even if the environment is fast changing. Indeed, in our model, the decision maker can learn about the prior distribution by observing a few successive draws from it, as seems to be the case for the participants in the foregoing “5-trial adaptation” experiment. This can occur because all aspects of the encoding function $m(x)$ need not be re-optimized each time a decision maker’s short-term context changes. Rather, $m(x)$ is re-optimized within a restricted family of possible functions, that vary in only one or two dimensions. Adaptation to a new context thus only requires adjustment of those one or two parameters. Given this, it can occur quite rapidly, after only a few observations drawn from the new frequency distribution.

To formalize this idea, we express the encoding function at any point in time in the form $m(x) = \tau(\phi(x))$, where

$$\phi(x) = \alpha + \beta x$$

is an affine transformation of the state space (with α an arbitrary constant and $\beta > 0$), and the “template function” $\tau(\phi)$ is a continuous, non-decreasing function with range $[\underline{m}, \bar{m}]$, that is assumed to adjust only slowly over time. Over any epoch over which the template function remains essentially constant, the above equation defines a two-parameter family of possible encoding rules (parameterized by α and β), all of which will be continuous, non-decreasing functions, and all of which will remain within the

¹⁵See, for example, Robson and Whitehead (2019) in the economics literature, or Bredenberg et al. (2020) and Aridor et al. (2020) in the literature on computational models of neural coding.

bounds $[\underline{m}, \bar{m}]$. The parameters α and β adjust quickly to a new short-run context, so that the expected probability of a correct choice (under the current short-run prior $f(x)$) is maximized, within the two-parameter family of encoding rules consistent with the current template function.

One implication of the model is that if the distribution of values of x shifts from $f(x)$ to $\tilde{f}(x) = f(\tilde{\phi}(x))$, where $\tilde{\phi}(x)$ is an affine transformation of the kind specified above, the optimal encoding rule for this new environment will be given by $\tilde{m}(x) = m(\tilde{\phi}(x))$. That is, only two parameters of the frequency distribution (the mean and standard deviation) need to be updated in order for the perceptual system to adjust to a shift in the statistics of the environment. Hence even a small sample of observations from a new environment might well be sufficient for adaptation to occur. Our experimental results in the 5-trial adaptation experiment provides behavioral evidence that this is indeed what happens. There is also neural evidence that patterns of neural activity evoked by given sensory stimuli appear to quickly adjust to the range of stimuli in a given experiment.¹⁶

In our proposed theory (explained further in Appendix B), the template function is also optimized in response to the decision maker's experience, but at a much lower frequency, allowing finer-grained statistics of the environment to be learned. This explains why the template function should have a sigmoid shape (as assumed in Figure 2) in the case of our application to the encoding of financial returns.

3.2 Tail Risk Neglect

A key practical implication of outlier blindness for financial decision-making is that it leads the decision maker to perceive extreme values as less extreme than the reality. To explain how outlier blindness leads to such "conservative bias," we need to complete our model by specifying how the investor estimates the magnitude of a given return based on its internal representation r .¹⁷

We assume that the brain produces an *optimal* estimate of the presented magnitude based on the information contained in r .¹⁸ Here optimality is judged with reference to a particular loss function associated with inaccurate estimates, and a particular prior distribution over the true magnitudes that may be encountered. Be-

¹⁶For example, Tobler et al. (2005), Padoa-Schioppa and Assad (2008), and Soltani et al. (2012).

¹⁷The key point to make here is that the noisy internal representation r is *not* itself the perceived return magnitude. Rather, it is the information that the investor's brain uses to produce an estimate of return magnitude. In the neuroscience literature, this is often referred to as the problem of "decoding" the noisy internal representation.

¹⁸In this last respect our model reflects discussions in Wei and Stocker (2015), Wei and Stocker (2017), Rustichini et al. (2017), and Polania et al. (2019), but differs from the discussion of the implications of adaptive coding in Bhui and Gershman (2018) and Robson and Whitehead (2019).

cause r is not perfectly informative about the true magnitude x , the resulting perceived values are not only generally incorrect, but often also biased on average.¹⁹

To be more specific, let \hat{x} denote the investor’s estimate of return magnitude x , on the basis of the internal representation r . The rule for optimal estimation—in the sense of minimizing the mean squared error of the estimate—makes the estimate equal to the posterior mean for x . That is,

$$\hat{x} = E[x|r],$$

where the conditional expectation is computed using the joint distribution for x and r implied by the frequency distribution $f(x)$ from which the return is drawn and the conditional probabilities $p(r|x)$ from the efficient coding scheme described above.

In the absence of any noise in perception (i.e., if the standard deviation of $p(r|x)$ is null), $\hat{x}(r) = m^{-1}(r)$, so that $\hat{x} = x$ for all x , even extreme outliers. In such a case, there is no bias in estimated magnitudes, regardless of the nature of the transformation $m(x)$. That is, the brain should be able to correct or “undo” any perceptual bias coming from $m(x)$.²⁰ Instead, if encoding is noisy as assumed in our model above, optimal estimates will not only be noisy, but will generally be biased. The bias is furthermore of a particular form (at least globally): it will be a “conservative bias,” in the sense that the bias

$$b(x) = E[\hat{x}|x] - x$$

will be negatively correlated with the value of x .

This conservative bias is familiar to financial economists from the theory of linear regression: if the prior $f(x)$ is Gaussian, and $m(x)$ is a linear function, so that the joint distribution of x and r is a bivariate Gaussian distribution, then $\hat{x}(r)$ will be just the fitted value from a linear regression of x on r , and the joint distribution of x and \hat{x} will also be bivariate Gaussian. The bias $b(x)$ is then a negative multiple of $x - \mu$; it grows as x is made more extreme in either direction, but only linearly in the distance from the mean.

However, the linear case is not the one of primary interest. In such a case, the degree of discriminability of nearby magnitudes (as measured by the ratio of $m'(x)$ to σ) remains constant over the entire range of possible magnitudes, and the bias

¹⁹See Woodford (2020) for further discussion of perceptual biases resulting from Bayesian decoding.

²⁰Through a similar argument, Rustichini et al. (2017) propose a model of efficient coding which involves range adaptation with a linear template function (of the kind discussed in the previous subsection), but in which optimal decoding “undoes” the effects of range adaptation, so that it has few behavioral consequences. We assume optimal decoding, as they do, but show that non-trivial biases remain in our case, owing to the nonlinearity of our encoding rule.

$b(x)$ grows as x is made more extreme, but only linearly with $x - \mu$. Instead, with a sigmoid encoding function of the kind predicted by our model (see Figure 2), the degree of discriminability falls to zero as x is made more extreme. This reduces the signal-to-noise ratio in the internal representation as one moves farther out into the tails of the prior distribution, and increases the extent to which the optimal estimate of x is biased toward the prior mean. As a consequence, the bias $b(x)$ grows more rapidly, and becomes much more important for magnitudes drawn from the tails of the prior distribution.

This is illustrated in Figure 8.²¹ The figure shows how the average estimated magnitude $E[\hat{x}|x]$ varies with x , when the prior $f(x)$ is again a Gaussian distribution $N(\mu, \omega^2)$, but the encoding function $m(x)$ is nonlinear (it is a sigmoid function with bounded range, see Figure 2).²² While the bias is again negatively correlated with x (it is a large negative quantity for all x more than one standard deviation greater than the prior mean, and a large positive quantity for all x more than one standard deviation less than the prior mean), it is no longer a linear function of x . The bias is negligible for values of x within one-half a standard deviation²³ of the prior mean, but is instead substantial in the case of outliers in either direction. Indeed, as Figure 8 illustrates, $E[\hat{x}|x]$ remains within a bounded range, no matter how extreme the true magnitude x .²⁴

²¹The figure is adapted from Woodford (2012b). For more computational details and an extensive discussion, see Woodford (2012b) and Polania et al. (2019).

²²The figure actually plots $E[\hat{z}|z]$ as a function of z , where $z = (x - \mu)/\omega$ is the “standardized” value of x . One can show that this normalized version of the figure is independent of the values of μ and ω , and thus applies to all Gaussian priors. The curve shown assumes a particular value of K , the bounded representational capacity of the neural system. As shown in Woodford (2012a), it has a similar sigmoid shape for other finite positive values of K .

²³The range of values of x for which the bias is approximately zero depends on the value of K ; it becomes wider as K is increased, eventually approaching the entire real line as K is made infinite (the “infinite capacity” case shown in Figure 8).

²⁴Again, these bounds depend on the value of K . They become wider as K increases, eventually exceeding any finite quantity for large enough K .

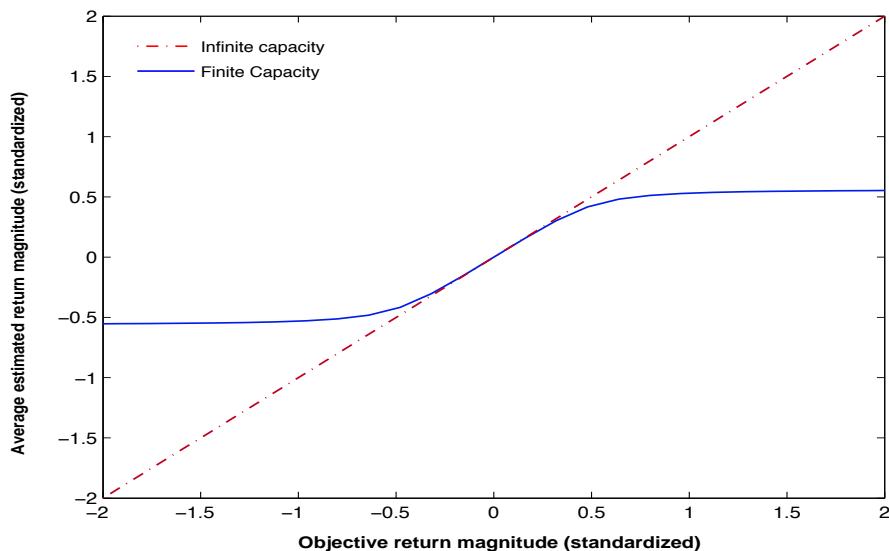


Figure 8: **Tail risk neglect.** The graph shows the objective (45-degree line) vs. average estimated (blue line) return magnitude. The investor fails to appreciate the extremeness of outliers on both tails, with extremely small values being estimated as larger than actual, and extremely large values being underestimated.

While the numerical results shown in Figure 8 assume that the estimate minimizes the mean squared error, it is important to note that it is not necessary that the model assumes this particular loss function to generate the conservative bias result. The more general point is that an optimal decision rule makes \hat{x} a function of the posterior distribution $p(x|r)$ implied by a given representation r , and extreme outliers will be perceived as much less extreme than they actually are. The reason is simply that the person's finite representational capacity does not allow truly extreme outliers to be perceived much differently than outcomes that are somewhat unusual, but not as far out in the tail of the prior distribution to which the person's perceptions have been adapted. Therefore, if the person's estimate of the risks associated with a particular situation is based on these perceived values, then the weight given to truly extreme events will be underestimated—not because the tail events are not included in the sample on the basis of which risks are estimated, but because they are perceived to be less extreme than they actually are.

As discussed in the introduction, this provides a potential explanation for the apparent under-estimation of tail risk by many investors in the period immediately prior to the global financial crisis of 2008. Large negative returns (points well to the left in Figure 8) are predicted to be estimated on average as less negative than they are. Moreover, the size of the bias depends not on the absolute size of the returns in question, but their standardized value: the number of standard deviations below the prior mean the contemplated return happens to be. A given size of negative return will

therefore be more of an outlier (and under-estimated more severely) when perceptions have adapted to a distribution of returns that is higher on average and/or less volatile. Thus the model explains why negative returns of a given (objective) magnitude would be perceived as less damaging following a prolonged period of macroeconomic and financial stability, as proposed by Minsky (1986).

3.3 Additional Factors Affecting the Perception of Tail Risk in Decisions from Description

An investor's mental representation of the distribution of possible outcomes to a particular investment has two elements: (1) their subjective assessment of the magnitude of each possible outcome (which is our focus above); and (2) their subjective assessment of the probability of each outcome. The nature of (2) differs depending on the kind of decision faced by the investor. As stressed in the introduction, there are two categories of decisions that an investor may face: on the one hand, those in which they are not told the probabilities of the different possible outcomes and must learn them through calling to mind a sample of possible outcomes ("*decisions from experience*"); on the other hand, the so-called "*decisions from description*" in which they are explicitly told the outcome probabilities. Although the former kind of situation is our main focus in this study, it is important to also consider what should happen in the case of a decision from description, given that we agree that both cases are relevant to financial decision-making, as stressed by Barberis (2013).

A key point to make is that for the case of decisions from description, (2) has to come from forming an internal representation of the numerical information provided about the probability of each possible outcome. It thus makes intuitive sense to formalize "outlier blindness" not only for the perception of outcome magnitude but also for the perception of outcome probability in this case.

We show in Appendix C.2 that applying our theory of outlier blindness to the perception of outcome probability—specifically, the joint hypothesis of efficient coding and optimal decoding of probability information—implies a conservative bias in the probability estimates of the same kind as the bias described above for the numerical information about outcome magnitudes. In the case of outcome probability, "conservative" estimates mean that small probabilities will not be perceived to be as small as they really are, while at the same time the probabilities of near-certain outcomes will not be perceived to be as close to one as they really are. This consideration provides a reason for extreme outcomes to receive excessive weight in the evaluation of a risky gamble in the case of decisions from description; as such, it provides a neurobiological foundation for the kind of biases reflected in the probability weighting function

posited by Prospect Theory.²⁵

Outlier blindness thus has two opposing effects on investor perception in the case of decisions from description: on the one hand, it leads to underweighting extreme events, as a result of the conservative bias in the perception of outcome magnitude; on the other hand, it leads to overweighting them, as a result of the conservative bias in the perception of outcome probability. Which of these factors outweighs the other will depend on fine details of the risks being evaluated and the prior distribution from which they are assumed to have been sampled. For example, the evidence documented in Kahneman and Tversky (1979) supports the idea that in the case of simple gambles involving a small probability of a relatively extreme outcome, the second factor dominates, leading to overweighting extreme events.

This contrasts with the case of decisions from experience, in which the assessment of outcome probability does *not* come from forming an internal representation of the numerical information provided about outcome probability; rather, it has to come from an entirely different process, namely learning (sampling outcomes). Therefore, in this case, outlier blindness only applies to the assessment of outcome magnitude, resulting in the underweighting of extreme events described above. And indeed, laboratory evidence suggests that in the case of decisions from experience, low-probability outcomes are underweighted rather than being overweighted (e.g., Hertwig et al. (2004), Ungemach et al. (2009), and Erev et al. (2010)).

In this way, our theory of outlier blindness can make sense of both the pattern of underweighting rare events in decisions from experience, and that of overweighting them in decisions from description. Thus it allows a unified understanding of tail risk neglect on the one hand, and investor overweighting of small-probability events on the other — two important ideas from prior work that at first glance may seem contradictory.

Another possible reason for overweighting of small-probability “outlier” outcomes in the case of decisions from description may be that greater attentional focus is directed to these particular outcomes, as postulated by the theories of Bordalo et al. (2012) and Koszegi and Szeidl (2013). Our theory can be extended to allow for effects of differential attentional focus on particular contingencies, if we suppose that the capacity K allocated to the encoding of returns is greater for the returns associated with certain contingencies. If the contingencies under which extreme returns occur are the ones to which more cognitive resources are devoted, this could be a further reason for extreme outcomes to be over-weighted in decisions from description.

By contrast, there is much less reason to expect salience effects of this kind in

²⁵It should be noted that in our theory, the degree of probability weighting should vary depending on the prior distribution to which preferences have been adapted.

the case of decisions from experience, where the investor simply happens to have observed different returns on different occasions, without having had reason to treat some outcomes as instances of a different category. Hence one would not expect salience effects to outweigh the effects of outlier blindness in the case of decisions from experience.

4 Discussion

To summarize, we propose a novel model of efficient coding which, relative to prior efficient coding models, is both simpler and more general in dimensions that are highly relevant for financial economics. It is simpler in that the model abstracts from any specific biophysical model of how neurons represent environmental features. At the same time, it is more general in that it considers a wide range of possible criteria for the efficiency of a coding scheme. In our benchmark case — the one that corresponds to the incentives provided in our experiment — efficiency is defined as maximizing the accuracy with which two return values are discriminated. But we show that the core prediction of our model also obtains when the reward for a correct decision varies across different cases, and when the decision is more complex than a simple choice between two options. It also holds whether or not the investment context is assumed to be stable, and even if it is assumed to be affected by frequent regime shifts (as is sometimes the case in financial trading for example).

The collection of experimental findings strongly supports the key prediction of our model, which is the idea that investors can make fine distinctions between values in the expected range but only coarse distinctions (if any at all) between outliers. Our model also suggests that investor misperception of extreme values may be particularly pronounced in the case of fat-tailed returns.

That said, it is worth emphasizing certain limitations of our study. One important limitation is that our analysis does not enable us to say how long a sudden new development should continue to be surprising—and hence misperceived because it is out of line with the distribution of values to which people’s perceptions are attuned—if the situation persists. How long does the initial misperception caused by outlier blindness last? We think it is plausible that under some circumstances, macroeconomic developments can be misperceived for months because they are outliers relative to people’s expectations. Likewise, we think it is plausible that traders can underreact to market shocks (e.g., large volatility movements) for several days due to outlier blindness. Yet our theory by itself does not answer this question. It must be augmented by a theory of learning, and above all, further empirical study of the determinants of the

speed with which the reallocation of limited perceptual capacities occurs in different contexts.²⁶

It would be interesting to study in the future whether outlier blindness causes distortions of asset prices, and to quantify its portfolio implications for practitioners. For one example, our results suggest that traders may initially underestimate the size of unexpected volatility movements, possibly leading them to stay on the wrong side of the trade when such moves happen. A recent study by Lochstoer and Muir (2021) provides suggestive evidence of such initial underreaction to volatility shocks.

Finally, our analysis does not preclude the possibility that financial agents will consciously correct for outlier blindness in their decision-making subsequent to dissemination of this research, just as the dissemination of research on stock market “anomalies” has reduced the extent to which they are observed (McLean and Pontiff, 2016). Some have indeed pointed out that while it is not possible to prevent our visual system from being subject to visual illusions, learning about these illusions makes it possible for people to consciously correct their decisions.²⁷ We do not mean that an intellectual understanding of outlier blindness should prevent outliers from being perceived as less extreme than they are; efficient coding should be understood as being applied subconsciously and optimized through evolution rather than through conscious reasoning. Nonetheless, we believe that a better understanding of characteristic misperceptions of the magnitude of financial returns can contribute to investors’ learning how to consciously improve their decisions, by realizing that extreme outcomes may well be more severe than one’s intuitive assessment of them would indicate.²⁸

²⁶Note that our theory is not alone in leaving this crucial question open. Any theory according to which people’s degree of sensitivity to some state depends on how it compares with expectations raises a similar question. The theory of reference-dependent preferences of Koszegi and Rabin (2006) and the diagnostic expectations of Gennaioli and Shleifer (2018) are two other examples.

²⁷For example, Ariely (2010).

²⁸We do not view the increasing automation of trading as a “remedy” for the case of trading, insofar as automation only concerns the mechanical aspects of trading; humans continue to control tasks that involve judgment (such as key trading decisions, forecasting future prices, and gauging whether price levels are correct).

Appendix

Here we explain our model in more detail. We begin with a discussion of the consequences of eventual fully optimal adaptation to a given prior distribution, without asking how long it might take for that to occur; we first show why outlier blindness should occur under the assumptions of our baseline model in section 1, and then derive the consequences of the generalizations discussed in section 3.1. We then discuss our extended model with two different speeds of adaptation, in which range adaptation occurs rapidly but adaptation of the “template function” takes longer.

Appendix A Optimal Adaptation to a Single Prior

We first discuss the implications of efficient coding under the assumption that the magnitudes to be encoded (shades of grey, in our experiment, or financial returns, in the application to an investment problem) are always drawn from a single prior distribution, and that the coding scheme is optimized for this prior. Let each of two magnitudes x_i (for $i = 1, 2$) be independent draws from a prior distribution with probability density function $f(x)$ and associated cumulative distribution function $F(x)$. The support of this distribution is assumed to be the real line (the magnitudes are scalar-valued). We further assume that the prior distribution is atomless, so that $F(x)$ is a continuous non-decreasing function.

We consider imprecise encoding schemes in which each magnitude x_i has a noisy internal representation r_i . Conditional on the magnitudes (x_1, x_2) , the distribution of each of the r_i is independent of the other, and is given by a Gaussian distribution

$$r_i \sim N(m(x_i), \sigma^2),$$

the mean of which depends on the true magnitude x_i , and the variance of which is the same for all possible magnitudes. The encoding function $m(x)$ is assumed to be a non-decreasing function,²⁹ defined on the entire support of the prior distribution, and taking values in a bounded interval $[\underline{m}, \bar{m}]$. As explained in the main text, it is this bound on the range of the function $m(x)$ (together with the positive value of σ) that implies a necessary degree of imprecision in internal representations.³⁰

²⁹In fact, there is no loss of generality involved in restricting our attention to non-decreasing encoding functions, since the use of a non-monotonic function would not allow one to achieve a lower value of the expected loss criterion discussed below. Because this result is intuitive, we simplify the discussion here by directly assuming a non-decreasing function.

³⁰This way of modeling the constraints on the precision of internal representations is also used by Steiner and Stewart (2016), Polania et al. (2019), Robson and Whitehead (2019), and Zhang et al. (2020), among others. The quantity K , proposed in the main text as a measure of the capacity

A.1 Optimal Encoding Under Our Baseline Objective

In our baseline model, the decision maker’s task (as in our experiment) is to determine which of the two magnitudes is larger; we assume that this decision must be based on the pair of internal representations (r_1, r_2) . Our theory of efficient coding proposes that both the encoding function $m(x)$ and the decision criterion (giving a response as a function of (r_1, r_2)) are chosen so as to maximize the probability of the decision maker’s making a correct judgment (that is, a judgment that item 1 is greater, when $x_1 > x_2$, or that item 2 is greater, when $x_2 > x_1$). (Equivalently, we can assume that the probability of an incorrect judgment is minimized.) Here the probability of a given outcome is computed under the assumption that x_1, x_2 are independent draws from the prior distribution, and the optimization is conditional upon the given bounds $[\underline{m}, \bar{m}]$ and the given standard deviation σ for the encoding noise.

Under our assumptions, it can be shown that in the case of any non-decreasing function $m(x)$, the response “magnitude 1 is greater” is more likely to be the correct response if and only if $r_1 > r_2$.³¹ Hence we can assume that this is the response rule (as already proposed in the main text), and consider the optimal choice of the function $m(x)$.

For any encoding function $m(x)$, it then follows that for any true magnitudes (x_1, x_2) with $x_2 > x_1$, the probability of an incorrect response (i.e., the response “magnitude 1 is greater”) will equal

$$\text{Prob}[r_1 > r_2 | x_1, x_2] = \Phi\left(\frac{m(x_1) - m(x_2)}{\sqrt{2}\sigma}\right), \quad (4)$$

where $\Phi(z)$ is the CDF of a standard normal distribution.³² More generally, for any (x_1, x_2) , the probability of an incorrect response will equal

$$\Phi\left(-\frac{|m(x_1) - m(x_2)|}{\sqrt{2}\sigma}\right).$$

The overall probability P^e of an erroneous judgment is then given by

$$P^e = \int \int f(x_1)f(x_2) \Phi\left(-\frac{|m(x_1) - m(x_2)|}{\sqrt{2}\sigma}\right) dx_1 dx_2. \quad (5)$$

for differentiation of alternative magnitudes, is called the “Thurstone invariant” by Zhang et al. (2020), who present evidence that it is a subject-specific magnitude with a similar value for different cognitive tasks.

³¹Here we neglect the case of an exact tie in which $r_1 = r_2$, as this will occur with probability zero, regardless of the encoding function chosen.

³²Note that since r_1 and r_2 are both normally distributed (conditional on the x_i), their difference $r_1 - r_2$ is normally distributed as well, with mean $m(x_1) - m(x_2)$ and variance $2\sigma^2$.

The efficient coding hypothesis requires that the function $m(x)$ be chosen (among all non-decreasing functions satisfying the specified bounds) so as to minimize P^e , for a given prior distribution $f(x)$.

Under our assumption that $F(x)$ is continuously increasing, we can alternatively parameterize a magnitude x_i by its quantile $q_i = F(x_i)$ in the prior distribution, and equivalently specify its noisy internal representation as

$$r_i = \underline{m} + (\bar{m} - \underline{m}) \cdot \tilde{r}_i,$$

where \tilde{r}_i is an independent draw from a distribution $N(n(q_i), \nu^2)$, and we define

$$n(q) \equiv \frac{m(F^{-1}(q)) - \underline{m}}{\bar{m} - \underline{m}}, \quad \nu \equiv \frac{\sigma}{\bar{m} - \underline{m}}. \quad (6)$$

Note that under this alternative measure of the magnitude, q_i necessarily lies in the interval $[0, 1]$, and is uniformly distributed over that interval under the prior; and $n(q)$ is a non-decreasing function defined on the interval $[0, 1]$ that also has a range equal to $[0, 1]$. Thus any encoding function $m(x)$ with range $[\underline{m}, \bar{m}]$ corresponds to a normalized encoding function $n(q)$ with range $[0, 1]$. In terms of the normalized encoding function, the probability of an error can be expressed as

$$P^e = \int \int \Phi \left(-\frac{|n(q_1) - n(q_2)|}{\sqrt{2}\nu} \right) dq_1 dq_2. \quad (7)$$

Note that expression (7) involves no parameters of the model other than ν (which is just the reciprocal of the ratio $K \equiv [\bar{m} - \underline{m}]/\sigma$ defined in the main text). Furthermore the set of admissible functions $n(q)$ is also independent of the model parameters. Thus the feasible degree of precision of comparative judgments depends only on the ratio K (or alternatively the normalized noise parameter ν), and not on the values of the parameters \underline{m} , \bar{m} , or σ individually, as stated in the main text.

It is even more noteworthy that the objective (7) can be expressed in a way that is independent of the prior distribution. This immediately gives us the following key result.

Proposition 1. *For a given value of the parameter $\nu > 0$, let $n^*(q)$ be the non-decreasing function mapping the interval $[0, 1]$ to itself that minimizes the value of P^e defined in (7). (Note that the solution to this problem is independent of the prior $F(x)$.) Then for any prior $F(x)$, the encoding rule $m(x)$ required for efficient coding is given by*

$$m(x) = \underline{m} + (\bar{m} - \underline{m}) \cdot n^*(F(x)), \quad (8)$$

Proof. It follows directly from the fact that $n^*(q)$ minimizes expression (7) that the

corresponding encoding function $m(x)$ given in (8) will minimize (5). Thus this function $m(x)$ solves our efficient coding problem. \square

A.2 Context-Sensitivity of the Efficient Encoding Rule

Proposition 1 shows that the efficient encoding rule $m(x)$ will depend on the prior distribution $F(x)$, as discussed in the main text. In particular, it implies that efficient coding implies range normalization.

Corollary. *Let $F(x)$ and $\tilde{F}(x)$ be the CDFs corresponding to two possible prior distributions for the magnitude x , such that there exists an increasing affine transformation of the state space,*

$$\phi(x) \equiv \alpha + \beta x \quad (9)$$

with $\beta > 0$, with the property that

$$\tilde{F}(x) = F(\phi(x))$$

for all x . And suppose that $m(x)$ and $\tilde{m}(x)$ are the optimal encoding functions in the case of priors $F(x)$ and $\tilde{F}(x)$ respectively. Then

$$\tilde{m}(x) = m(\phi(x)) \quad (10)$$

for all x . That is, $\tilde{m}(x)$ can be obtained from $m(x)$ by shifting and stretching the state space in the same way as is required in order to obtain $\tilde{F}(x)$ from $F(x)$.

This result generalizes the familiar idea of range normalization (more below). The conventional concept assumes a stimulus distribution with a bounded range, and assumes that a second distribution is obtained from an original distribution by stretching the range. Range normalization asserts that when the range of the distribution is stretched, the encoding function should be correspondingly stretched, so that the new range of stimuli are mapped into the same range of internal representations. But as our result shows, the idea can be extended to cases in which the support of the distribution need not be bounded. We can still assert that if a second distribution is obtained from some first distribution through an affine transformation of the state space, the appropriate encoding function in the case of the second distribution of magnitudes should be obtained from the appropriate function in the case of the first distribution through that same affine transformation.

As a simple example, our Corollary implies that if a second prior distribution is a simple horizontal translation of some first distribution, as in the case of the prior distributions $f(x)$ and $\tilde{f}(x)$ shown in the upper right panel of Figure 2, then the

optimal encoding function in the case of the second prior will also be a corresponding horizontal translation of the optimal encoding function for the first prior, as shown in the lower right panel of the figure.

A further step in the reasoning is required to explain the prediction illustrated in Figure 2, which is in turn the key prediction tested (and confirmed) in our experiment. This is the assumption in the figure that a unimodal prior distribution $f(x)$ is associated with an optimal encoding function $m(x)$ with a sigmoid shape, with the steepest part of the sigmoid curve around the peak of the function $f(x)$. We consider this issue next.

Let x_q be the magnitude with quantile q in the prior distribution, i.e., such that $F(x_q) = q$. And suppose that $n^*(q)$ is differentiable at this value of q .³³ It then follows directly from (8) that the slope of the optimal encoding function at this point will be given by

$$m'(x_q) = (\bar{m} - \underline{m}) \cdot n^{*\prime}(F(q)) \cdot f(x_q). \quad (11)$$

It is this slope that determines the discriminability of different magnitudes near x , as discussed in the main text (and illustrated in Figure 2). If we compare the optimal encoding rules for two different prior distributions, and in each case consider the slope $m'(x_q)$ at the value corresponding to quantile q , we find that $m'(x_q)$ will be smaller for the distribution that has the smaller prior density $f(x_q)$ at that quantile. Furthermore, the slope will be smaller exactly in proportion to the extent that the prior density is smaller.

Now let q be a quantile that is relatively extreme (so that x_q is an outlier). The greater the extent to which the prior distribution has a long tail, so that the density $f(x_q)$ is small for all sufficiently extreme values of q , the clearer it is that the optimal encoding rule must be relatively flat at those points in the prior distribution.

A.3 A Small-Noise Approximation

But we can go farther, and show that except when the capacity K is quite small, optimal encoding requires that the slope $n^{*\prime}(q)$ not vary much across different quantiles. This result can be stated in a particularly strong (and simple) form in the case that the encoding noise parameter ν is small (that is, in the limiting case in which relatively precise discrimination is possible).

Suppose that we fix a function $n(q)$, which we shall assume is monotonically increasing and continuously differentiable, but vary the size of ν . Let us first consider the value of the inner integral in (7), as a function of q_2 . For small enough values of

³³This will be true of most q , though it may fail to be true on a set of measure zero.

ν , the value of this integral can be approximated as follows:

$$\begin{aligned}
\int \Phi\left(-\frac{|n(q_1) - n(q_2)|}{\sqrt{2\nu}}\right) dq_1 &= \int q'(n_1) \Phi\left(-\frac{|n_1 - n_2|}{\sqrt{2\nu}}\right) dn_1 \\
&= \sqrt{2\nu} \cdot \int_{-n_2/\sqrt{2\nu}}^0 q'(n_2 + \sqrt{2\nu}\xi) \Phi(\xi) d\xi \\
&\quad + \sqrt{2\nu} \cdot \int_0^{(1-n_2)/\sqrt{2\nu}} q'(n_2 + \sqrt{2\nu}\xi) \Phi(-\xi) d\xi \\
&\approx \sqrt{2\nu} q'(n_2) \cdot \left\{ \int_{-n_2/\sqrt{2\nu}}^0 \Phi(\xi) d\xi + \int_0^{(1-n_2)/\sqrt{2\nu}} \Phi(-\xi) d\xi \right\} \\
&\approx 2\sqrt{2\nu} q'(n_2) \int_0^\infty \Phi(-\xi) d\xi = \frac{2\nu}{\sqrt{\pi}} \frac{1}{n'(q_2)} \tag{12}
\end{aligned}$$

Here the first line follows from a change of the variable of integration, defining the function $q(n)$ as the inverse of the function $n(q)$, and letting $n(q_1) = n_1, n(q_2) = n_2$; and the second line follows from another change of variables, in which define $\xi \equiv (n_1 - n_2)/\sqrt{2\nu}$, and break the integral into two parts, corresponding to the domains over which ξ is negative or positive. The third line then follows from a local approximation of the function $q'(n)$, in which we omit terms smaller than order ν ; and the fourth line from another approximation (extending the range of integration to infinity) that also contributes only terms that are smaller than order ν . The resulting expression (12) is therefore accurate to order ν , for any value of q_2 .

Integrating expression (12) over q_2 , we then obtain an approximation to the objective (7),

$$P^e \approx \frac{2\nu}{\sqrt{\pi}} \int_0^1 \frac{dq}{n'(q)}$$

again accurate to order ν . Thus the normalized encoding function $n(q)$ that will achieve the lowest possible error probability, at least in the case of small enough values of the noise parameter ν , will be the function with the lowest value of the integral

$$\int_0^1 \frac{dq}{n'(q)}. \tag{13}$$

Because of the bounded range of the function $n(q)$, its derivative $n'(q)$ must satisfy the integral equation

$$\int_0^1 n'(q) dq = 1. \tag{14}$$

Thus the optimal slope function $n'(q)$ (in the small-noise approximation) will be the one that maximizes (13) subject to the constraint (14). It is easily verified that the solution to this problem is given by $n'(q) = 1$ for all q . Given that $n(q)$ is restricted

to the range $[0, 1]$, we see that the optimal encoding function is given by

$$n^*(q) = q \tag{15}$$

for all $0 \leq q \leq 1$. Combining this result with (8) then yields the following.

Proposition 2. *Let the prior $F(x)$ be fixed, and consider the implications of efficient coding in the case of each of sequence of progressively smaller values for ν (measuring the degree of encoding noise). The sequence of optimal encoding functions $m(x)$ for different values of ν will satisfy*

$$\lim_{\nu \rightarrow 0} m(x) = \underline{m} + (\bar{m} - \underline{m}) \cdot F(x). \tag{16}$$

That is (up to an affine transformation that does not affect the information content of the neural representation), the optimal encoding function is given by the cumulative distribution function for the prior to which the encoding is adapted.

This result — the optimality of encoding magnitudes by an imprecise record of their quantile in the prior distribution — explains the formula (1) given in the main text. The asymptotic formula (16) is what we assume in Figure 2: in the figure, the prior density functions $f(x)$ and $\tilde{f}(x)$ are two Gaussian distributions (with the same variance but different means), while the associated optimal encoding rules are both affine transformations of the CDFs of these Gaussian distributions. And it immediately gives rise to the prediction of outlier blindness. Equation (16) implies that $m'(x)$ will be proportional to $f(x)$, the prior density at point x in the state space. If the prior distribution is unimodal, with $f(x)$ monotonically decreasing as x moves farther away from the prior mode in either direction, it follows that $m'(x)$ — which determines the accuracy with which nearby magnitudes can be distinguished at that part of the state space — should also be monotonically decreasing the farther x is from the prior mode, i.e., the greater the extent to which x is an outlier.

A.4 Comparison with the Literature

A number of papers in the neuroscience literature discuss an implication of efficient coding that is closely related to our equation (1), though not expressed in terms of the derivative of a function $m(x)$ that maps stimulus magnitudes to a “Thurstone scale.”³⁴ It is common to consider the implications of efficient coding for the acuity

³⁴While this device has long been used in the psychophysics literature as a way of modeling variation in the degree of discriminability of nearby stimuli, it is not commonly used (as we propose here) as a way of modeling the resource constraint in models of efficient coding.

of perception as measured by the “discrimination threshold” — the size of change in the stimulus magnitude required in order for a second stimulus to be recognized to be larger than some reference stimulus with a certain probability. In our model, the probability that a magnitude x will be judged greater than some reference magnitude x_0 is equal to

$$\Phi\left(\frac{m(x) - m(x_0)}{\sqrt{2}\sigma}\right),$$

using (4). Thus the discrimination threshold is the size of increase in x required to increase $m(x)$ by some fixed amount; it varies inversely with the slope of the encoding function $m(x)$. Hence equation (1) implies that measured discrimination thresholds should vary inversely with the prior density in a given region of the stimulus space.

This inverse relationship between the discrimination threshold and the frequency distribution of stimulus magnitudes in the natural environment has been shown to be an implication of efficient coding in a number of biologically-motivated models of efficient coding (e.g., Laughlin (1981), Wei and Stocker (2015), and Ganguli and Simoncelli (2016)). These papers, however, obtain the result under different assumptions than ours. In the case of Laughlin (1981), the internal representation r is a single real number, as in our model, but it is assumed to be a deterministic function of x , rather than random; the limited precision of the representation is instead imposed by assuming that only a finite number of different values of r can be read out. The models of Wei and Stocker (2015) and Ganguli and Simoncelli (2016) instead assume a high-dimensional internal representation, indicating the degree of activation of different members of a large number of neurons in a region of the brain involved in processing the sensory feature in question. However, the information that allows different stimulus magnitudes to be discriminated can largely be summarized by variation in one dimension, as in our model.

Besides modelling the constraint on feasible internal representations differently from what we do here, neuroscience models typically assume that encoding is optimized for a different objective, namely, maximization of the Shannon mutual information between the objective magnitude x and its internal representation r . This is a property of the joint distribution of the random variables x and r that measures how informative r is about the value of x (and vice versa). This objective allows an efficient coding problem to be defined without reference to any intended use of the internal representation in making some further decision, which is convenient in applications to early stages of sensory processing, where similar initial processing is often assumed to take place regardless of the eventual higher-level decisions that will be made on the basis of the sensory information. We instead assume that encoding is optimized to maximize a decision maker’s success at a particular task, under the fur-

ther assumption that optimal use is made of the information contained in the internal representation. (In the analysis above, the measure of success that is maximized is the probability of a correct comparative judgment, which corresponds to maximizing expected reward in our experiment. In the next section, we discuss a variety of alternative criteria that might be maximized instead, corresponding to alternative types of decision problem.)

Moreover, while the neuroscience papers conclude, like us, that the discriminability of nearby states should vary monotonically with the prior probability density, the outlier blindness result that we stress has not been a focus of that literature. One reason is that the probability distributions of sensory features that are studied in the neuroscience literature often do not have the unimodal form assumed in Figure 2; hence the regions of the stimulus space where discrimination is poorer are not necessarily “outliers.” For example, one much-studied case is the accuracy of discrimination between different orientations of lines or edges in the visual environment. It is well-established that more oblique orientations are less precisely discriminated than orientations that are nearly vertical or nearly horizontal (orientations near the so-called “cardinal” orientations). This is consistent with efficient coding, given that near-vertical and near-horizontal orientations occur much more frequently than oblique orientations, in both natural and man-made environments (Girshick et al. (2011), Wei and Stocker (2015), Ganguli and Simoncelli (2016)); but while oblique orientations occur less frequently, they cannot be considered more “extreme” than near-cardinal orientations. The kind of unimodal distribution that we assume in Figure 2 is instead characteristic of financial returns (the application of particular interest to us here), as well as being the kind of distribution used in our experiment.

And finally, the neuroscience papers have mainly focused on differences in the discriminability of different types of stimuli that are thought to be relatively stable (such as the poorer discrimination of oblique orientations, just mentioned), and to explain these in terms of relatively invariant features of the sensory environment. In Figure 2 we have instead emphasized the way in which the theory implies that the degree of discriminability of two objective magnitudes can vary depending on the context in which they are encountered, and this is also the implication that our experiment is designed to test. In this respect as well, our emphases are different from those of much of the neuroscience literature,³⁵ because of our interest in drawing out potential implications of the theory for financial decision-making.

³⁵Exceptions include Wei and Stocker (2017), which contains discussions of the way in which the theory of efficient coding might also explain relatively high-frequency changes in the precision of discrimination when the statistics of the sensory environment change.

A.5 Efficient Coding with Alternative Objectives

Above we have assumed that the encoding function is optimized to maximize the probability of a correct choice, which corresponds to the incentive structure in our experiment. But as discussed in the main text, other objectives are of particular relevance for financial applications.

Maximizing the expected value of the chosen item Suppose instead that encoding is optimized to maximize the average value (i.e., the mean value of x) of the item that is chosen. If the true magnitudes of the items offered are x_1 and x_2 , the probability of an incorrect choice is again given by $\Phi(-|m(x_1) - m(x_2)|/\sigma)$, but now the loss from an incorrect choice is given by $|x_1 - x_2|$, rather than by a positive constant. (An incorrect choice is more costly when the difference in the two amounts is greater.) The expected loss in the case of an encoding function $m(x)$ is then equal to

$$E[\text{Loss}] = \int \int f(x_1)f(x_2) |x_1 - x_2| \Phi\left(-\frac{|m(x_1) - m(x_2)|}{\sqrt{2}\sigma}\right) dx_1 dx_2, \quad (17)$$

instead of (5). In this case, efficient coding requires that the function $m(x)$ be chosen (among all non-decreasing functions satisfying the specified bounds) so as to minimize (17) for a given prior distribution $f(x)$.

Just as above, we can use a small-noise approximation to obtain an explicit solution for the optimal encoding function for a given prior. When σ is small, we obtain

$$\begin{aligned} & \int f(x_1) |x_1 - x_2| \Phi\left(-\frac{|m(x_1) - m(x_2)|}{\sqrt{2}\sigma}\right) dx_1 \\ &= \int_{\xi_{\min}}^0 \frac{\sqrt{2}\sigma}{m'(x_1(\xi))} f(x_1(\xi))(x_2 - x_1(\xi)) \Phi(\xi) d\xi \\ & \quad + \int_0^{\xi_{\max}} \frac{\sqrt{2}\sigma}{m'(x_1(\xi))} f(x_1(\xi))(x_1(\xi) - x_2) \Phi(-\xi) d\xi \\ &\approx \frac{2\sigma^2 f(x_2)}{(m'(x_2))^2} \left\{ \int_{-\infty}^0 (-\xi)\Phi(\xi) d\xi + \int_0^{\infty} \xi\Phi(-\xi) d\xi \right\} \\ &= \frac{\sigma^2 f(x_2)}{(m'(x_2))^2}. \end{aligned} \quad (18)$$

Here the first line uses the change of variable

$$x_1(\xi) \equiv m^{-1}(m(x_2) + \sqrt{2}\sigma\xi)$$

for values of ξ varying between $\xi_{min} \equiv -(m(x_2) - \underline{m})/\sqrt{2}\sigma$ and $\xi_{max} \equiv (\bar{m} - m(x_2))/\sqrt{2}\sigma$; the second line is the small-noise approximation, retaining only terms of order σ^2 or larger; and the final line evaluates the definite integrals.

Integrating (18) over the possible values of x_2 (weighted by their probability of occurrence under the prior), we find that

$$E[\text{Loss}] \approx \sigma^2 \int \frac{(f(x_2))^2}{(m'(x_2))^2} dx_2, \quad (19)$$

if we retain only terms of order σ^2 or larger. The efficient coding problem is then to choose a function $m'(x)$ to minimize (19), subject to the constraint that

$$\int m'(x) dx \leq \bar{m} - \underline{m}, \quad (20)$$

as required by the bounded range of the encoding function.

It is easily seen that the constraint (20) must bind, and that the first-order condition requires that

$$m'(x) \sim f(x)^{2/3} \quad (21)$$

for all x in the support of $f(x)$, which is just the case $\alpha = 2/3$ of condition (2) in the main text. The constant of proportionality, which corresponds to the Lagrange multiplier associated with constraint (20), must be chosen so as to ensure that (20) holds with equality. This determines the function $m'(x)$. Integration of this function then determines the encoding function up to a constant of integration, which can be chosen so as to make the range of the encoding function precisely the interval $[\underline{m}, \bar{m}]$. Thus we obtain

$$\lim_{\sigma \rightarrow 0} m(x) = \underline{m} + (\bar{m} - \underline{m}) \cdot \frac{\int_{-\infty}^x f(\tilde{x})^{2/3} d\tilde{x}}{\int_{-\infty}^{\infty} f(\tilde{x})^{2/3} d\tilde{x}} \quad (22)$$

as an alternative to (16). As explained in the main text, we find that efficient coding implies outlier blindness in this case as well, since the slope of the optimal encoding function is again smaller wherever the prior density is smaller.

Maximizing expected utility of the chosen item We can generalize this result, to consider the case in which the encoding rule maximizes not the expected financial return on the investment that is chosen, but the expected utility from that financial return, where utility is assumed to be some smoothly increasing function $u(x)$ of the financial return x . Let us again suppose that the available financial returns are drawn from a prior distribution $f(x)$. This implies a prior distribution $g(u)$ for the utility

that may be obtained by choosing different investments, where

$$g(u(x)) = f(x)/u'(x). \quad (23)$$

Furthermore, any encoding rule $m(x)$ can be expressed as $m(x) = k(u(x))$, for some increasing function $k(u)$ that satisfies the bounds $\underline{m} \leq k(u) \leq \bar{m}$ for all values of u . We can then define the efficient coding problem as the choice of a function $k(u)$ to maximize the expected value of u from the chosen investment, subject to the bounds on the range of the function.

This problem has exactly the same mathematical form as the one just considered, if we replace the variable x by u , the prior $f(x)$ by $g(u)$, and the encoding function $m(x)$ by $k(u)$. The same argument as the one used to derive (21) implies in this case that the optimal encoding rule must satisfy

$$k'(u) \sim g(u)^{2/3}.$$

Then noting that $m'(x) = k'(u(x))u'(x)$ and using (23), we see that in this case, efficient coding requires that

$$m'(x) \sim [f(x)/u'(x)]^{2/3} \cdot u'(x) = u'(x)^{1/3} \cdot f(x)^{2/3}. \quad (24)$$

This is just the case $\alpha = 2/3$ of condition (3) in the main text. It remains the case that, given any value for the marginal utility $u'(x)$, the optimal degree of sensitivity to alternative financial returns near x will be lower the lower is the prior density $f(x)$ at that value. And as discussed in the main text, it can easily be the case that the optimal value of $m'(x)$ decreases for more extreme values of x , even if $u'(x)$ grows sharply for more extreme values of x .

Minimizing estimation error Another case of interest is one in which an investor must make a decision on the basis of an estimate of the value of some state variable x , with losses (from a less accurate decision) that will be proportional to the squared error of the investor's estimate \hat{x} . In the case the efficient coding problem is to choose an encoding function $m(x)$ and an estimation rule $\hat{x}(r)$ so as to minimize the value of

$$\mathbb{E}[\text{Loss}] = \int \int f(x) \frac{1}{\sigma} \phi\left(\frac{r - m(x)}{\sigma}\right) (\hat{x}(r) - x)^2 dr dx, \quad (25)$$

where $\phi(z) \equiv \Phi'(z)$ is the probability density function of the standard normal distribution.

Again we can obtain an explicit solution using a small-noise approximation. For

any encoding function $m(x)$, the optimal estimation rule $\hat{x}(r)$ is implicitly defined by

$$\int f(x) \phi\left(\frac{r - m(x)}{\sigma}\right) (\hat{x}(r) - x) dx = 0$$

for each value of r . (This means that $\hat{x}(r)$ should be the mean of the posterior distribution for x , given the prior $f(x)$ and the noisy internal representation r .) In the small-noise limit, the solution to this equation can be approximated by

$$\hat{x}(r) \approx m^{-1}(r),$$

where we keep only terms of order σ or larger. Then if we let $\xi \equiv (r - m(x))/\sigma$, and express the loss $(\hat{x}(r) - x)^2$ as a function of x and ξ rather than x and r , we obtain

$$(\hat{x} - x)^2 \approx \frac{\xi^2}{(m'(x))^2} \sigma^2,$$

keeping only terms of order σ^2 or larger.

We thus obtain a small-noise approximation to the loss function (25), given by

$$\begin{aligned} \text{E[Loss]} &\approx \sigma^2 \cdot \int \int f(x) \phi(\xi) \frac{\xi^2}{(m'(x))^2} d\xi dx \\ &= \sigma^2 \cdot \int \frac{f(x)}{(m'(x))^2} dx. \end{aligned} \tag{26}$$

The efficient coding problem is then to choose a function $m'(x)$ to minimize (26), subject to the constraint (20). The first-order condition for this problem is easily seen to require that

$$m'(x) \sim f(x)^{1/3}, \tag{27}$$

which is the case $\alpha = 1/3$ of condition (2) in the main text. Once again we find that efficient coding implies outlier blindness, since the slope of the optimal encoding function is again smaller wherever the prior density is smaller.

Alternatively, suppose that x represents a financial return in some state of the world, and that losses depend not on the squared error of the investor's estimate of this return, but on the squared error of the investor's estimate of the utility implied by this return, where utility is a nonlinear function $u(x)$ of the return. Using the same kind of argument as is used above to derive (24), we can show that in the case of a nonlinear utility function, (27) takes the more general form

$$m'(x) \sim u'(x)^{2/3} \cdot f(x)^{1/3}.$$

This is just the case $\alpha = 1/3$ of condition (3) in the main text. Once again we see that it continues to be optimal to have less precise discrimination between values of x in a region of the support of the prior distribution where the prior density is smaller, so that our prediction of outlier blindness remains valid (with modest qualifications) in this case as well.

Appendix B Fast and Slow Adaptation to Changing Environments

A key implication of our results above is that efficient coding (in the sense of optimal adaptation to some single prior) implies range adaptation. We have already noted this in the case of our baseline model (the Corollary above); in fact, a similar corollary can be established in the case of each of the other two objectives discussed above, in which the objective for which perceptions are optimized depends only on the true magnitude x (rather than on some nonlinear utility derived from x).³⁶ Here we discuss our proposal (in section 3.2 of the main text) that this is not just one of the many ways in which the encoding of returns should adjust to the frequency distribution that investors encounter, but one that should occur more rapidly than other kinds of adaptation.

In our two-speed model of adaptation to environmental frequencies, range adaptation occurs more rapidly than other aspects of the adjustment of the encoding rule required by the theory of efficient coding. Here we explain this idea in the context of our baseline model (i.e., under the assumption that the objective is maximization of the probability of a correct binary comparison). This is purely for convenience; a similar two-speed model of adaptation can be formulated in either of the other cases as well. The key idea of the more general theory is the observation that the prediction of range adaptation — which is a crucial step in the argument for the prediction that we experimentally test — can be obtained under a much weaker (and more plausible) assumption than fully optimal adaptation of the encoding rule to each relatively transitory environment that is encountered.

³⁶We can also establish a version of the Corollary for the cases in which the objective depends on $u(x)$ rather than x itself; in this case, the affine transformation must be an affine transformation of the utility scale rather than an affine transformation of the financial return. We omit the details, which are relatively straightforward.

B.1 Fast Range Adaptation

Consider a parametric family of possible encoding functions of the form

$$m(x) = \tau(\phi(x)),$$

where the “template function” $\tau(\phi)$ is some given continuously increasing function with a bounded range, and $\phi(x)$ is any affine transformation of the form (9) with $\beta > 0$. Taking the template function as given, this defines a two-parameter family of possibilities (with parameters α and β).

For example, suppose that the template function is given by $\tau(y) = \underline{m}$ for all $y \leq 0$, $\tau(y) = \underline{m} + (\bar{m} - \underline{m}) \cdot y$ for all $0 \leq y \leq 1$, and $\tau(y) = \bar{m}$ for all $y \geq 1$. Then the two-parameter family of possible encoding functions consists of the family of affine functions, truncated so that the value of $m(x)$ remains within the bounds $[\underline{m}, \bar{m}]$. This is the class of alternative encoding functions assumed to be possible in models of context-sensitive encoding like that of Soltani et al. (2012), as discussed further below. However, one might well consider other template functions. The adaptation shown in Figure 2 is an example of range adaptation in which the template function has a sigmoid shape (an affine transformation of the normal CDF).

Let us use the notation $\mathcal{F}(\tau)$ for the two-parameter family of possible encoding rules associated with a given template function τ . Then the hypothesis of *fast range adaptation* assumes that in each environment, if magnitudes in that environment are drawn from a prior with density function $f(x)$, the encoding function $m(x)$ used in that environment will be chosen from among the parametric family $\mathcal{F}(\tau)$ so as to minimize P^e , where (5) is evaluated using the environment-specific density $f(x)$. That is, this hypothesis requires only that the parameters α and β of the transformation $\phi(x)$ be optimized for each environment; the template function $\tau(\phi)$ is assumed to remain unchanged (or to adjust only over some much lower frequency).

This is a weaker optimality criterion than the one assumed in the theory of optimal adaptation presented above, since it is only required that $m(x)$ be optimal within the restricted family $\mathcal{F}(\tau)$, rather than that it be optimal within some much broader class of functions, such as the class of all non-decreasing functions with a given bounded range. Nonetheless, even this weaker optimality criterion suffices to deliver the result of the Corollary.

Proposition 3. *Let $F(x)$ and $\tilde{F}(x)$ be the CDFs corresponding to two possible prior distributions for the magnitude x , such that there exists an increasing affine transformation $\tilde{\phi}(x)$ with the property that $\tilde{F}(x) = F(\tilde{\phi}(x))$ for all x . And suppose that $m(x)$ and $\tilde{m}(x)$ are the optimal encoding functions in the case of priors $F(x)$ and $\tilde{F}(x)$ re-*

spectively, from within the parametric family of encoding functions $\mathcal{F}(\tau)$ associated with some template function $\tau(\phi)$. Then regardless of the template function used to define the parametric family, it must be the case that $\tilde{m}(x) = m(\tilde{\phi}(x))$ for all x . That is, $\tilde{m}(x)$ can be obtained from $m(x)$ by shifting and stretching the state space in the same way as is required in order to obtain $\tilde{F}(x)$ from $F(x)$.

Proof. Just as in the proof of Proposition 1, the value of P^e depends only on the normalized encoding function $n(q)$ implied by a given encoding function (which in the present context means, by a given affine transformation) and a given prior; the normalized encoding function implied by a given prior $F(x)$ and affine transformation $\phi(x)$ is defined as

$$n(q) \equiv \frac{\tau(\phi(F^{-1}(q))) - \underline{m}}{\bar{m} - \underline{m}}. \quad (28)$$

Let $\bar{n}(q)$ be the normalized encoding function that minimizes P^e (for the given template function $\bar{m}(x)$).

Now suppose that in the environment with prior $F(x)$, the optimal encoding function within the parametric family is given by $m(x) = \tau(\phi(x))$. The corresponding optimal normalized encoding function $\bar{n}(q)$ is then given by the right-hand side of (28). Since the assumption that $\tilde{F}(x) = F(\tilde{\phi}(x))$ implies that $F^{-1}(q) = \tilde{\phi}(\tilde{F}^{-1}(q))$, we must also have

$$\bar{n}(q) = \frac{\tau(\phi(\tilde{\phi}(\tilde{F}^{-1}(q)))) - \underline{m}}{\bar{m} - \underline{m}}.$$

This implies that in the environment with prior $\tilde{F}(x)$, the affine transformation $\phi(\tilde{\phi}(x))$ results in a normalized encoding function $\bar{n}(q)$, and hence that this affine transformation minimizes P^e . Hence the optimal encoding function for this environment will be

$$\tilde{m}(x) = \tau(\phi(\tilde{\phi}(x))) = m(\tilde{\phi}(x)),$$

as the Proposition asserts. □

Thus the weaker hypothesis of fast range adaptation in each of the environments encountered by a decision maker suffices to imply that if the prior shifts from the distribution $f(x)$ shown in Figure 2 to the alternative distribution $\tilde{f}(x)$ (a horizontal translation of $f(x)$), and the optimal encoding function for the first environment is given by the function $m(x)$ shown in the figure, then the optimal encoding function for the second environment will be the horizontal translation $\tilde{m}(x)$, as shown. The fact that this prediction (the basis for our experimental test) requires only a hypothesis of optimal range adaptation makes it more plausible that adaptation of the kind required can occur relatively rapidly. For the adaptation in question requires only estimation

of optimal values for the two parameters α and β ; and even a relatively small number of observations of a new environment can suffice to allow at least a rough estimate of the way in which these two parameters should be adjusted.

B.2 Comparison with the Literature

Examples of optimal range-adaptation models of this kind include Rustichini et al. (2017), Khaw et al. (2020), and Zhang et al. (2020). Note that the assumed form of the template function is different in each of these cases. For example, Rustichini et al. (2017) assume a truncated linear template function. Khaw et al. (2020) instead propose a theory of optimal range adaptation in which the template function is a logarithmic transformation, of a kind often assumed in sensory domains to which “Weber’s Law” applies. Zhang et al. (2020) assume a template function given by a truncated log-odds transformation. This illustrates the need for a theory of the determination of the template function, of the kind that we propose below (section 5.2.3).

Optimal range adaptation in the sense proposed above represents a generalization of a more common conception of range adaptation (dating at least to the work of Parducci (1965)), in which it is assumed that the frequency distribution $f(x)$ has finite support $[\underline{x}, \bar{x}]$ (the “range” of stimulus magnitude). In this case, the distribution of internal representations r associated with a given stimulus x depends only on its “range adapted” value $\hat{x} \equiv (x - \underline{x})/(\bar{x} - \underline{x})$, indicating where x lies relative to the boundaries of the range. This idea corresponds to a truncated-linear template function, together with the hypothesis that the parameters α and β are chosen, for any given frequency distribution $f(x)$, so that the affine function $\phi(x)$ maps the stimulus range $[\underline{x}, \bar{x}]$ into the interval $[m, \bar{m}]$.

The difference between our theory and this classic formulation is that we do not assume that the template function must be of the truncated-linear form. Instead, the shift in the encoding function shown in Figure 2 can be viewed as an example of fast range adaptation only on the assumption that the template function has a sigmoid shape, and the sigmoid shape is important for our prediction of outlier blindness; it is because of the sigmoid shape that the encoding function is less steeply sloped between the points x_1 and x_2 when those points are in a tail of the prior distribution than when they are near the prior mode.³⁷

³⁷Note however that under the assumption of a truncated-linear template function, as assumed for example in Soltani et al. (2012), outlier blindness would also be observed: there should be no ability to discriminate between stimuli outside the range beyond which the encoding function is truncated. And in the case of a prior distribution with long tails (say, a Gaussian distribution), the tails of the distribution will necessarily be outside the finite range over which the encoding function

In our model, the template function is endogenously determined as an optimal adaptation to the agent's environment, within a much more flexible family, but only through an adaptive process that occurs much more slowly. In the short run, the template function is fixed, and only the parameters α and β of the normalization adjust to the current frequency distribution of magnitudes to be encoded; but eventually the template function adjusts as well. Thus the template function is also assumed to be optimized to maximize the expected probability of a correct choice, but with the expectation calculated by averaging over all of the different short-run contexts that are encountered over the long run.

B.3 Long-Run Adaptation of the Template Function

In our complete model, adaptation of the encoding rule occurs on two separate time scales. Given the current template function $\tau(\phi)$, the affine transformation $\phi(x)$ adjusts relatively quickly in a way that is optimal for the current statistics of the environment. But the template function is also adjusted in a way that adapts it to the statistics of the magnitudes that are encountered, only at a much lower frequency. Thus $\tau(\phi)$ should be optimal relative to the statistics of the decision maker's environment as well; but because $\tau(\phi)$ adapts only slowly, the relevant statistics are average frequencies over a fairly long run, while $\phi(x)$ adjusts more rapidly and so is at each point in time optimal (or nearly optimal) relative to the frequencies with which different magnitudes are encountered in a particular, more transitory environment.

To formalize this idea, we extend the theory of efficient coding presented above in the following way. Suppose that a decision maker transitions among environments (indexed by a parameter θ), with the distribution of magnitudes in each environment given by a distribution $F(x|\theta)$. Each short-run environment lasts long enough for range adaptation to that environment to occur; thus there is an affine transformation $\phi(x|\theta)$ for each possible θ , that is adapted to the distribution $F(x|\theta)$. Finally, suppose that over the long run, different environments θ are encountered with frequencies corresponding to a distribution $p(\theta)$. We assume that the template function adjusts slowly enough that it is the same in each of the environments, but that it is optimized with respect to the long-run frequency distribution $p(\theta)$.

For any environment, let $P^e(F, \phi, \tau)$ be the value of the objective (7) for that environment, if magnitudes are drawn from the distribution $F(x)$, and encoded using an affine transformation of the state space $\phi(x)$ and a template function $\tau(\phi)$, and the normalized encoding function is $n(q)$ as defined in (28). The complete theory of efficient coding with two time scales can then be expressed as follows. There is

is increasing.

assumed to be a single template function $\tau(\phi)$, and an affine transformation $\phi(x|\theta)$ for each possible environment θ , such that (i) for each environment θ , $\phi(x|\theta)$ minimizes

$$P^e(F(\cdot|\theta), \phi(\cdot|\theta), \tau)$$

given the distribution of magnitudes associated with that environment; and (ii) $\tau(\phi)$ minimizes

$$\int P^e(F(\cdot|\theta), \phi(\cdot|\theta), \tau) p(\theta) d\theta,$$

from among the class of all non-decreasing functions with range $[\underline{m}, \bar{m}]$.

Condition (ii) of this definition determines the shape of the optimal template function. Under some (not implausible) conditions, a sigmoid shape of the kind assumed in Figure 2 will be optimal. For example, consider the case in which for each possible environment θ , $F(x|\theta)$ is a Gaussian distribution $N(\mu_\theta, \omega_\theta^2)$, where the mean and standard deviation depend on the environment. In this case, each of the priors can be obtained from any of the others through an affine transformation of the state space. It then follows from Proposition 3 that the optimal encoding function for each environment can be obtained from the function for some single environment, through this same affine transformation.

Thus the optimal encoding function for each environment θ must be of the form

$$m(x|\theta) = \tau\left(\frac{x - \mu_\theta}{\omega_\theta}\right),$$

for some non-decreasing function $\tau(z)$ that is the same for all θ . It follows from this that P^e will have the same value for each environment θ , equal to the value of P^e in the case in which $F(x)$ is a standard normal distribution, and $\phi(x)$ is the identity mapping ($\phi(x) = x$ for all x). Condition (ii) then reduces to the requirement that $\tau(z)$ be chosen so as to minimize P^e (as defined in (5) for the case of a prior over z that is a standard normal distribution). This is the kind of efficient coding problem discussed in the previous section. The asymptotic result (16) implies that at least in the small-noise limit, the optimal template function $\tau(z)$ will be a sigmoid function (in fact, approaching the function $\Phi(z)$ as the noise is made small).

This provides another justification for an encoding function with a sigmoid shape, as shown in the lower right panel of Figure 2, and hence for the prediction that we test in our experiment. Thus the prediction that we test can be obtained as a prediction of either of two possible theories of efficient coding. The first theory is one in which the encoding function $m(x|\theta)$ is fully optimal for the distribution of magnitudes $F(x|\theta)$ for each environment θ ; since this theory implies range adaptation (see the Corollary

above), a change in the prior of the kind shown in the figure should result in a shift of the encoding function of the kind that is shown. The second possible theory is one in which only the affine transformation $\phi(x|\theta)$ is optimized for each environment, while the template function remains the same across environments. On the assumption that the long-run optimal template function (i.e., the one that is optimal for the full range of environments encountered over the long run) has a sigmoid shape, this theory will also imply that a shift in the distribution of magnitudes of the kind shown in the figure should result in a shift in the encoding function of the kind shown in the figure.

It might be wondered why we bother to present this more complex version of our theory, given that the basic version of our model already provides a reason for outlier blindness to be observed. But the conclusion that adaptation occurs relatively quickly is important for financial applications of our theory, inasmuch as investment contexts are often not fixed over time but shifting.

If the encoding rule adapts only to the frequency distribution of monetary amounts encountered over some very long run, one might suppose that this distribution will be quite dispersed, so that none of the values that can be encountered in the current short-run context may be much more unlikely than any of the others; in such a case, one might suppose that none of the values that are encountered should really be “outliers” relative to the range of values to which perception has been adapted, and that the accuracy with which different magnitudes can be discriminated should be approximately constant. If, instead, there is fast range adaptation to the current short-run frequency distribution, it is easier to believe that some values that can occur with non-trivial probability are nonetheless far enough from the center of the short-run prior to be more poorly distinguished than values nearer the mode of that distribution. Thus it is in this case that imprecision in the internal representation of monetary amounts is likely to lead not just to noise in subjective valuations, but also to systematic bias (i.e., outlier blindness).

Appendix C Optimality of Conservative Bias

Thus far, we have analyzed only optimal adaptation of the encoding function $m(x)$; these results suffice to explain the predictions of our model for the experiments discussed in section 2. However, our discussion of the implications of our theory for tail risk neglect in section 3.3 invokes further results about the optimal “decoding” of the internal representation r to produce a subjective estimate $\hat{x}(r)$ of the magnitude that has been observed. We turn now to this issue.

C.1 Bayesian Decoding

Suppose that an investor knows (either on the basis of past experience, or description of the possible outcomes) that a given investment will yield one or another of a set of possible returns $\{x_s\}$, and that each of the possible returns x_s (associated with some state of the world s) is separately encoded, resulting in an internal representation $r_s \sim N(m(x_s), \sigma^2)$. In order to choose the investment with the highest expected return, the investor's decision rule will be a function of the conditional means $\hat{x}_s \equiv E[x_s | r_s]$ for each of the possible states of the world; specifically, the investor should choose the investment for which the weighted average of the estimates \hat{x}_s is highest, where the weights applied to the different states are given by the investor's subjective estimate of the probability of each of these outcomes.³⁸

Thus an optimal decision rule requires that the internal representation be "decoded" to obtain an estimate \hat{x}_s of the return associated with each state, where the optimal decoding rule associates with an internal representation r the implied conditional mean $E[x | r]$. This is often called "Bayesian decoding," since the estimate $\hat{x}(r)$ is the mean of the Bayesian posterior distribution for x , if the prior distribution is updated using Bayes' rule following observation of the noisy representation r .

If the prior distribution is specified by a density function $f(x)$ and the encoding function is $m(x)$, then the Bayesian posterior mean estimate implied by any internal representation r is given by

$$\hat{x}(r) = \frac{\int \exp \left\{ -\frac{1}{2} \left(\frac{r-m(x)}{\sigma} \right)^2 \right\} x f(x) dx}{\int \exp \left\{ -\frac{1}{2} \left(\frac{r-m(x)}{\sigma} \right)^2 \right\} f(x) dx}. \quad (29)$$

Here we assume that $f(x)$ integrates to 1 (as required in order for it to represent a probability distribution), and that $f(x)$ has thin enough tails for the integrals

$$\int_{-\infty}^0 x f(x) dx, \quad \int_0^{\infty} x f(x) dx$$

both to be finite (so that the prior distribution has a first moment); then both the numerator and denominator of (29) are well-defined.

We can show quite generally that Bayesian decoding implies that the estimate \hat{x} will not be unbiased, and more specifically that there will be an overall conservative

³⁸This will be true regardless of whether these weights are formed by sampling from memory, in the case of a decision from experience, or the weights are derived from a noisy internal representation of the probabilities that the decision maker has been told, in the case of a decision from description. The latter case is discussed further below.

bias. If we define the bias in the case of a true magnitude x as $b(x) \equiv E[\hat{x} | x]$, then it must be the case that

$$\begin{aligned}
 \text{cov}(b(x), x) &= E[b(x) \cdot x] = E[E[b(x) \cdot x | r]] \\
 &= E[E[E[x | r] \cdot x - x^2 | r]] \\
 &= E[E[x | r]^2 - E[x^2 | r]] \\
 &= -E[\text{var}(x | r)] \leq 0.
 \end{aligned}$$

Moreover, the final inequality must be strict, unless perfectly precise identification of the true magnitude is possible with probability one. Thus the bias must be non-zero with positive probability, and it must be negatively correlated with the true magnitude (conservatism).

This result holds quite generally, as a simple consequence of noisy encoding combined with Bayesian decoding. The kind of relationship between the average estimate $E[\hat{x} | x]$ and the true magnitude x that is implied by an efficient coding rule is illustrated by Figure 8 in the main text, for the case of a Gaussian prior distribution of returns.

C.2 Biased Probability Weights in the Case of Decisions from Description

In the discussion above, we have supposed that the magnitude x that must be encoded with noise and then decoded in order to make a decision is some monetary payoff or rate of return (whether observed or described). However, in the case of so-called “decisions from description,” the probability information that is presented to the decision maker must also be encoded and then decoded in order to evaluate the gamble. The finiteness of the decision maker’s cognitive resources implies that information about stated (or calculated) probabilities should also be encoded with noise, so that the problem of optimal decoding of the noisy internal representation arises in the case of information about probabilities as well.

Consider for example the case of choice between simple gambles, each of which offers some payoff x with probability p , and zero otherwise. The characteristics of each gamble are then specified by the two quantities x and p . If we suppose that x and p are separately represented, by internal representations r_x and r_p respectively, and furthermore that under the prior distribution x and p are distributed independently of one another, then the decision rule that maximizes the decision maker’s expected financial gain from their decision will choose the gamble with the highest

value of $\hat{p} \cdot \hat{x}$, where $\hat{p}(r_p) \equiv E[p | r_p]$ and $\hat{x} \equiv E[x | r_x]$.³⁹ Thus optimal decoding of the noisy internal representation of the probability information will again correspond to Bayesian decoding.

It then follows from the discussion above that the estimate \hat{p} will be biased, and furthermore that the bias $b(p)$ will satisfy

$$\text{cov}(b(p), p) = -E[\text{var}(p | r_p)] < 0.$$

Thus our theory predicts that the probability weight used in choosing between gambles will differ on average from the true probability of the gamble paying off, and the bias will be conservative (a tendency to underestimate larger probabilities and to overestimate smaller ones).

Now let us assume a more specific model of the noisy encoding of probability information,

$$r_p \sim N(\tilde{m}(p), \sigma_p^2), \tag{30}$$

by analogy with our model of the encoding of monetary payoffs, where $\tilde{m}(p)$ is an increasing function that must satisfy the bounds

$$\underline{m} \leq \tilde{m}(p) \leq \bar{m} \tag{31}$$

for all $0 < p < 1$. Then we can obtain a stronger result about the degree to which small probabilities will be over-estimated on average, using (29). Note that in our theory, the function

$$w(p) \equiv E[\hat{p} | p]$$

plays a role analogous to the probability weighting function in prospect theory (Kahneman and Tversky (1979)), indicating the average weight that is placed on an outcome with true probability p . It is thus useful to study the behavior of this function for small values of p .

To apply (29) to the case of probability estimation, we replace the variable x in the formula by p , the standard deviation σ by σ_p , and so on. We then observe that the exponential factor appearing in this formula must satisfy the lower bound

$$\exp \left\{ -\frac{1}{2} \left(\frac{r_p - \tilde{m}(p)}{\sigma_p} \right)^2 \right\} \geq \min \left\{ \exp \left(-\frac{1}{2} \left(\frac{r_p - \underline{m}}{\sigma_p} \right)^2 \right), \exp \left(-\frac{1}{2} \left(\frac{r_p - \bar{m}}{\sigma_p} \right)^2 \right) \right\}$$

for any real number r . Substituting this expression for the exponential factor in the

³⁹See Khaw et al. (2020) for further explanation of this model of choice under risk in the case of decisions from description.

numerator of (29), we obtain a positive lower bound for the numerator.

The exponential factor must also be bounded above by 1 for all r ; hence the denominator of (29) can be no greater than 1. The lower bound for the numerator and the upper bound for the denominator can then be used to establish the bound

$$\hat{p}(r_p) \geq \bar{p} \cdot \exp \left\{ -\frac{1}{2} \left(\frac{r_p - \bar{m}}{\sigma_p} \right)^2 \right\}$$

for all $r_p \leq m^* \equiv (\underline{m} + \bar{m})/2$, and the bound

$$\hat{p}(r_p) \geq \bar{p} \cdot \exp \left\{ -\frac{1}{2} \left(\frac{r_p - \underline{m}}{\sigma_p} \right)^2 \right\}$$

for all $r_p \geq m^*$, where $\bar{p} \equiv E[p]$ is the prior mean for the probability p .

We can then compute a lower bound for the average estimate \hat{p} , conditional on any true probability p , by integrating this bound for $\hat{p}(r_p)$ over the distribution (30) of possible values for r_p . We obtain the bound

$$\begin{aligned} E[\hat{p} | p] &\geq \int_{-\infty}^{m^*} \frac{1}{\sqrt{2\pi}\sigma_p} \exp \left(-\frac{1}{2} \left(\frac{r_p - \tilde{m}(x)}{\sigma_p} \right)^2 \right) \cdot \bar{p} \cdot \exp \left(-\frac{1}{2} \left(\frac{r_p - \bar{m}}{\sigma_p} \right)^2 \right) dr_p \\ &\quad + \int_{m^*}^{\infty} \frac{1}{\sqrt{2\pi}\sigma_p} \exp \left(-\frac{1}{2} \left(\frac{r_p - \tilde{m}(x)}{\sigma_p} \right)^2 \right) \cdot \bar{p} \cdot \exp \left(-\frac{1}{2} \left(\frac{r_p - \underline{m}}{\sigma_p} \right)^2 \right) dr_p \\ &= \frac{\bar{p}}{\sqrt{2}} \left\{ \exp \left(-\frac{1}{2} \bar{z}(p)^2 \right) \Phi(-\underline{z}(p)) + \exp \left(-\frac{1}{2} \underline{z}(p)^2 \right) \Phi(-\bar{z}(p)) \right\} \\ &\equiv \underline{w}(p), \end{aligned}$$

where we define

$$\underline{z}(p) \equiv \frac{\tilde{m}(p) - \underline{m}}{\sqrt{2}\sigma_p}, \quad \bar{z}(p) \equiv \frac{\bar{m} - \tilde{m}(p)}{\sqrt{2}\sigma_p}.$$

The lower bound $\underline{w}(p)$ is easily seen to be strictly positive for all $\tilde{m}(p)$ satisfying the bounds (31). Furthermore, we see that $\underline{w}(p)$ remains bounded away from zero as $\tilde{m}(p)$ approaches its lower bound \underline{m} . Since this is a lower bound, it follows that $w(p)$ is also bounded away from zero as p approaches zero. Hence for all small enough probabilities $p > 0$, the average decoded value $\hat{p}(r_p)$ will necessarily be larger than p on average. Yet the same is not true for all probabilities; indeed, using a similar argument, we can show that for all probabilities $p < 1$ close enough to 1, the average decoded value will necessarily be less than p on average. Thus our model predicts over-estimation of low probabilities and over-estimation of large probabilities.

It follows that in the case of decisions from description, we should expect low-probability gambles to be evaluated as if the probability of the gamble's paying off were greater than it really is, fairly often. In the case of a gamble that combines a low positive value of p with a value for x relatively far in the right tail of the prior distribution, we should expect \hat{x} to be less than x on average (the conservative bias discussed above), but at the same time we should expect \hat{p} to be larger than p on average. It can then easily be the case that the subjective estimate $\hat{p} \cdot \hat{x}$ of the expected value of the gamble exceeds its true expected value more often than not. It is furthermore quite possible that the subjective estimate of the expected value of the gamble will exceed the subjective estimate of the value of receiving a payment of px with certainty, given that there is no opportunity for probability distortion in the latter case. Thus one could observe apparently risk-seeking behavior, despite under-estimation of the magnitude of right-tail payoffs.⁴⁰

References

- Guy Aridor, Francesco Grechi, and Michael Woodford. Adaptive efficient coding: A variational autoencoder approach. *pre-print posted on bioRxiv*, <https://www.biorxiv.org/content/10.1101/2020.05.29.124453v1>, May 2020.
- Dan Ariely. *Predictably Irrational: The Hidden Forces That Shape Our Decisions*. HarperCollins, 2010.
- Nicholas Barberis. A model of casino gambling. *Management Science*, 58:35–51, 2011.
- Nicholas C. Barberis. The psychology of tail events: Progress and challenges. *American Economic Review: Papers Proceedings*, 103:611–616, 2013.
- Nicholas C. Barberis and Ming Huang. Stocks as lotteries: The implications of probability weighting for security prices. *American Economic Review*, 98:2066–2100, 2008.
- Levon Barseghyan, Francesca Molinari, Ted O'Donoghue, and Joshua C. Teitelbaum. The nature of risk preferences: Evidence from insurance choices. *American Economic Review*, 103:2499–2529, 2013.
- Rahul Bhui and Samuel J. Gershman. Decision by sampling implements efficient coding of psychoeconomic functions. *Psychological Review*, 125(6):985–1001, 2018.

⁴⁰For similar reasons, of course, imprecise probability estimates can easily lead to more risk-averse choices in the case of gambles with negative tail risk. This is indeed observed in the case of decisions from description (e.g., the experiments of Kahneman and Tversky), though not for decisions from experience.

- Pedro Bordalo, Nicola Gennaioli, and Andrei Shleifer. Salience theory of choice under risk. *The Quarterly Journal of Economics*, 127:1243–1285, 2012.
- Colin Bredenberg, Eero P. Simoncelli, and Cristina Savin. Learning efficient task-dependent representations with synaptic plasticity. *pre-print posted on bioRxiv*, <https://www.biorxiv.org/content/10.1101/2020.06.19.162172v1.full.pdf>, June 2020.
- Matteo Carandini and David J. Heeger. Normalization as a canonical neural computation. *Nature Reviews Neuroscience*, 13:51–62, 2012.
- Avinash K. Dixit and Barry J. Nalebuff. *Thinking Strategically: The Competitive Edge in Business, Politics, and Everyday Life*. W. W. Norton Company, 1991.
- Catherine Donnelly and Paul Embrechts. The devil is in the tails: Actuarial mathematics and the subprime mortgage crisis. *ASTIN Bulletin*, 40:1–33, 2010.
- Ido Erev, Eyal Ert, Alvin E. Roth, Ernan Haruvy, Stefan M. Herzog, Robin Hau, Ralph Hertwig, Terrence Stewart, Robert West, and Christian Lebiere. A choice prediction competition: Choices from experience and from description. *Journal of Behavioral Decision Making*, 23:15–47, 2010.
- Eugene F. Fama. The behavior of stock market prices. *Journal of Business*, 38:34–105, 1965.
- Franz Faul, Edgar Erdfelder, Albert-Georg Lang, and Axel Buchner. G*power 3: a flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavioral Research Methods*, 39:175–191, 2007.
- Cary Frydman and Lawrence J. Jin. Efficient coding and risky choice. *Working Paper*, 2020.
- Xavier Gabaix and David Laibson. Myopia and discounting. *NBER Working Paper No. 23254*, 2017.
- Xavier Gabaix, Parameswaran Gopikrishnan, Vasiliki Plerou, and H. Eugene Stanley. A theory of power-law distributions in financial market fluctuations. *Nature*, 423:267–70, 2003.
- Xavier Gabaix, Parameswaran Gopikrishnan, Vasiliki Plerou, and H. Eugene Stanley. Institutional investors and stock market volatility. *Quarterly Journal of Economics*, 121:461–504, 2006.
- Deep Ganguli and Eero P. Simoncelli. Neural and perceptual signatures of efficient sensory coding. Working paper, arXiv: 1603.00058v1, February 2016.
- Nicola Gennaioli and Andrei Shleifer. A crisis of beliefs. Investor psychology and financial fragility. *Princeton University Press*, 2018.
- Stefano Giglio, Matteo Maggiori, Johannes Stroebel, and Stephen Utkus. Inside the mind of a stock market crash. Working paper, 2020.
- Ahna R. Girshick, Michael S. Landy, and Eero P. Simoncelli. Cardinal rules: Vi-

- sual orientation perception reflects knowledge of environmental statistics. *Nature Neuroscience*, 14:926–932, 2011.
- Paul Glimcher. Understanding the hows and whys of decision-making: From expected utility to divisive normalization. *Cold Spring Harbor Symposia on Quantitative Biology*, 79:169–176, 2014.
- Ralph Hertwig, Greg Barron, Elke U. Weber, and Ido Erev. Decisions from experience and the effect of rare events in risky choices. *Psychological Science*, 15(8):534–539, 2004.
- Daniel Kahneman. Maps of bounded rationality: Psychology for behavioral economics. *The American Economic Review*, 93(5):1449–1475, 2003.
- Daniel Kahneman and Amos Tversky. Prospect Theory: An analysis of decision under risk. *Econometrica*, 47(2):263–291, 1979.
- Bryan Kelly and Hao Jiang. Tail risk and asset prices. *The Review of Financial Studies*, 27:2841–2871, 2014.
- Mel W. Khaw, Paul W. Glimcher, and Kenway Louie. Normalized value coding explains dynamic adaptation in the human valuation process. *Proceedings of the National Academy of Sciences*, 114:12696–12701, 2017.
- Mel W. Khaw, Ziang Li, and Michael Woodford. Cognitive imprecision and small-stakes risk aversion. *Review of Economic Studies*, forthcoming 2020.
- Botond Koszegi and Matthew Rabin. A model of reference-dependent preferences. *Quarterly Journal of Economics*, 121(4):1133–1165, 2006.
- Botond Koszegi and Adam Szeidl. A model of focusing in economic choice. *Quarterly Journal of Economics*, 128:53–104, 2013.
- Alok Kumar. Who gambles in the stock market? *The Journal of Finance*, LXIV(4):1889–1933, 2009.
- Alok Kumar, Jeremy Page, and Oliver Spalt. Religious beliefs, gambling attitudes, and financial market outcomes. *The Journal of Financial Economics*, 102(3):671–708, 2011.
- Simon B. Laughlin. A simple coding procedure enhances a neuron’s information capacity. *Zeitschrift für Naturforschung C*, 36:910–912, 1981.
- Andrew Lo. Adaptive markets: Financial evolution at the speed of thought. *Princeton University Press*, 2017.
- Lars A. Lochstoer and Tyler Muir. Volatility expectations and returns. *NBER Working Paper No. w28102*, 2021.
- Benoit B. Mandelbrot. *Fractales, Hasard et Finance*. Paris: Flammarion, 1957.
- R. David McLean and Jeffrey Pontiff. Does academic research destroy stock return predictability? *Journal of Finance*, 71:5–32, 2016.
- Hyman P. Minsky. *Stabilizing an Unstable Economy*. New York: McGraw-Hill, 1986.

- Camillo Padoa-Schioppa and John A. Assad. The representation of value in the orbitofrontal cortex is invariant for changes of menu. *Nature Neuroscience*, 11(1): 95–102, 2008.
- Camillo Padoa-Schioppa and Aldo Rustichini. Rational attention and adaptive coding: A puzzle and a solution. *American Economic Review*, 104(5):507–513, 2014.
- Allen Parducci. Category judgment: A range-frequency model. *Psychological Review*, 72:407–418, 1965.
- Elise Payzan-LeNestour. Can people learn about ‘black swans’? experimental evidence. *Review of Financial Studies*, 31(12), 2018.
- Elise Payzan-LeNestour, Lionnel Pradier, and Talis J. Putnins. The “waterfall illusion” in financial markets: How risk perception is distorted after exposure to extreme risk. Working paper, November 2020.
- Rafael Polania, Michael Woodford, and Christian C. Ruff. Efficient coding of subjective value. *Nature Neuroscience*, 22:134–142, 2019.
- Matthew Rabin. An approach to incorporating psychology into economics. *American Economic Review*, 103:617–622, 2013.
- Arthur J. Robson and Lorne A. Whitehead. Adaptive cardinal utility. *unpublished paper*, August 2019.
- Aldo Rustichini, Katherine E. Conen, Xinying Cai, and Camillo Padoa-Schioppa. Optimal coding and neuronal adaptation in economic decisions. *Nature Communications*, 8(1208), 2017.
- Alireza Soltani, Benedetto De Martino, and Colin Camerer. A range-normalization model of context-dependent choice: A new model and evidence. *PLoS Computational Biology*, 8(7), 2012.
- Jakub Steiner and Colin Stewart. Perceiving prospects properly. *American Economic Review*, 106:1601–1631, 2016.
- Christopher Summerfield and Konstantinos Tsetsos. Do humans make good decisions? *Trends in Cognitive Sciences*, 19:27–34, 2015.
- Nassim N. Taleb. *Foiled By Randomness: The Hidden Role of Chance in Life and in the Markets*. Penguin Books, London, 2004.
- Leon L. Thurstone. A law of comparative judgment. *Psychological Review*, 34:273–286, 1927.
- Philippe N. Tobler, Christopher D. Fiorillo, and Wolfram Schultz. Adaptive coding of reward value by dopamine neurons. *Science*, 307:1642–1645, 2005.
- Leon Tremblay and Wolfram Schultz. Relative reward preference in primate orbitofrontal cortex. *Nature*, 398:704–708, 1999.
- Christoph Ungemach, Nick Chater, and Neil Stewart. Are probabilities overweighted or underweighted when rare outcomes are experienced (rarely)? *Psychological*

- Science*, 20:473–479, 2009.
- Xue-Xin Wei and Alan A. Stocker. A Bayesian observer model constrained by efficient coding can explain ‘anti-bayesian’ percepts. *Nature Neuroscience*, 18:1509–1517, 2015.
- Xue-Xin Wei and Alan A. Stocker. Lawful relation between perceptual bias and discriminability. *Proceedings of the National Academy of Sciences*, 114:10244–10249, 2017.
- Michael Woodford. Inattentive valuation and reference-dependent choice. Working paper, Columbia University, 2012a.
- Michael Woodford. Prospect theory as efficient perceptual distortion. *American Economic Review*, 102:41–46, 2012b.
- Michael Woodford. Modeling imprecision in perception, valuation and choice. *Annual Review of Economics*, 12:579 – 601, 2020.
- Hang Zhang, Xiangjuan Ren, and Laurence T. Maloney. The bounded rationality of probability distortion. *Proceedings of the National Academy of Sciences*, 36:22024 – 22034, September 2020.

Online Appendix (not for publication)

We provide supplementary information about the experiments documented in Section 2 of the main text. The code to generate the experimental task used in each experiment, as well as the code to run the analyses reported in Section 2, can be downloaded at <https://supplementarymat.weebly.com>. The experimental data for each experiment are also provided there.

Experimental Paradigm

On each of 1,128 trials, the participant is to discriminate between the shades of grey in two rectangles by clicking on the rectangle that looks darker, or on an “=” icon displayed in the middle of the screen if the rectangles look of the same color (see Figure 3 in the main text). The imparted time to reply is 2 sec in the main, placebo, and 5-trial adaptation experiments, and 4 sec in the experiment with double response time (run as a robustness check, as explained in the main text). In the majority (87%) of the trials, the shades of the two rectangles are adjacent, i.e., shade values x and $x + 1$ are displayed, where x is randomly drawn from a scale of grey with 12 different shades (1: very light; 12: extremely dark).⁴¹ In the remaining trials, the

⁴¹Nuances of grey and colors are coded in HTML format `rgb(r,g,b)`, with `r` = Red, `g` = Green, and `b` = Blue, each one having a value comprised between 0 and 255. Greys are coded with `r=g=b`,

shades of the two rectangles are the same.

Stochastic Structure of the Task The core of our experimental strategy is summarized in Figure 4 of the main text. The participant goes through an adaptation phase whose duration is varied across experiments (main and double response time experiments: 40 trials; placebo test: 3 trials; 5-trial adaptation experiment: 5 trials) in which on each trial shade value x is drawn from a normal distribution with mean m and standard deviation s ,⁴² immediately followed by a “test trial” with shade x' randomly drawn in the range of values located at least three standard deviations from m . Each value x' in turn defines an adaptation phase in which the shade value presented to the participant is normally distributed around x' (standard deviation is either 1 or 2; choice is random), followed by a “control trial” in which the shade value presented to the participant is exactly like in the test trial (i.e., either x' and $x' + 1$, or x' in both rectangles).

We randomly select 12 values for m and s in the sets $\{3, 4, \dots, 10\}$ and $\{1, 2\}$ respectively. With such a design, our main statistic of interest (the difference between the accuracy level in the control vs. test trials across pairs, see the main text) would have 12 observations per participant. To double the number of observations without increasing the total number of trials, we impose that the shade value presented to the participant in the last trial of the adaptation phase around m be also presented in the trial immediately following the control trial. By design, the value is an outlier in the later case, whereas it is not in the former. We compare accuracy in the two cases across the 12 pairs. This gives us overall 24 observations per participant for our main statistic of interest.

We impose that the values drawn for m be evenly distributed across the possible values $\{3, \dots, 10\}$ and constrain our choice of values such that the shade presented at the beginning of an adaptation phase around m is never an outlier for the participant. We also ensure that the proportion of same-shade trials is the same (13%, as noted above) across all trial types (the trials within the adaptation phases and test/control trials), to allow comparisons across trial types.⁴³ Apart from these constraints, the design is fully randomized. For example, the control sequence depicted in Figure 4 of the main text (adaptation phase around x' followed by control trial) does not have to follow the sequence depicted in the top graph (adaptation phase around m followed

from 0 (black) to 255 (white). In the experiment, shade of grey $\text{rgb}(a,a,a)$ is compared to shade $\text{rgb}(a+N,a+N,a+N)$ or to shade $\text{rgb}(a-N,a-N,a-N)$, with $N=21$.

⁴²The drawn value is rounded. For instance, if the number drawn is 4.3, the 4th-shade on our scale of grey is chosen.

⁴³One expects the base rate accuracy to be decreased in the same-shade trials (irrespective of trial type) inasmuch as same-color trials occur with low probability in this task, so in principle participant replies should be biased against the “=” reply.

by test trial). The order of appearance of the sequences is random.

Rationale for this Experimental Design Our main motivation for choosing 2 seconds as the allowed time to provide a reply on each trial in the main experiment is to maximize statistical power in the analysis, by maximizing the number of trials for each participant under the standard duration limits for this kind of task (30 minutes should not be exceeded), while keeping in mind task feasibility (excessive time pressure leads to random choice in the participants).⁴⁴ Participant behavior in the main experiment validates this parameter choice inasmuch as participants usually reply well within the allowed time, the frequency of missed trials is very low overall, and doubling the allowed time to provide a reply does not change our main findings (see Section 2 of the main text and the supplementary statistics in Section 6.2 below).

Following the same logic (maximizing statistical power), we prefer the current paradigm over the 2-choice discrimination version of the task with no same-color trials. The current task is reportedly less boring for the participants (so participant data are of better quality).⁴⁵ Moreover, in pilot sessions, some “ceiling effect” in accuracy level emerges in the 2-choice version because it is too easy. Our choice of a 12-point scale for our scale of grey follows similar logic. Using a coarser scale would potentially decrease power by making the task too easy. Using a finer scale would increase randomness in participant replies.⁴⁶

Another aspect of reducing randomness in participant replies is to have participants pay attention on each trial. To that purpose, we use a special experimental procedure—spelled out next.

Experimental Procedure Undergraduate students from the University of New South Wales register online to participate in the experiment.⁴⁷ Upon arrival at the lab, the participants seat in front of an individual computer. The experimenters tell them that they can potentially earn a significant amount of money from the task, but only if they familiarize themselves with the task instructions and the task interface.

The participants then watch the following instructions for the task for 15 minutes (Figure 9). These instructions point out that each correct reply yields \$0.1, each incorrect reply leads to a loss of \$0.25, each missed trial leads to a loss of \$1, and that

⁴⁴We tested different time parameter values in pilot sessions; both participant choice data and participant oral reports after the pilot suggest that allowing 2 seconds to provide a reply is a reasonable trade-off between sample size and task feasibility.

⁴⁵When asked in a post-task debrief of a pilot to compare the 2-choice and 3-choice versions of the task, participants report the 2-choice version to be quite boring.

⁴⁶We do not claim that our choice is optimal but our findings show it is good enough to identify the outlier blindness effect which is our main purpose in this experiment.

⁴⁷Registration is done through the ORSEE recruitment system (<http://www.orsee.org/web/>).

the participant receives at the end of the experimental session all the net accumulated outcomes from the task.⁴⁸

At the end of the instructions, the participants are asked if they wish to complete a 3-minute training session in which they play a few trials of the experimental task to familiarize themselves with the task interface. Given the experimenters' strong emphasis on the importance of familiarizing oneself with the task interface to have a chance to fare well in the task, the preference to skip the training in a given participant signals that the participant is not particularly interested in doing well in the task. Six participants choose to skip the training. The experimenters kindly ask them if they would be content with the \$5 show-up reward and not performing the task, as their data will not be used in any case.⁴⁹ All six participants are keen to take the \$5 straight away and accordingly leave the lab.

Just before the participants start performing the task, they are briefed again by the experimenters on the distinctive nature of the experiment, the payment rule that is used in particular. The experimenters also stress that the luminance setting has been adjusted on each machine before the beginning of the experimental session and that it has been locked so that luminance cannot be changed during the experiment.⁵⁰ Subsequently, the participants complete one run of the task, which lasts for approximately 25 minutes.

⁴⁸In case of negative earnings, the participant ends up with the \$5 show-up reward which is given irrespective of task performance (as per the lab protocol).

⁴⁹The experimenters explain why in a short debriefing. ("The fact that you skipped the training makes it unlikely that you will be able to perform well in the task, and it is important that the data be of high-quality for scientific purposes. Given this, we thought you may be content with getting the \$5 show-up reward straightaway and leaving the lab?")

⁵⁰The monitors in the lab are the HP E272q. We use the factory default settings (brightness: 90; contrast: 80; dynamic contrast: off; black stretch: off). To prevent participants from changing settings during the experiment, the experimenters enable the button lockout feature of the monitors beforehand.

Thanks for accepting to participate in “The Hue Task”!

On each of 1,128 trials, you will be asked to discriminate between two hues of grey, see picture below. In some of the trials the two rectangles will be of the same color. In some others, they will be of different color. If the two rectangles look of the same color, click on the "=" icon in the middle of the screen. If the rectangles look of different color, click on the rectangle that looks darker.

Which one is darker?

=

Trial: 1 / 1128

We (the experimenters) want to reward very significantly high performance in this task. You will earn \$0.10 per correct reply and will lose \$0.25 per incorrect reply. Since you will be playing 1,128 trials overall, this means you can earn a lot of money potentially—more than \$100—if you perform well in the task.

Please Note:

- You will lose \$1 if you fail to reply within the imparted time on a given trial (2 sec — note the pace of the game is high).
- The task is quite long (about 30 minutes overall). There will be a short break after the first 15 minutes.

The task therefore requires you to keep the pace and pay attention for a prolonged period of time. To familiarize you with the task settings before performing the task, we offer you the opportunity to do a 3-minute training session in which you will be playing a few trials of the task; note these trials will NOT be counting for your final payment (your replies won't be recorded).

Please indicate your choice:

Figure 9: **Instructions for the task used in the main experiment** The text of the instructions for the placebo and 5-trial adaptation experiments is the same except for the information concerning the total number of trials (there are 960 trials in the placebo experiment and 1,000 trials in the 5-trial adaptation experiment). The text of the instructions for the experiment with double response time is the same except for the information about the time allowed to provide a reply on each trial (which is 4 seconds in the double-response time experiment).

Supplementary Statistics

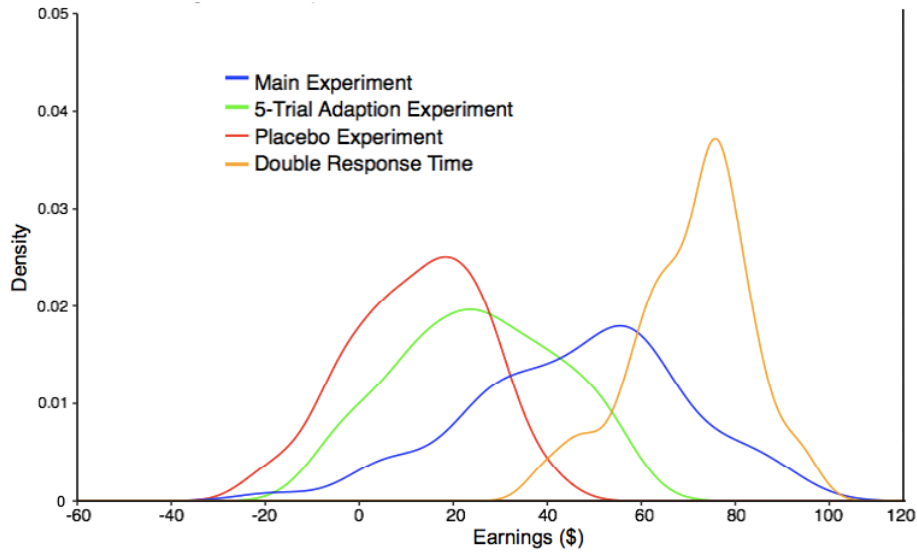


Figure 10: **Earning distribution across participants in each experiment.** For each participant the earnings are computed as the net accumulated outcomes at the end of the task. Note: The lower earnings in placebo and 5-trial adaptation experiments (relative to main and double response time experiments) is expected given the smaller number of trials in those experiments (960 and 1,000 vs. 1,128, see the legend of Figure 9).

Table 3: **Number of missed trials in each experiment.** The mean, standard deviation, min and max are derived across participants, for each experiment.

	Missed Trials			
	Mean	Std. Dev.	Minimum	Maximum
Main Experiment	6.17	10.04	0.00	50.00
5-Trial Adaptation Experiment	7.52	10.32	0.00	47.00
Placebo Experiment	9.29	5.82	0.00	27.00
Experiment with Double Response Time	1.52	1.89	0.00	7.00

Table 4: **Mean accuracy according to trial type.** Numbers in parenthesis: standard error of the mean (sem). The main thing to note is that accuracy is lower in the test trials than in the control trials in all the experiments except for the placebo experiment, as predicted by outlier blindness theory (see the main text for the statistical tests). Also of note, accuracy is higher in the trials preceding test trials than in the trials preceding control trials. This simply reflects the fact that by design, the shade values presented to the participants are on average more extreme in the latter than in the former [due to programming constraints related to the definition of an outlier in our test: a value at least 3 standard deviations from m , see above], and it is harder to discriminate among shade values located at the extremes of the scale of grey relative to middle range values. This decreased sensitivity of perception for extreme values is a well-known fact in the field of psychophysics (see, for example, the celebrated “Weber’s Law”). Our test of outlier blindness fully controls for it—recall that by design the value presented to the participant is *the same* at the test and control trials.

	Mean Accuracy						
	Overall	Same Color Trials	Different Color Trials	Trials within Adaptation Phase			
				Before Test Trials	Before Control Trials	Test Trials	Control Trials
Main Experiment	0.83 (0.01)	0.44 (0.02)	0.88 (0.01)	0.87 (0.01)	0.81 (0.02)	0.61 (0.02)	0.79 (0.01)
5-Trial Adaptation Experiment	0.80 (0.01)	0.41 (0.03)	0.83 (0.01)	0.86 (0.01)	0.84 (0.01)	0.72 (0.01)	0.75 (0.01)
Placebo Experiment	0.77 (0.01)	0.29 (0.03)	0.79 (0.01)	0.86 (0.01)	0.85 (0.01)	0.73 (0.01)	0.73 (0.01)
Experiment with Double Response Time	0.90 (0.01)	0.43 (0.03)	0.95 (0.01)	0.92 (0.01)	0.89 (0.01)	0.72 (0.02)	0.84 (0.01)

Table 5: **Mean response time according to trial type.** Numbers in parenthesis: standard error of the mean (sem).

	Response Time (in seconds)						
	Overall	Same Color Trials	Different Color Trials	Trials within Adaptation Phase			
				Before Test Trials	Before Control Trials	Test Trials	Control Trials
Main Experiment	0.90 (0.02)	1.08 (0.02)	0.87 (0.02)	0.86 (0.02)	0.92 (0.02)	1.13 (0.02)	0.89 (0.02)
5-Trial Adaptation Experiment	1.03 (0.03)	1.16 (0.03)	1.02 (0.03)	0.99 (0.03)	1.01 (0.03)	1.14 (0.03)	1.02 (0.03)
Placebo Experiment	1.09 (0.03)	1.24 (0.02)	1.08 (0.03)	1.03 (0.03)	1.04 (0.03)	1.16 (0.02)	1.07 (0.03)
Experiment with Double Response Time	1.20 (0.04)	1.81 (0.07)	1.12 (0.04)	1.15 (0.04)	1.23 (0.04)	1.59 (0.05)	1.21 (0.04)