"Outlier Blindness": Efficient Coding Generates an Inability to Represent Extreme Values

By Elise Payzan-LeNestour and Michael Woodford[*]

Draft: December 20, 2018

How do people perceive outliers? Building on a well-established theory from neuroscience, we conjecture that people are inherently hampered in the way they perceive outliers because the human brain has been designed to devote neural activity to representing the most probable values at the expense of the improbable ones. We find support for this conjecture in a series of controlled laboratory experiments. (JEL: C91, D87)

Keywords*: Neuroeconomics, Tail risk, Efficient coding, Adaptation, Decision making under uncertainty.*

[*] 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.

Introduction

How do economic agents perceive extreme values or outliers—so-called "black swans" (Taleb, 2007)? This question is crucial given the ubiquity of tail risk in modern financial markets, which has been long been established.[1] Under tail risk, black swans are not rare enough for the agents to be able to ignore them without negative consequences. In this study, we conjecture that people are inherently hampered in the way they perceive black swans because the human brain has been designed to devote neural activity to representing the most probable values at the expense of the improbable ones. We document strong experimental evidence for this conjecture, which builds on a well-established theory from neuroscience.

According to neuroscientists, devoting neural activity to representing probable values at the expense of improbable values is one key aspect of minimizing uncertainty in perception under the fundamental constraints imposed by Nature. The representational capacities of the brain are indeed limited (we have a finite number of neurons and a finite number of possible spike outputs of each neuron), and if a neuron's limited outputs were allocated evenly to represent the potentially infinite number of possible values of a stimulus, then that neuron's activity would allow for little if any discrimination between values. The optimal solution to the problem—"efficient coding"—consists of ensuring that neurons have access to information indicating that some values are more likely to occur than others, and allocating most of their spike outputs to representing the most probable values at the expense of the improbable values.[2] This way of allocating neural activity may generate an epiphenomenon, which we refer to as "outlier blindness," in which economic agents are unable to properly represent values that are highly improbable relative to the range of values they are expecting.

Of course a prerequisite for the occurrence of outlier blindness is that the agents hold expectations. From the perspective of the agents, the range of expected values is the range to which the agents have been exposed—or, in the language of neuroscientists, "adapted"—in the recent past. In our modern environments which are notoriously unstable owing to the frequent occurrence of regime shifts or jumps in values, a precondition for outlier blindness is thus that the human brain learns sufficiently quickly the statistics

---

[1]For example, Mandelbrot (1957), Fama (1965), Gabaix et al. (2003), Gabaix et al. (2006), Kelly and Jiang (2014).

[2]For example, Laughlin (1981), Tobler et al. (2005), Wei and Stocker (2015), and Ganguli and Simoncelli (2016). Similar principles have been applied to models of imprecision in economic decision making by Robson (2001), Netzer (2009), Woodford (2012), and Polania et al. (2019).

of the environment, so that the agents hold expectations even though those expectations are to change all the time.[3] Whether the brain learns quickly and what "quickly" means is an empirical question that we attempt to answer here.

To test our conjecture that efficient coding generates a phenomenon of outlier blindness in economic decision-makers (henceforth, "outlier blindness hypothesis"), we ask university students to perform for significant amounts of money a perceptual task in which on each trial, they have to discriminate between two adjacent shades of grey. The core of our experimental strategy consists of assessing the accuracy of subject perception of outliers, which we define as shade values that fall at least three standard deviations away from the range of values to which the subject has been exposed (adapted) in the previous 40 trials, relative to subject accuracy in control test trials in which the subjects are presented with the exact same shade values but this time they're not outliers as they fall within the range of values to which the subject has been adapted in the previous 40 trials. This way we control for important confounds inherent in the presentation of extreme values. For instance, in light of the classical "*Weber-Fechner's law*" of diminishing sensitivity (in which a change in stimulus value is harder to perceive if it occurs at a higher level), one expects perceptual accuracy to be inherently decreased for extreme values simply because it is harder to discriminate between values located at the extremes on the scale of grey. Therefore, merely comparing subject accuracy when discriminating between extreme shade values to their mean accuracy when discriminating between the other possible shades of grey on the scale features an increased risk of Type I error. In contrast, our test, by controlling for the diminishing sensitivity phenomenon, is specifically designed to isolate the outlier blindness effect caused by efficient coding.

An important aspect of the method used in the present study thus consists of designing a non-economic task that involves the perception of a simple sensory stimulus in order to learn about how economic agents perceive extreme values in the field. We refrain from designing an economic task as our motivation is to isolate the nature of perception per se, and in an economic task, separating purely perceptual factors from behavioral factors related to risk attitude is impossible. Our design choice also reflects a growing consensus in neuroscience that efficient coding is a universal neural principle.[4] As such, it is *stimulus invariant*, i.e., it applies similarly to the perception of a wide range of stimuli, from simple sensory stimuli such as the one used in the present task, to complex stimuli

---

[3]Prior field research has identified the presence of regime shifts or jumps as a major feature of the returns generating process. See Ang and Timmermann (2011) for a survey of the literature.

[4]For example, Carandini and Heeger (2012) and Glimcher (2014).

such as the value of economic options (Khaw et al., 2017). This implies that one can learn about how economic agents perceive financial returns from the findings obtained in purely perceptual tasks involving simple stimuli.

We find that perception accuracy is markedly decreased for outliers, as predicted by outlier blindness hypothesis. To strengthen the evidence for the latter and to rule out any possibility that such decrease reflect some experimental artefact, we run a placebo test in which the experimental task is the same as the original task except for the duration of the adaptation phase prior to the presentation of the outliers, which lasts for only 3 trials in the placebo test (versus 40 trials in the original experiment). So by design, in the placebo test the brain does not have time to learn which values are more likely to be presented, and hence outlier blindness should *not* occur. Consistent with this prediction, the original effect vanishes in the placebo test.

We further find that increasing the length of adaptation to 5 trials is enough for the outlier blindness effect to reappear, which shows that prolonged adaptation is not a prerequisite for the emergence of the effect. This finding suggests that the human brain is a quick learner—put it differently, our perceptual system is quick to adapt—so in relation to the foregoing note, one should not expect market instability to protect financial decision-makers against the outlier blindness bias.

One striking aspect of our finding that people are unable to properly discriminate between extreme values owing to the outlier blindness bias is that this bias runs in *opposite* direction to what one may consider the most standard effect associated with the perception of improbable values, namely the classic "*oddball effect*." In the oddball effect, information processing of unexpected stimulus values is improved—and hence perception accuracy is increased—relative to the perception of expected values because the unexpected is salient ("attention-grabbing") per se due to its novelty or special significance to the subject.[5] Our findings show that the outlier blindness effect dominates the oddball effect.

By changing stimulus distribution every 40 trials, our task is meant to capture the instability of real-world financial markets which arises from the frequent occurrence of regime shifts and the like, both at short and long timescales. While the short timescale feature of our design (the imparted response time on each trial is 2 sec, see Section I.A.) directly maps into contexts of short timescale economic decision-making such as intra-day trading, there is a growing consensus in neuroscience that common principles apply over

---

[5]For example, Squires et al. (1975), Tse et al. (2004), Ferrari et al. (2010).

different (short and long) timescales.[6] Therefore, having the scale invariance feature of human psychology in mind, one expects the phenomenon of outlier blindness to be similar over different timescales, suggesting that our findings apply similarly to the perception of jumps at short (intra-day or daily) horizons and the perception of black swans in a macroeconomic context.

This paper relates to several strands of literature. First, a growing literature in neuroeconomics seeks to understand the implications for economic decision making of the hypothesis that decisions are based on internal representations that, while imprecise, are structured in a way that makes the best possible use of finite cognitive resources. For instance, Steiner and Stewart (2016) show how imprecise perception exhibits a key feature of Prospect Theory's probability weighting function (namely, overweighting of small probability events and corresponding underweighting of high probability events). In the same vein, Gabaix and Laibson (2017) show how imprecise perception generates apparent temporal discounting in perfectly patient agents. Within that literature and closest to the present study, recent work attempts to understand the implications of efficient coding theory for economic decision-making. Khaw et al. (2017) provide direct evidence that efficient coding applies not only to the perception of natural images (e.g., Simoncelli (2003)) but also to complex stimuli such as the value of economic options. Payzan-LeNestour et al. (2018) report findings consistent with the idea that efficient coding generates a perceptual bias ("after-effect") in risk perception. Frydman and Jin (2018) show that efficient coding provides an explanation for the variation of risk taking across environments.

The outlier blindness effect described in this study is reminiscent of the well-known phenomenon of "inattentional blindness" (Mack and Rock, 1998) in which the decision maker fails to notice the unexpected (see Simons and Chabris (1999) for a classic example). The source of the two biases is different: outlier blindness is a perceptual bias generated by efficient coding whereas inattentional blindness is an attentional bias reflecting sparse attention (e.g., Sims (2003), Mackowiak and Wiederholt (2009), and Gabaix (2014)).

Last but not least, our study adds to the finance literature on tail risk. As stressed above, tail risk has long been established as an essential feature of financial markets. However, the focus on the implications of tail risk for the decision-making of financial

---

[6]See Brown et al. (2007) for a survey of the literature showing that many memory phenomena appear similar over a wide range of timescales.

agents is recent.[7] Of particular importance, the literature on neglected risks provides profound insights into how agents undestimate the probability of black swans as a result of both the availability heuristic (extreme events that occurred far away are harder to recall from memory; see, e.g., Gennaioli et al. (2012), Gennaioli et al. (2015), and Jin (2015)) and the representativeness heuristic (upon receiving the news that an extreme event has taken place, agents recall black swan outcomes that such news renders more likely than before, i.e., the news is "diagnostic" of black swan outcomes; see Gennaioli and Shleifer (2018)). By sheding light on how agents respond to regime shift news, our study may help in building a narrative of the 2008 financial crisis within Gennaioli and Shleifer (2018)'s diagnostic expectations framework. Diagnostic expectations theory can readily account for both the tail risk neglect prior to the news of Lehman collapse in 2008, and the investor panic that ensued from such news: the left tail of outcome distribution was unrepresentative before Lehman and became representative thereafter. However, the complete lack of reaction of investors to the regime shift news of the housing bubble burst in 2006 and ensuing "domino effect" during the pre-Lehman period is puzzling at first glance.[8] During this period, which has been referred to as "the quiet period" (Gennaioli and Shleifer, 2018), the evidence for a major regime shift was compelling, yet investors ignored it and continued to neglect tail risk (including when daily realized volatility started to rise above 20% from the summer of 2007, with peaks above 40%). There are two explanations for this apparent paradox. The first is that investors chose to ignore the regime shift news in 2006-2007 (Lo, 2017). A second explanation is that they missed it, as a result of outlier blindness, because their perceptual system was not yet adapted to the new regime. However they didn't miss Lehman collapse news because the news came after the aforementioned series of warning signs in 2006-2007, so the occurrence of an extreme event such as Lehman collapse was not totally unexpected. This would explain why Lehman collapse news was diagnostic of the left tail whereas the extreme events that took place during the pre-Lehman period were not.

---

[7]For example, Gennaioli et al. (2012), Gennaioli et al. (2015), Jin (2015), Goetzmann et al. (2017), Gennaioli and Shleifer (2018), Payzan-LeNestour (2018).

[8]In Chapter 9, Lo (2017) describes how the regime shift that started the financial crisis happened as early as in June 2006, when home prices peaked and then started to fall. Rising interest rates and decreasing prices led to defaults in homeowners, and the ensuing stop in the flow of money that supported the pools of MBS securities was a disaster leading the "dominos" of the economy to fall one by one, starting by the bankrupcy of risk-taking mortgage lenders.

I.   Lab experiment

A.   *Experimental Design*

Task Description

The principle of the task is that on each of 1,128 trials, the subject is to discriminate between the shades of grey of two rectangles by clicking within 2 seconds 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 1). We construct a scale of grey with 12 different shades (1: very light; 12: extremely dark). In a typical trial, the shades of the two rectangles are adjacent, i.e., shade values x and $x + 1$ are displayed, where x is randomly drawn from the scale of grey as described next. The two shades are the same in 13% of the trials. The subjects are told the frequency of same-shade trials in the task instructions (see Appendix IV.A).

The logic of our design is to have each subject go through a 40-trial "adaptive sequence," in which on each trial shade value x is drawn from a normal distribution with mean m and standard deviation s (the values of m and s are randomly selected in the sets $\{3, 4, ..., 10\}$ and $\{1, 2\}$ respectively),[9] immediately followed by an "adaptive test trial" in which the shade (x') is randomly drawn in the range of values located at least three standard deviations from m. So by design, the shade presented at the adaptive test trial is an outlier as it falls outside the range of values to which the subject has been adapted in the adaptive sequence. Each outlier value x' defines a 40-trial "control sequence" in which the shade value on each trial is normally distributed around x' (standard deviation is either 1 or 2; choice is random). The control sequence systematically ends with a "control test trial" in which x' is presented. So each adaptive sequence is paired with a control sequence and both sequences are immediately followed by the same test trial.

There are 12 adaptive sequences (so 12 sequence pairs). The order of appearance of all sequences is randomized. So generally, the adaptive and control sequences of a given pair do not follow each other, and we do not impose either that the adaptive sequence be presented before the control sequence.[10]

For each subject, we compute the mean accuracy level in the control test trials averaged

---

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

[10]The only constraints that we impose are that the means of two successive sequences be as close as possible (to ensure the absence of outliers at the beginning of a sequence) and that each m value from the set of possible values $\{3, 4, ..., 10\}$ be used in at least two of the adaptive sequences.

across the 12 sequence pairs, as well as the mean accuracy level in the adaptive test trials. Accuracy is defined as the fraction of correct replies. Our main statistic of interest is the difference between the accuracy level in the control vs. adaptive test trials. To double the number of observations for our main statistic without increasing the total number of sequence pairs, we add to our design a symmetrical feature whereby within each sequence pair, the shade value presented at the last trial of the adaptive sequence is also presented immediately after the control test trial. When presented within the adaptive sequence, the value is not an outlier whereas when presented after the control test trial, it is an outlier as by design it falls at least three standard deviations away from the mean value to which the subject has been adapted during the control sequence. The last trial of the adaptive sequence thus serves as a control test trial w.r.t the adaptive test trial that is put immediately after the control test trial. So for each subject, we have 24 observations for our main statistic of interest.

Outlier blindness hypothesis predicts a significantly decreased accuracy in the adaptive test trials compared to the control test trials. The differential accuracy comes from the fact that by design, the shade value presented at the adaptive test trial is improbable from the perspective of the subject inasmuch as it is well outside the range of values the subject is expecting, whereas when presented in the control test trial, the same shade value is probable from the perspective of the subject (it is within the range of values the subject is expecting). As explained in the introduction, efficient coding theory predicts that perception accuracy is decreased for improbable values (e.g., Tobler et al. (2005) and Woodford (2012)).

Of note, the proportion of same-shade trials is the same (13%, as noted above) across all trial types (the trials within the adaptive sequences, the trials within the control sequences, and the test trials). This is important to allow comparisons across trial types, as one expects the base rate accuracy to be decreased in the same-shade trials (irrespective of trial type) inasmuch as the subjects know that same-color trials occur with low probability in this task, so in principle their replies should be biased against the "=" reply.

Why the current task specification

As stressed in the Introduction, our motivation for designing a non-economic task in order to study the perception of financial returns in economic agents is motivated by a growing consensus in neuroscience that efficient coding applies similarly to the
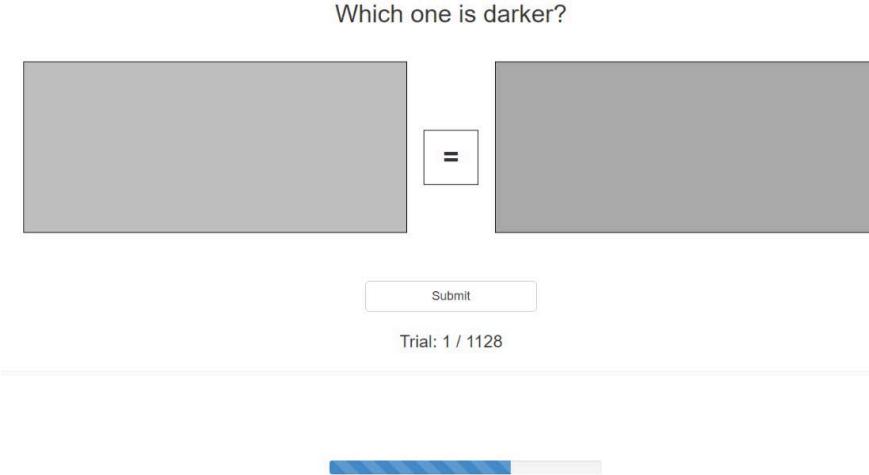
FIGURE 1. **User interface of the experimental task.** THE TASK IS A 3-CHOICE DISCRIMINATION TASK IN WHICH ON EACH TRIAL THE SUBJECT IS TO CLICK ON THE RECTANGLE THAT LOOKS DARKER, OR ON AN "=" ICON DISPLAYED IN THE MIDDLE OF THE SCREEN IF THE RECTANGLES LOOK OF THE SAME COLOR. THE TIME ALLOWED TO MAKE A REPLY IS 2 SECONDS (REMAINING TIME IS INDICATED THROUGH A TIMER AT THE BOTTOM OF THE SCREEN).

perception of simple sensory stimuli and the perception of complex stimuli including economic valuations. This suggests that one can apply findings from the former context to the latter. By using a simple perceptual task rather than an economic decision-making task, we avoid potential confounds related to risk attitude which could not be easily controlled for in an economic decision-making task (fully disentangling risk preferences and beliefs is not possible).

Our main statistic of interest allows us to isolate the outlier blindness effect by ruling out any confound related to the extreme-value effect as explained in the Introduction (the fact that outliers are more likely to be located at the extremes of the scale, and it is harder to discriminate between shade values located at the extremes). If accuracy is inherently decreased when discriminating between extreme values, this affects both adaptive and control test trials in the exact same way in this task, since adaptive and control test trials feature the same shade value by design. Therefore, any differential accuracy between the adaptive and control test trials cannot be attributed to the extreme-value effect.

Our main motivation for choosing the current task settings is to maximize statistical power in the analysis. For instance, our choice of 2 seconds as the allowed time to make a choice on each trial is to maximize the number of trials for each subject 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 subjects). We tested different time parameter values in pilot sessions; both subject choice data and subject oral reports on the task after the pilot suggest that allowing 2 seconds to make a choice is a reasonable trade-off between sample size and task feasibility. Subject behavior in the experiment also validates this parameter choice inasmuch as subjects usually reply well within the allowed time, the frequency of missed trials is fairly low overall, and doubling the allowed time to make a choice does not change our main findings (more in the Results section). All this suggests time pressure is not an issue in this experiment.

Likewise, our motivation for designing a 3-choice discrimination task (the default would be a 2-choice discrimination task with no same-color trials) is to maximize statistical power by averting some "ceiling effect" in accuracy level which would emerge if the task were too easy.[11] Our choice of a 12-point scale for our scale of grey follows similar logic. Using a coarser scale would potentially decrease our power by making the task too easy. Using a finer scale would increase randomness in choice at the test trials (the finer the scale, the harder it is to discriminate between extreme values). We tried different degrees of coarseness in pilot sessions and arrived at the present scale.[12]

Last but not least, one important aspect of reducing noise in the data is to ensure that the subjects do pay attention on each trial. To that purpose, we use a special experimental procedure—spelled out next.

B.   Experimental procedure

Sixty-nine undergraduate students from the University of New South Wales register online to participate in the experiment.[13] Upon arrival at the lab, they watch the online instructions for the experiment for 15 minutes (see Appendix IV.A). At the very end of the instructions, they complete a 3-minute training session in which they play a few trials of the experimental task to familiarize themselves with the task interface (more on this below), after which they are briefed again by the experimenter on the distinctive nature of the experiment, the payment rule that is used in particular. The experimenter re-emphasizes 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 the subject receives at the end of the

---

[11] We ran a pilot session of the 2-choice variant of the task with no same-color trials and found the overall accuracy was very high.

[12] We 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 focus in this study.

[13] Registration is done through the ORSEE recruitment system (http://www.orsee.org/web/).

experimental session all the net accumulated outcomes from the task.[14] The experimenter also stresses 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.[15] Subsequently, the subjects complete one run of the task, which lasts for approximately 25 minutes.

To minimize noise in our data by ensuring that the subjects do pay attention on each trial, we take two steps. First we provide our subjects with high monetary incentives by rewarding high performance in the task through high payoffs. As indicated above, each correct reply yields $0.1 and the subject performs 1,128 trials overall, so high performers can potentially end up with more than $100. This feature of the experiment is well emphasized in the task instructions. Eight subjects earn more than $70, 32 earn more than $50 (mean: 43.8; median: 50.10; mode: 45.65; std: 36.7; see Appendix IV.B). In addition to rewarding high performance in the task through high payoffs, an important aspect of incentivizing the subjects consists of using a "pay all" payment rule (we pay for the outcome from every decision made) rather than the alternative "pay one" approach, which pays for the outcome of only a subset of the choices made. The pay all approach averts the issue of diluted incentives which arises if, for instance, the subjects are susceptible to the "disjunction effect" or the like.[16]

The second key feature of the experimental procedure used in this experiment consists of using the online instructions not only in the usual way (i.e., to acquaint the subjects with the experimental task) but also as a screening device in order to screen out subjects that are not suitable for the task. Specifically, at the very end of the online instructions, the subjects are to decide whether they wish to do the aforementioned short training session to get familiar with the task interface before proceeding to the test. The decision to skip the training, which is implemented by clicking on a button "I wish to skip the training session and go directly to the task" (see Appendix IV.A), results in being immediately excluded from the experiment.[17] The logic here is that the decision of a given subject to skip the training session provides a reliable signal that the subject is not highly motivated to do well in the task. Six subjects choose to skip the training and are therefore excluded from the experiment.

---

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

[15]The monitors in the lab are the HP E272q. We used the factory default settings (brightness: 90; contrast: 80; dynamic contrast: off; black stretch: off). To prevent participants from changing settings during the experiment, the experimenter enabled the button lockout feature of the monitors beforehand.

[16]For example, Tversky and Shafir (1992), Shafir and Tversky (1992), Shafir (1994).

[17]The experimenter gives the excluded subjects the $5 show-up reward.

II.   Results

A.   Findings of the Original Experiment

Across subjects, we find that accuracy is significantly lower in the adaptive test trials than in the control test trials, as predicted by outlier blindness hypothesis. Our statistic of interest (the accuracy difference between control and adaptive test trials, as explained above) averaged across subjects is significantly positive according to our statistical tests (e.g., paired t-test based on the individual statistics of interest as defined above: $t = 11.7$; $p \sim 0$, two-tailed); see Figure 2 and Figure 11 in Appendix IV.C. Strikingly, the statistic is positive for almost all (60 out of 63) the subjects; see Figure 3.

The mean accuracy level across all subjects and trial types is 0.83. It is is markedly decreased in the same-shade trials (0.44), which as explained above is expected given the nature of the task. Accuracy is significantly higher on average within the adaptive sequences than within the control sequences (0.87 vs. 0.81), which comes as no surprise in light of the aforementioned extreme-value effect.[18] 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). The percentage of missed trials is very low on average (0.005%). Appendix IV.B provides more details on the main descriptive statistics related to missed trials, response times, subject accuracy, and subject earnings. Appendix IV.C provides more details on the main statistical tests.

B.   Follow-Up Experiments

Placebo Experiment

An obvious prerequisite for the outlier blindness effect is that the agent holds expectations. An "outlier" is indeed defined with respect to the expectation held by the agent. If the agent expects nothing in particular, the notion is meaningless (there are no outliers from the perspective of the agent). According to efficient coding theory, expectations are set through adaptation. Applied to the current experiment, what this means is that following each 40-trial sequence, the subjects expect the range of shade values to which

---

[18]As explained above, the extreme-value effect refers to the fact that it is harder to discriminate between shades located at the extremes of the scale. By design, the shades that are presented within the control sequences are more extreme on average than those presented within the adaptive sequences. So the extreme-value effect implies that accuracy be decreased on average within the control sequences. See Figure 13 in Appendix IV.C for more details.
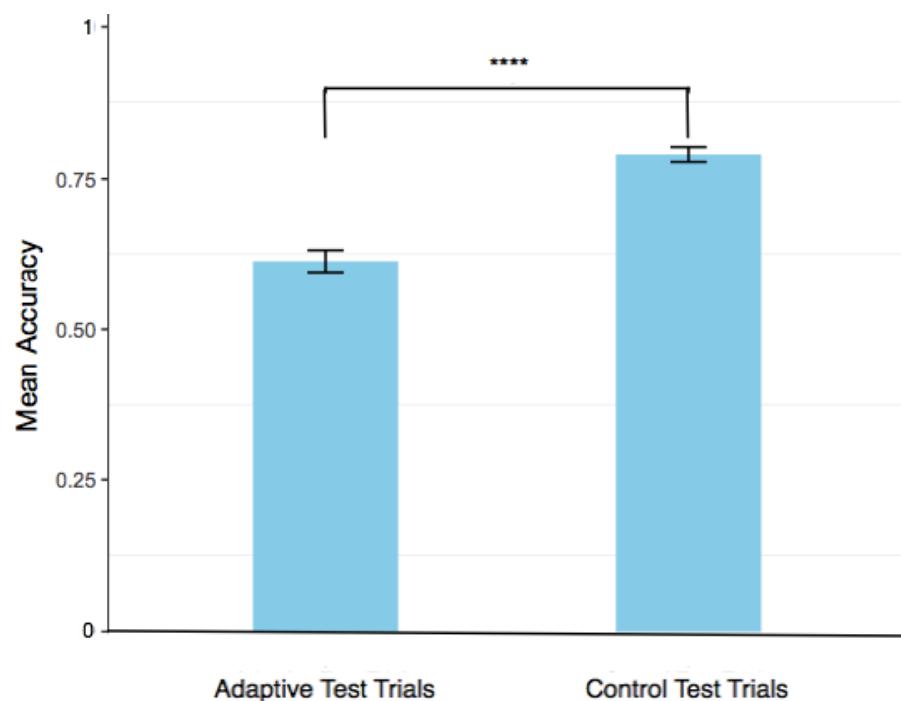
FIGURE 2. **Accuracy in adaptive test trials (left) and control test trials (right).** HEIGHTS OF BARS INDICATE THE MEAN ACCURACY AVERAGED ACROSS THE 63 SUBJECTS. ACCURACY IS DEFINED AS THE FRACTION OF CORRECT REPLIES. LINE SEGMENTS INDICATE SEM. **** $p < .00001$.
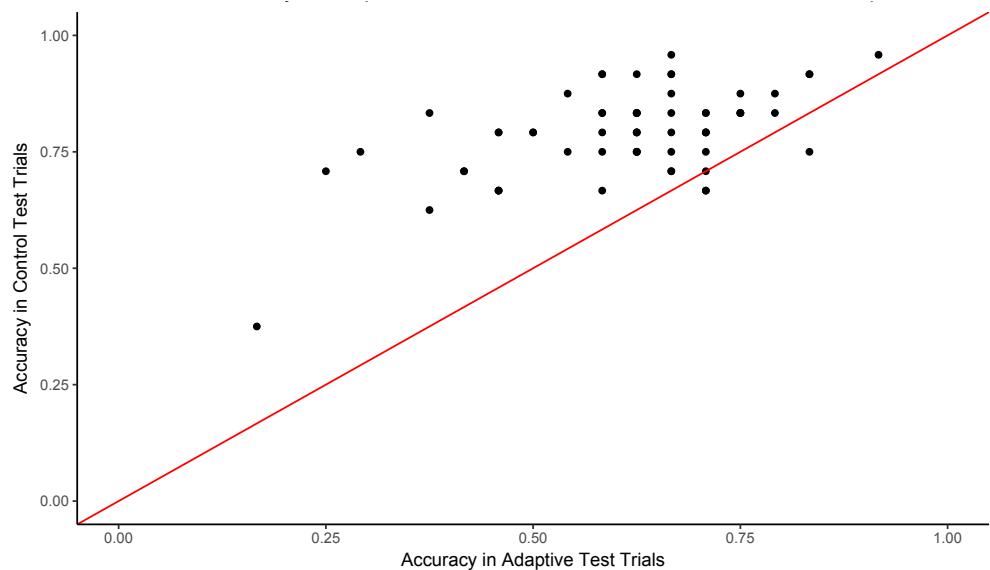
FIGURE 3. **Comparative accuracy levels in adaptive versus control test trials.** EACH DATA POINT CORRESPONDS TO ONE SUBJECT (N=63; 24 OBSERVATIONS PER SUBJECT). X AXIS: ACCURACY IN THE ADAPTIVE TEST TRIALS. Y AXIS: ACCURACY IN THE CONTROL TEST TRIALS. ACCURACY IS DEFINED AS THE FRACTION OF CORRECT REPLIES. DATA POINTS ABOVE THE 45 DEGREE LINE CORRESPOND TO SUBJECTS FOR WHOM ACCURACY IS DECREASED IN THE ADAPTIVE TEST TRIALS, AS PREDICTED BY OUTLIER BLINDNESS HYPOTHESIS.

they have been exposed during the sequence. Therefore, a litmus test of outlier blind-
ness hypothesis consists of suppressing adaptation in a follow-up "placebo experiment"
in which the length of each sequence is reduced to a very low level, so that the subjects
do not have time to form expectations about shade values, and hence there should be no
outlier blindness effect.

In the placebo experiment (N= 35; same cohort as in the original experiment: under-
graduate students at the University of New South Wales), each sequence lasts for only 3
trials. There are 160 sequences (960 trials overall), so 320 observations for our statistic of
interest. Except for sequence duration (and total number of sequences), the experimental
task is the same as the one used in the original experiment. Appendix IV.B reports the
main descriptive statistics with regard to response time, earnings, and accuracy per trial
type in the placebo experiment.

Figure 4 displays the distribution of accuracy in the adaptive and control test trials
across subjects, for both the original and placebo experiments. The main effect docu-
mented above for the original experiment is apparent on this graph (blue curves: the
mean accuracy level at the adaptive test trials is shifted to the left relative to the mean
accuracy at the control test trials). In contrast, the effect is absent in the placebo experi-
ment (red curves: the distributions of adaptive and control test trials accuracy overlap).[19]
One cannot reject the null hypothesis that in the placebo experiment, the mean accuracy
level is the same in the adaptive and control test trials, as predicted by outlier blindness
hypothesis (see Figure 11 in Appendix IV.C).

<center>Shortened Adaptation</center>

Our finding that the outlier blindness effect is suppressed in the placebo experiment
rules out the possibility that our statistic of interest in the original experiment be driven
by some hidden experimental artefact (since the task used in both experiments is the same
except for the adaptation length factor, as explained above). As such, it strengthens the
evidence for outlier blindness hypothesis.

Next we ask how long it takes for the human brain to hold expectations. On the basis
of prior work showing that adaptation is fairly quick in mammals,[20] we conjecture that
it does not take long and that increasing adaptation length by a few trials is enough

---

[19]Note that variance is lower for the placebo experiment than for the original experiment (the density
width is lower for the placebo experiment, for each trial type). This is because sample size is much
higher, namely 320 observations per subject vs. only 24 in the original experiment (no time to do more
to ensure task duration does not exceed 30 minutes).
[20]For example, Ohzawa et al. (1985), Bayer and Glimcher (2005), Fairhall (2014).
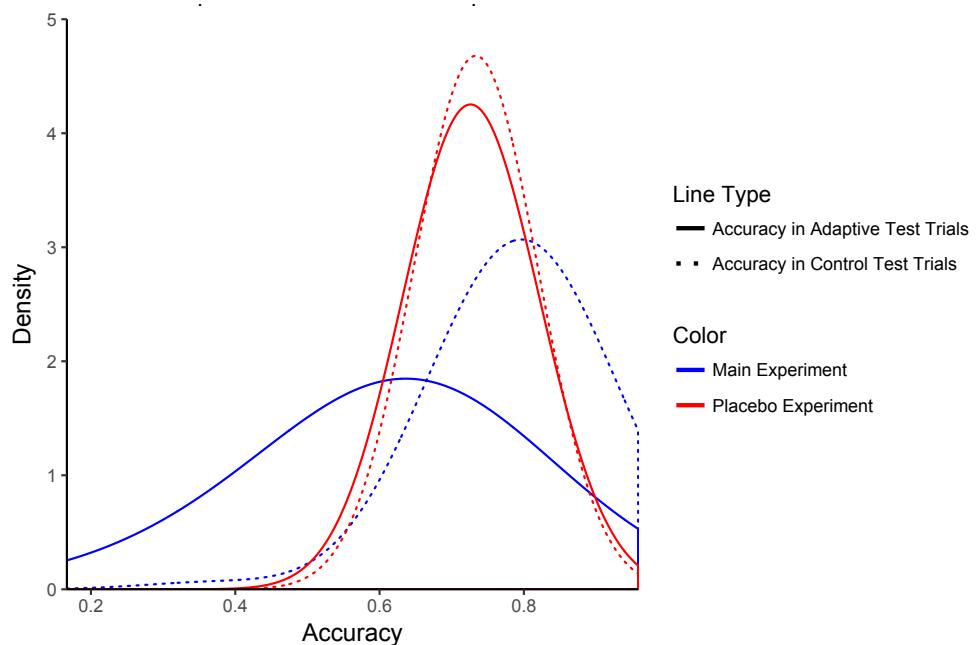
FIGURE 4. **Density plot of accuracy in the adaptive test trials and control test trials, for both the original experiment (blue) and the placebo experiment (red).** THE DENSITY IS DERIVED ACROSS SUBJECTS, BASED ON 63 SUBJECTS, 24 OBSERVATIONS PER SUBJECT (ORIGINAL EXPERIMENT), AND 35 SUBJECTS, 320 OBSERVATIONS PER SUBJECT (PLACEBO EXPERIMENT).

to restore the outlier blindness effect. To test our conjecture, we thus run the "5-trial adaptation experiment" (N=31; same cohort as in the original experiment: undergraduate students at the University of New South Wales) in which each sequence lasts for 5 trials. There are 100 sequences in that experiment (1,000 trials overall).[21] Consistent with our conjecture, we find the outlier blindness effect is significant in the 5-trial adaptation experiment. The null hypothesis that accuracy is the same in the adaptive and control test trials is rejected in each of our statistical tests (e.g., paired t-test: $t = 4.2$, $p < .001$, two-tailed; see Figure 11 in Appendix IV.C).

This finding shows that prolonged adaptation is not a precondition for the emergence of the outlier blindness effect. 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 original effect is only partially recovered: the outlier blindness effect is four times stronger in the original experiment than in the 5-trial adaptation experiment (two sample t test to compare the mean statistic of interest in the two experiments: $t = 8.8$, $p < .001$; see Figure 12 in Appendix IV.C). This again conforms to efficient coding theory which predicts that the magnitude of the outlier blindness effect increases with adaptation length; adaptation length is eight times bigger in the original experiment than in the 5-trial adaptation experiment.

<div align="center">ROBUSTNESS CHECKS</div>

We check that the foregoing findings are robust to excluding the data from the beginning of the task, to account for the possibility that subjects need some time to become familiar with the task interface.[22] The main accuracy results are also unchanged when we exclude from the analysis the missed trials. See the supplementary table in Appendix 14.

In the final stage of the study, we investigate whether the outlier blindness effect is affected by time pressure. We run a final experiment (N= 33; same cohort as in the original experiment: undergraduate students at the University of New South Wales) in which the time allowed to provide a reply on each trial is doubled relative to the original experiment (4 seconds vs. 2 seconds). The average response time in the experiment with double response time is 1.20 seconds on average across subjects, which is significantly

---

[21] See Appendix IV.B for the main descriptive statistics regarding response time, earnings, and accuracy per trial type in the 5-trial adaptation experiment.

[22] Note that a priori the subjects are familiar with the task interface from the start, since all of them complete a 3-minute training session prior to the task as per the experimental procedure described in Section I.B.

above the mean response time in the original experiment (two sample t-test: $p \sim 0$).[23]
However the magnitude of the outlier blindness effect is unchanged in the experiment
with double response time compared to that in the original experiment. This is already
apparent from Figure 5 which displays the distribution of accuracy in the adaptive and
control test trials across subjects, for both the original experiment and the experiment
with double response time. The distributions for each trial type look very similar across
the two experiments. Formal statistical tests confirm that doubling the response time
allowed on each trial does not affect the outlier blindness effect, which remains highly
significant (see Figure 11 in Appendix IV.C). We cannot reject the null hypothesis that
the outlier blindness effect is similar in the two experiments (two sample t-test to compare
the mean statistic of interest in the two experiments: $t = 1.9$; $p = 0.056$; see Figure 12
in Appendix IV.C).
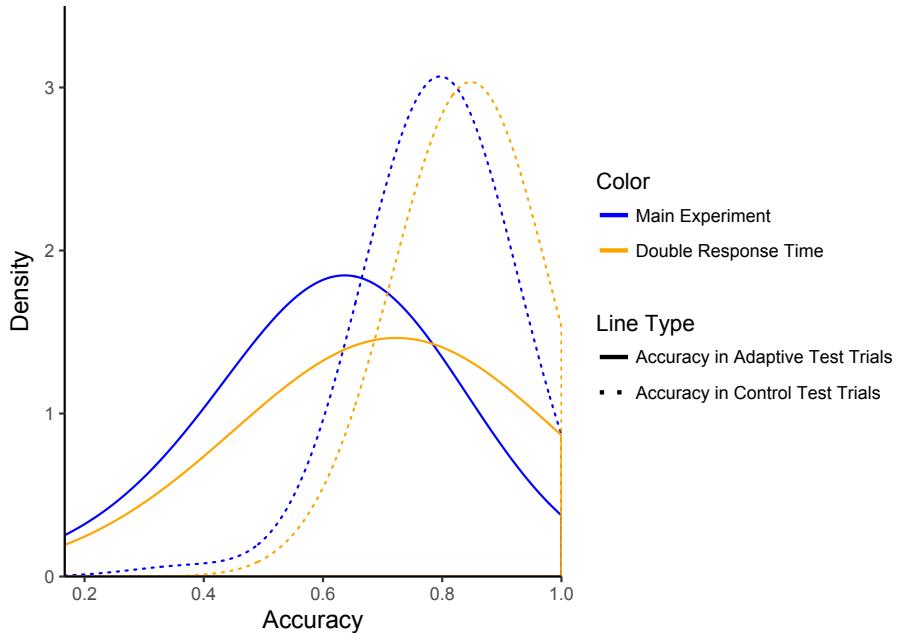
FIGURE 5. **Density plot of accuracy in the adaptive test trials and control test trials, for the
original experiment (blue) and the experiment with double response time (orange).** THE
DENSITY IS DERIVED ACROSS SUBJECTS, BASED ON 63 SUBJECTS, 24 OBSERVATIONS PER SUBJECT (ORIGINAL
EXPERIMENT), AND 33 SUBJECTS, 24 OBSERVATIONS PER SUBJECT (PLACEBO EXPERIMENT).

---

[23]See Appendix IV.B for the main descriptive statistics regarding response time, earnings, and accuracy
per trial type in the experiment with double response time.

III. Discussion

In this study we aim to test outlier blindness hypothesis, which is the novel idea that humans have difficulties in perceiving extreme values because the human brain evolved to draw sharper distinctions between magnitudes that occur more frequently, while making only coarse distinctions (if any at all) between magnitudes that are improbable. To test our hypothesis, we use a special experimental paradigm that allows us to separate the extreme-value effect (which simply reflects the fact that it is harder to see at the extremes on a scale of values, as described by the well-known Weber-Fechner law) from the outlier blindness effect which is the novel focus in this study. Also, to ensure that our findings are not related to some hidden factor related to risk preferences, we refrain from designing an economic decision-making task and instead use a purely perceptual task that involves the perception of a simple sensory stimulus.

We provide strong evidence that the outlier blindness phenomenon biases people's perception of outliers. Our findings further support the idea that the magnitude of the bias is highly significant in stable environments and still significant in very unstable environments (in our setting, when payoff distributions shift every 5 trials). Only in extremely unstable contexts where things change almost continuously does the bias vanish. In regard to human psychology, these findings suggest that humans are quick to form expectations. In regard to practical implications for financial agents, our findings suggest the outlier blindness phenomenon may be pervasive, prevailing over different kinds of environments both in terms of timescale (cf. the scale invariance property mentioned in the Introduction) and in terms of instability levels (no one should expect to be immune to the bias including those operating in very unstable environments).

The neglect of rare events that arises from the outlier blindness phenomenon is maladaptive in modern financial markets, where rare events are not rare enough for their neglect to be innocuous for economic agents (Taleb, 2004).[24] However it should be noted that this same neglect was adaptive ("efficient") in the ancestral human lifestyle, where the distribution of payoffs was most often unimodal. As such, the efficient coding principle and corollary outlier blindness phenomenon reflect the fact that our brains evolved to optimize survival chances in the environmental context in which they emerged. Taken out of their proper environment, they become maladaptive (see Robson (2001) and Chapter

---

[24]An exception is the case of outliers that should be ignored altogether as they are pure aberrations caused by leptokurtic noise (Bossaerts and d'Acremont, 2016). The 1987 stock market crash is one example of such "aberrations." Traders overreacted to it, thinking that something fundamental had changed, which was not the case (within 1 year, the stock market had more than recovered).

6 of Lo (2017)).

IV.   Appendix

A.   Instructions for the Task

The reader will find below the text of the instructions provided to the subjects in the original 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 make a choice on each trial (which is 4 seconds in the double-response time experiment).

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?

=

Submit
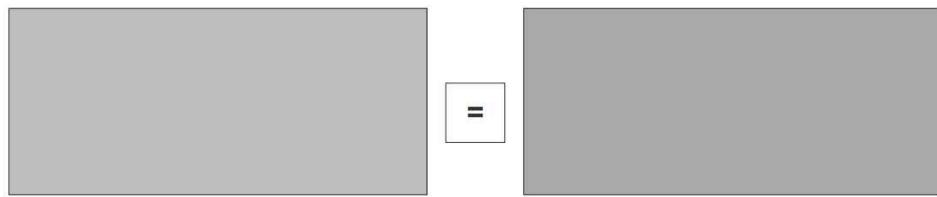
Trial: 1 / 1128

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:

I wish to do the training session before performing the task

I wish to skip the training session and go directly to the task

FIGURE 6. **Instructions for the task.**

B.   *Descriptive Statistics for Each Experiment*

Experiment \ Missed Trials	Mean	Standard Deviation	Minimum	Maximum
Main Experiment	6.17	10.04	0	50
5-Trial Adaption Experiment	7.52	10.32	0	47
Placebo Experiment	9.29	5.82	0	27
Experiment with Double Response Time	1.52	1.89	0	7

FIGURE 7. **Number of missed trials in each experiment.** THE MEAN, STANDARD DEVIATION, MIN AND MAX ARE DERIVED ACROSS SUBJECTS, FOR EACH EXPERIMENT.

Experiment \ Mean Accuracy	Overall	Same Color Trials	Different Color Trials	Adaptive Trials	Control Trials	Adaptive Test Trials	Control Test 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 Adaption 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)

FIGURE 8. **Mean accuracy according to trial type.** "ADAPTIVE TRIALS": TRIALS WITHIN THE ADAPTIVE SEQUENCES. "CONTROL TRIALS": TRIALS WITHIN THE CONTROL SEQUENCES. NUMBERS IN PARENTHESIS: SEM.

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

FIGURE 9. **Mean response time according to trial type.** "ADAPTIVE TRIALS": TRIALS WITHIN THE ADAPTIVE SEQUENCES. "CONTROL TRIALS": TRIALS WITHIN THE CONTROL SEQUENCES. NUMBERS IN PARENTHESIS: SEM.
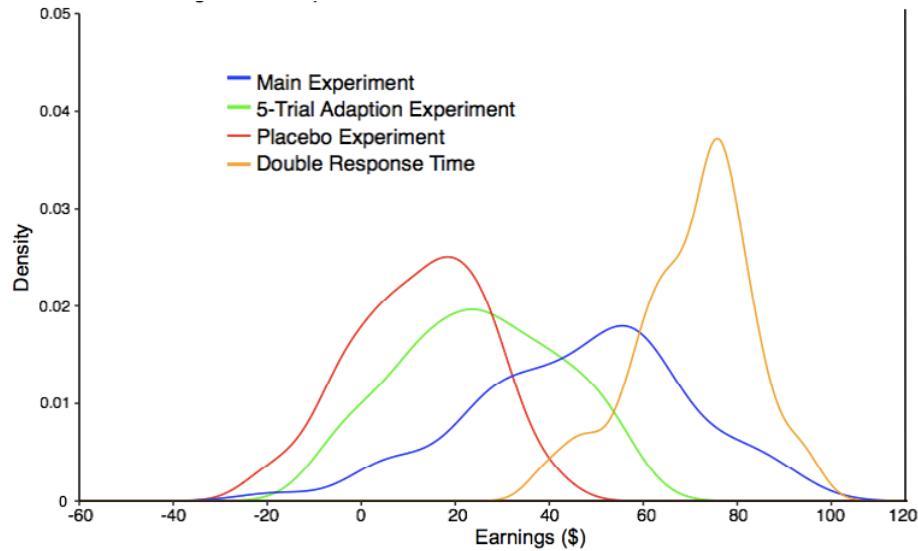

FIGURE 10. **Earning distribution across subjects in each experiment.** FOR EACH SUBJECT THE EARNINGS ARE COMPUTED AS THE NET ACCUMULATED OUTCOMES AT THE END OF THE TASK.

C.   Main Statistical Tests

Experiment	Mean Accuracy in Adaptive Test Trials	Mean Accuracy in Control Test Trials	TWO TAILED PAIRED T TEST	WILCOXON SIGNED 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 Adaption 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

FIGURE 11. **Tests of the outlier blindness effect, for each experiment.** THE OUTLIER BLINDNESS EFFECT IS MEASURED BY THE DIFFERENTIAL ACCURACY IN THE CONTROL TEST TRIALS VS. ADAPTIVE TEST TRIALS AS EXPLAINED IN THE MAIN TEXT. NUMBERS IN PARENTHESIS: SEM.

Experiment	(Mean Accuracy in Control Test Trials) – (Mean Accuracy in Adaptive 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 Adaption Experiment (N=31)	0.13 (0.02)	

FIGURE 12. **Comparison of the magnitude of the outlier blindness effect across experiments.** NUMBERS IN PARENTHESIS: SEM.

Experiment	Mean Accuracy in Adaptive Trials	Mean Accuracy in Control Trials	TWO TAILED PAIRED T TEST
Main Experiment (N=63)	0.87 (0.01)	0.81 (0.01)	t = 8.85, p ~ 0
5-Trial Adaption Experiment (N=31)	0.86 (0..01)	0.84 (0.01)	t = 2.75, p = 0.010
Placebo Experiment (N=35)	0.86 (0.01)	0.85 (0.01)	t = 0.68, p = 0.502
Experiment with Double Response Time (N=33)	0.92 (0.01)	0.89 (0.01)	t = 5.16, p < 0.0001

FIGURE 13. **Measure of the extreme-value effect.** "ADAPTIVE TRIALS": TRIALS WITHIN THE ADAPTIVE SEQUENCES. "CONTROL TRIALS": TRIALS WITHIN THE CONTROL SEQUENCES. THE EXTREME-VALUE EFFECT REFERS TO DISCRIMINATION BEING POTENTIALLY HAMPERED FOR VALUES THAT ARE AT THE EXTREMES OF THE SCALE RELATIVE TO THE OTHER VALUES. BY DESIGN, THE VALUES PRESENTED DURING THE CONTROL SEQUENCES ARE ON AVERAGE MORE EXTREME THAN THE VALUES PRESENTED DURING THE ADAPTIVE SEQUENCES. (THE MEAN SHADE VALUE FOR THE CONTROL SEQUENCES IS EITHER $3 - 4$ OR $9 - 10$ IN THE MAJORITY OF THE CASES WHEREAS FOR THE ADAPTIVE SEQUENCES, THE MEAN SHADE VALUE IS EVENLY DISTRIBUTED ACROSS ALL POSSIBLE VALUES $\{3, ..., 10\}$.) THEREFORE, ONE EXPECTS THE AVERAGE ACCURACY TO BE HIGHER IN THE ADAPTIVE SEQUENCES THAN IN THE CONTROL SEQUENCES. NUMBERS IN PARENTHESIS: SEM.

D.   Supplementary Table

Mean Accuracy Experiment	First 4 Trials	Missed Trials	Overall	Same Color Trials	Different Color Trials	Adaptive Trials	Control Trials	Adaptive Test Trials	Control Test Trials
Main Experiment	Include	Include	0.83 (0.01)	0.44 (0.02)	0.88 (0.01)	0.87 (0.01)	0.81 (0.01)	0.61 (0.01)	0.79 (0.01)
5-Trial Adaption Experiment	Include	Include	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	Include	Include	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	Include	Include	0.90 (0.01)	0.43 (0.03)	0.95 (0.01)	0.92 (0.01)	0.89 (0.01)	0.72 (0.01)	0.84 (0.01)
Main Experiment	Exclude	Include	0.83 (0.01)	0.44 (0.02)	0.88 (0.01)	0.87 (0.01)	0.81 (0.01)	0.61 (0.02)	0.79 (0.01)
5-Trial Adaption Experiment	Exclude	Include	0.80 (0.01)	0.41 (0.02)	0.83 (0.01)	0.86 (0.01)	0.84 (0.01)	0.72 (0.02)	0.75 (0.01)
Placebo Experiment	Exclude	Include	0.77 (0.01)	0.29 (0.03)	0.79 (0.01)	0.85 (0.01)	0.85 (0.01)	0.73 (0.01)	0.73 (0.01)
Experiment with Double Response Time	Exclude	Include	0.90 (0.01)	0.43 (0.03)	0.95 (0.01)	0.92 (0.01)	0.89 (0.01)	0.72 (0.01)	0.84 (0.01)
Main Experiment	Include	Exclude	0.84 (0.01)	0.45 (0.02)	0.89 (0.01)	0.87 (0.01)	0.82 (0.01)	0.62 (0.02)	0.79 (0.01)
Experiment with Double Response Time	Include	Exclude	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)
Placebo Experiment	Include	Exclude	0.78 (0.01)	0.30 (0.03)	0.80 (0.01)	0.86 (0.01)	0.86 (0.01)	0.74 (0.01)	0.74 (0.01)
5-Trial Adaption Experiment	Include	Exclude	0.81 (0.01)	0.42 (0.03)	0.84 (0.01)	0.86 (0.01)	0.85 (0.01)	0.73 (0.01)	0.75 (0.01)
Main Experiment	Exclude	Exclude	0.84 (0.01)	0.45 (0.02)	0.89 (0.01)	0.87 (0.01)	0.82 (0.01)	0.62 (0.02)	0.79 (0.01)
5-Trial Adaption Experiment	Exclude	Exclude	0.81 (0.01)	0.42 (0.03)	0.84 (0.01)	0.86 (0.01)	0.85 (0.01)	0.73 (0.01)	0.75 (0.01)
Placebo Experiment	Exclude	Exclude	0.78 (0.01)	0.30 (0.03)	0.80 (0.01)	0.86 (0.01)	0.86 (0.01)	0.74 (0.01)	0.74 (0.01)
Experiment with Double Response Time	Exclude	Exclude	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)

FIGURE 14. **Mean accuracy according to trial type for each experiment.** "ADAPTIVE TRIALS": TRIALS WITHIN THE ADAPTIVE SEQUENCES. "CONTROL TRIALS": TRIALS WITHIN THE CONTROL SEQUENCES. NUMBERS IN PARENTHESIS: SEM.

REFERENCES

Andrew Ang and Allan G. Timmermann. Regime changes and financial markets. Technical report, Netspar Discussion Paper No. 06/2011-068, 2011.

Hannah M. Bayer and Paul W. Glimcher. Midbrain dopamine neurons encode a quantitative reward prediction error signal. *Neuron*, 47:129–141, 2005.

Peter Bossaerts and Mathieu d'Acremont. Neural mechanisms behind identification of leptokurtic noise and adaptive behavioral response. *Cerebral Cortex*, 26:1818–1830, 2016.

Gordon D.A. Brown, Nick Chater, and Ian Neath. A temporal ratio model of memory. *Psychological Review*, 114:539–576, 2007.

Matteo Carandini and David J. Heeger. Normalization as a canonical neural computation. *Nature Reviews Neuroscience*, 13:51–62, 2012.

Adrienne Fairhall. *The Cognitive Neurosciences*, chapter Adaptation and Natural Stimulus Statistics. MIT Press, London, UK, 2014.

Eugene F. Fama. The behavior of stock market prices. *Journal of Business*, 38:34–105, 1965.

Vera Ferrari, Margaret M. Bradley, Maurizio Codispoti, and Peter J. Lang. Detecting novelty and significance. *Journal of Cognitive Neuroscience*, 22:404–411, 2010.

Cary Frydman and Lawrence J. Jin. Efficient coding and risky choice. *Working Paper*, 2018.

Xavier Gabaix. A sparsity-based model of bounded rationality. *The Quarterly Journal of Economics*, 129:1661–1710, 2014.

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 fragiity. *Princeton University Press*, 2018.

Nicola Gennaioli, Andrei Shleifer, and Robert Vishny. Neglected risks, financial innovation, and financial fragility. *Journal of Financial Economics*, 104:452–468, 2012.

Nicola Gennaioli, Andrei Shleifer, and Robert Vishny. Neglected risks: The psychology of financial crises. *American Economic Review*, 105:310–314, 2015.

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.

William N. Goetzmann, Dasol Kim, and Robert J. Shiller. Crash beliefs from investor surveys. *The National Bureau of Economic Research*, 2017.

Lawrence J. Jin. A speculative asset pricing model of financial instability. *SSRN*, 2015.

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.

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.

Arien Mack and Irvin Rock. *Inattentional Blindness*. MIT Press, 1998.

Bartosz Mackowiak and Mirko Wiederholt. Optimal sticky prices under rational inattention. *American Economic Review*, 99(3):769–803, 2009.

Benoit B. Mandelbrot. *Fractales, Hasard et Finance*. Paris: Flammarion, 1957.

Nick Netzer. Evolution of time preferences and attitudes toward risk. *American Economic Review*, 99:937–955, 2009.

Izumi Ohzawa, Gary Sclar, and Ralph D. Freeman. Contrast gain control in the cat's visual system. *Journal of Neurophysiology*, 54:651–667, 1985.

Elise Payzan-LeNestour. Can people learn about 'black swans'? experimental evidence. *Review of Financial Studies*, Forthcoming, 2018.

Elise Payzan-LeNestour, Lionnel Pradier, and Tālis J. Putniņš. The waterfall illusion, and how prior exposure to extreme risk distorts risk perception. *SSRN*, 2018.

Rafael Polania, Michael Woodford, and Christian C. Ruff. Efficient coding of subjective value. *Nature Neuroscience*, 22:134–142, 2019.

Arthur J. Robson. The biological basis of economic behavior. *Journal of Economic Literature*, 39:11–33, 2001.

Eldar Shafir. Uncertainty and the difficulty of thinking through disjunctions. *Cognition*, 50:403–430, 1994.

Eldar Shafir and Amos Tversky. Thinking through uncertainty: Nonconsequential reasoning and choice. *Cognitive Psychology*, 24(4):449–474, 1992.

Eero P. Simoncelli. Vision and the statistics of the visual environment. *Current Opinion in Neurobiology*, 13:144–149, 2003.

Daniel J Simons and Christopher F. Chabris. Gorillas in our midst: Sustained inattentional blindness for dynamic events. *Perception*, 28(9):1059–74, 1999.

Christopher A. Sims. Implications of rational inattention. *Journal of Monetary Economics*, 50(3):665–690, 2003.

Nancy K. Squires, Kenneth C. Squires, and Steven A. Hillyard. Two varieties of long-latency positive waves evoked by unpredictable auditory stimuli in man. *Electroencephalography and Clinical Neurophysiology*, 38:387–401, 1975.

Jakub Steiner and Colin Stewart. Perceiving prospects properly. *American Economic Review*, 106:1601–1631, 2016.

Nassim N. Taleb. Bleed or blowup: What does empirical psychology tell us about the preference for negative skewness? *Journal of Behavioral Finance*, 5, 2004.

Nassim N. Taleb. *The Black Swan: The Impact of the Highly Improbable*. Random House, NY, 2007.

Philippe N. Tobler, Christopher D. Fiorillo, and Wolfram Schultz. Adaptive coding of reward value by dopamine neurons. *Science*, 307:1642–1645, 2005.

Peter Ulric Tse, James Intriligator, Josée Rivest, and Patrick Cavanagh. Attention and the subjective expansion of time. *Perception and Psychophysics*, 66:1171–1189, 2004.

Amos Tversky and Eldar Shafir. The disjunction effect in choice under uncertainty. *Psychological Science*, 3(5):305–309, 1992.

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.

Michael Woodford. Inattentive valuation and reference-dependent choice. *Working paper*, 2012.