

SALIENCE THEORY OF CHOICE UNDER RISK

PEDRO BORDALO
NICOLA GENNAIOLI
ANDREI SHLEIFER

We present a theory of choice among lotteries in which the decision maker's attention is drawn to (precisely defined) salient payoffs. This leads the decision maker to a context-dependent representation of lotteries in which true probabilities are replaced by decision weights distorted in favor of salient payoffs. By specifying decision weights as a function of payoffs, our model provides a novel and unified account of many empirical phenomena, including frequent risk-seeking behavior, invariance failures such as the Allais paradox, and preference reversals. It also yields new predictions, including some that distinguish it from prospect theory, which we test. *JEL* Codes: D03, D81.

I. INTRODUCTION

Over the past several decades, social scientists have identified a range of important violations of expected utility theory, the standard theory of choice under risk. Perhaps at the most basic level, in both experiments and everyday life, people frequently exhibit both risk-loving and risk-averse behavior, depending on the situation. As first stressed by Friedman and Savage (1948), people participate in unfair gambles, pick highly risky occupations (including entrepreneurship) over safer ones, and invest without diversification in individual risky stocks, while simultaneously buying insurance. Attitudes toward risk are unstable in this very basic sense.

This systematic instability underlies several paradoxes of choice under risk. As shown by Allais (1953), people switch

We are grateful to Nicholas Barberis, Gary Becker, Colin Camerer, John Campbell, Tom Cunningham, Xavier Gabaix, Morgan Grossman-McKee, Ming Huang, Jonathan Ingersoll, Emir Kamenica, Daniel Kahneman, Botond Koszegi, David Laibson, Pepe Montiel Olea, Drazen Prelec, Matthew Rabin, Josh Schwartzstein, Jesse Shapiro, Jeremy Stein, Tomasz Strzalecki, Dmitry Taubinsky, Richard Thaler, Georg Weiszacker, George Wu and three referees of this journal for extremely helpful comments, and to Allen Yang for excellent research assistance. Gennaioli thanks the Spanish Ministerio de Ciencia y Tecnologia (ECO 2008-01666 and Ramon y Cajal grants), the Barcelona Graduate School of Economics Research Network, and the Generalitat de Catalunya for financial support. Shleifer thanks the Kauffman Foundation for research support.

© The Author(s) 2012. Published by Oxford University Press, on behalf of President and Fellows of Harvard College. All rights reserved. For Permissions, please email: journals.permissions@oup.com

The Quarterly Journal of Economics (2012), 1243–1285. doi:10.1093/qje/qjs018.

Advance Access publication on April 13, 2012.

from risk-loving to risk-averse choices among two lotteries after a common consequence is added to both, in contradiction to the independence axiom of expected utility theory. Another form of instability is preference reversals (Lichtenstein and Slovic 1971): in comparing two lotteries with a similar expected value, experimental subjects *choose* the safer lottery but are willing to *pay more* for the riskier one. Camerer (1995) reviews numerous attempts to amend the axioms of expected utility theory to deal with these findings, but these attempts have not been conclusive.

We propose a new psychologically founded model of choice under risk, which naturally exhibits the systematic instability of risk preferences and accounts for the puzzles. In this model, risk attitudes are driven by the salience of different lottery payoffs. Psychologists view salience detection as a key attentional mechanism enabling humans to focus their limited cognitive resources on a relevant subset of the available sensory data. As Taylor and Thompson (1982) put it: “Salience refers to the phenomenon that when one’s attention is differentially directed to one portion of the environment rather than to others, the information contained in that portion will receive disproportionate weighting in subsequent judgments.” According to Kahneman (2011, 324), “our mind has a useful capability to focus on whatever is odd, different or unusual.” We call the payoffs that draw the decision maker’s attention “salient.” The decision maker is then risk-seeking when a lottery’s upside is salient and risk-averse when its downside is salient. More generally, salience allows for a theory of context-dependent choice consistent with a broad range of evidence.

We build a model of decision making in which salient lottery payoffs are overweighted. Our main results rely on three assumptions. Two of them, which we label *ordering* and *diminishing sensitivity*, formalize the salience of payoffs. Roughly speaking, a lottery payoff is salient if it is very different in percentage terms from the payoffs of other available lotteries (in the same state of the world). This specification of salience captures the ideas that (1) we attend to differences rather than absolute values (Kahneman 2003), and (2) we perceive changes on a log scale (Weber’s law). Our third assumption states that the extent to which decision weights are distorted depends on the salience of the associated payoffs, and not on the underlying probabilities. This assumption implies (see Proposition 1) that low probabilities

are relatively more distorted than high ones, in accordance with Kahneman and Tversks's (1979) observation that people have "limited ability to comprehend and evaluate extreme probabilities." We describe how, under these assumptions, the decision maker develops a context-dependent representation of each lottery. Aside from replacing objective probabilities with decision weights, the decision maker's valuation of payoffs is standard.

At a broad level, our approach is similar to that pursued by Gennaioli and Shleifer (2010) in their study of the representativeness heuristic in probability judgments. The idea of both studies is that decision makers do not take into account fully all the information available to them, but overemphasize the information their minds focus on.¹ Gennaioli and Shleifer (2010) call such decision makers local thinkers, because they neglect potentially important but unrepresentative data. Here, analogously, in evaluating lotteries, decision makers overweight states that draw their attention and neglect states that do not. We continue to refer to such decision makers as local thinkers. In both models, the limiting case in which all information is processed correctly is the standard economic decision maker.

Our model describes factors that encourage and discourage risk seeking, but also leads to an explanation of the Allais paradoxes. The strongest departures from expected utility theory in our model occur in the presence of extreme payoffs, particularly when these occur with a low probability. Due to this property, our model predicts that subjects in the Allais experiments are risk-loving when the common consequence is small and attention is drawn to the highest lottery payoffs, and risk-averse when the common consequence is large and attention is drawn to the lowest payoffs. We explore the model's predictions by describing, and then experimentally testing, how Allais paradoxes can be turned on and off. We also show that preference reversals can be seen as a consequence of lottery evaluation in different contexts that affect salience, rather than the result of a fundamental difference between pricing and choosing. The model thus provides a unified explanation of risk preferences and invariance violations based on a psychologically motivated mechanism of salience.

It is useful to compare our model to the gold standard of behavioral theories of choice under risk, Kahneman and Tversky's

¹Other models in the same spirit are Mullainathan (2002), Schwartzstein (2009), and Gabaix (2011).

(1979) (henceforth KT) prospect theory. Like prospect theory, our model incorporates the assumption that decision makers focus on payoffs, rather than on absolute wealth levels, when evaluating risky alternatives. Prospect theory also incorporates the assumption that the probability weights people use to make choices are different from objective probabilities. But the idea that these weights depend on the actual payoffs and their salience is new here. In some situations, our decision weights look very similar to KT's, but in other situations—for instance when small probabilities are *not* attached to salient payoffs or when lotteries are correlated—they are very different. We conduct multiple experiments, both of simple risk attitudes and of Allais paradoxes with correlated states, which distinguish our predictions from KT's, and uniformly find strong support for our model of probability weighting.

The article proceeds as follows. In Section II, we provide the basic intuition for how the salience of lottery payoffs shapes risk attitudes in the context of Allais' "common consequence" paradox. In Section III, we present a salience-based model of choice among two lotteries. In Section IV, we use this model to study risk attitudes, derive from first principles prospect theory's weighting function for a class of choice problems where it should apply, and provide experimental evidence for our predictions. In Section V we show that our model accounts for the Allais paradoxes, as well as for preference reversals, a phenomenon that prospect theory cannot accommodate. We obtain further predictions for context effects (which prospect theory also cannot accommodate), such as turning the Allais paradoxes or preference reversals on and off depending on the description of payoff states, and find experimental support for these predictions. We then describe how the model deals with losses and addresses reflection and framing effects. In Section VI, we take stock of the model's predictions, and compare it to alternative models of choice under risk. Section VII concludes. Proofs of the results in the text can be found in the Appendix. As supplementary material, Online Appendix 1 presents additional results on preference reversals and failures of transitivity, addresses mixed lotteries, and extends the model to choice among many lotteries. Online Appendix 2 provides a detailed account of the experimental procedures and results.

II. SALIENCE AND THE ALLAIS PARADOX

The Allais (1953) paradoxes are the best known and most discussed instances of failure of the independence axiom of expected utility theory. KT's version of the "common consequence" paradox asks experimental subjects to choose among two lotteries $L_1(z)$ and $L_2(z)$:

$$(1) \quad L_1(z) = \begin{cases} \$2500 & \text{with prob. } 0.33 \\ \$0 & 0.01 \\ \$z & 0.66 \end{cases}, \quad L_2(z) = \begin{cases} \$2400 & \text{with prob. } 0.34 \\ \$z & 0.66 \end{cases}$$

for different values of the payoff z . By the independence axiom, an expected utility maximizer should not change his choice as the common consequence z is varied, since z cancels out in the comparison between $L_1(z)$ and $L_2(z)$.

In experiments, for $z = 2400$, most subjects are risk-averse, preferring $L_2(2400)$ to $L_1(2400)$:

$$(2) \quad L_1(2400) = \begin{cases} \$2500 & \text{with prob. } 0.33 \\ \$0 & 0.01 \\ \$2400 & 0.66 \end{cases} < L_2(2400) = \begin{cases} \$2400 & \text{with prob. } 1. \end{cases}$$

However, when $z = 0$, most subjects are risk seeking, preferring $L_1(0)$ to $L_2(0)$:

$$(3) \quad L_1(0) = \begin{cases} \$2500 & \text{with prob. } 0.33 \\ \$0 & 0.67 \end{cases} > L_2(0) = \begin{cases} \$2400 & \text{with prob. } 0.34 \\ \$0 & 0.66 \end{cases}$$

In violation of the independence axiom, z affects the experimental subjects' choices, causing switches between risk-averse and risk-seeking behavior. Prospect theory (Kahneman and Tversky 1979; Tversky and Kahneman 1992) explains these switches as follows. When $z = 2400$, the low 0.01 probability of getting zero in $L_1(2400)$ is overweighted, generating risk aversion. When $z = 0$, the extra 0.01 probability of getting zero in $L_1(0)$ is not overweighted, generating risk seeking. This effect is directly built into the probability weighting function $\pi(p)$

by the assumption of subcertainty, for example, $\pi(0.34) - \pi(0) < 1 - \pi(0.66)$.²

Our explanation of the Allais paradox does not rely on a fixed weighting function $\pi(p)$. Rather, it relies on how decision weights change as the payoff z alters the salience of different lottery outcomes. Roughly speaking, in the choice between $L_1(2400)$ and $L_2(2400)$, the downside of \$0 feels a lot lower than the sure payoff of \$2400. The upside of \$2500, however, feels only slightly higher than the sure payoff. Because the lottery's downside is more salient than its upside, the subjects focus on the downside when making their decisions. This focus triggers the risk-averse choice.

In contrast, in the choice between $L_1(0)$ and $L_2(0)$, both lotteries have the same downside risk of zero. Now the upside of winning \$2500 in the riskier lottery $L_1(0)$ is more salient and subjects focus on it when making their decisions. This focus triggers the risk-seeking choice. The analogy here is to sensory perception: a lottery's salient payoffs are those that differ most from the payoffs of alternative lotteries. The decision maker's mind then focuses on salient payoffs, inflating their weights when making a choice. Section V provides a fuller account of the Allais experiment, which also highlights the role played by the level of objective probabilities.

III. THE MODEL

A choice problem is described by (1) a set of states of the world S , where each state $s \in S$ occurs with objective and known probability π_s such that $\sum_{s \in S} \pi_s = 1$, and (2) a choice set $\{L_1, L_2\}$, where the L_i are risky prospects that yield monetary payoffs x_s^i in each state s . For convenience, we refer to L_i as lotteries.³ Here we focus

²In cumulative prospect theory (Tversky and Kahneman 1992) the mathematical condition on probability weights is slightly different but carries the same intuition: the common consequence is more valuable when associated with a sure rather than a risky prospect.

³Formally, L_i are acts, or random variables, defined over the choice problem's probability space (S, F_S, π) , where S is assumed to be finite and F_S is its canonical σ -algebra. However, as we will see in equation (11), the decision maker's choice depends only on the L_i 's joint distribution over payoffs and not on the exact structure of the state space. Thus we use the term lotteries in a slight abuse of nomenclature relative to the usual definition of lotteries as probability distributions over payoffs.

on choice between two lotteries, leaving the general case of choice among $N > 2$ lotteries to Online Appendix 1.

The decision maker uses a value function v to evaluate lottery payoffs relative to the reference point of zero.⁴ Through most of the article, we illustrate the mechanism generating risk preferences in our model by assuming a *linear* value function v (in Online Appendix 1, when we focus on mixed lotteries, we consider a piece-wise linear value function featuring loss aversion, as in prospect theory). Absent distortions in decision weights, the local thinker evaluates L_i as:

$$(4) \quad V(L_i) = \sum_{s \in S} \pi_s v(x_s^i).$$

The local thinker (LT) departs from equation (4) by overweighting the lottery's most salient states in S . Saliency distortions work in two steps. First, a salience ranking among the states in S is established for each lottery L_i . Second, based on this salience ranking, the probability π_s in equation (4) is replaced with a transformed, lottery-specific decision weight π_s^i . To formally define salience, let $\mathbf{x}_s = (x_s^i)_{i=1,2}$ be the vector listing the lotteries' payoffs in state s and denote by x_s^{-i} the payoff in s of lottery L_j , $j \neq i$. Let x_s^{\min}, x_s^{\max} respectively denote the largest and smallest payoffs in \mathbf{x}_s .

DEFINITION 1. The salience of state s for lottery L_i , $i = 1, 2$, is a continuous and bounded function $\sigma(x_s^i, x_s^{-i})$ that satisfies three conditions:

1. **Ordering.** If for states $s, \tilde{s} \in S$ we have that $[x_s^{\min}, x_s^{\max}]$ is a subset of $[x_{\tilde{s}}^{\min}, x_{\tilde{s}}^{\max}]$, then

$$\sigma(x_s^i, x_s^{-i}) < \sigma(x_{\tilde{s}}^i, x_{\tilde{s}}^{-i}).$$

2. **Diminishing sensitivity.** If $x_s^j > 0$ for $j = 1, 2$, then for any $\epsilon > 0$,

$$\sigma(x_s^i + \epsilon, x_s^{-i} + \epsilon) < \sigma(x_s^i, x_s^{-i}).$$

⁴This is a form of narrow framing, also used in prospect theory. Koszegi and Rabin (2006, 2007) build a model of reference point formation and use it to study shifts in risk attitudes. Their model cannot account for situations where expectations and thus reference points are held fixed (such as lab experiments we consider here). Our approaches are complementary, as one could combine our model of decision weights with Koszegi and Rabin's two-part value function.

3. Reflection. For any two states $s, \tilde{s} \in S$ such that $x_s^j, x_{\tilde{s}}^j > 0$ for $j = 1, 2$, we have

$$\sigma(x_s^i, x_s^{-i}) < \sigma(x_{\tilde{s}}^i, x_{\tilde{s}}^{-i}) \text{ if and only if } \sigma(-x_s^i, -x_s^{-i}) < \sigma(-x_{\tilde{s}}^i, -x_{\tilde{s}}^{-i}).$$

Section III.A discusses the connection between these properties and the cognitive notion of salience. The key properties driving our explanations of anomalies are ordering and diminishing sensitivity. The reflection property only plays a role in Section V.C when we consider lotteries that yield negative payoffs. To illustrate Definition 1, consider the salience function:

$$(5) \quad \sigma(x_s^i, x_s^{-i}) = \frac{|x_s^i - x_s^{-i}|}{|x_s^i| + |x_s^{-i}| + \theta},$$

where $\theta > 0$. According to the ordering property, the salience of a state for L_i increases in the distance between its payoff x_s^i and the payoff x_s^{-i} of the alternative lottery. In (5), this is captured by the numerator $|x_s^i - x_s^{-i}|$. Diminishing sensitivity implies that salience decreases as a state's average (absolute) payoff gets farther from zero, as captured by the denominator term $|x_s^1| + |x_s^2|$ in (5). Finally, according to reflection, salience is shaped by the magnitude rather than the sign of payoffs: a state is salient not only when the lotteries bring sharply different gains, but also when they bring sharply different losses. In (5), reflection takes the strong form $\sigma(x_s^i, x_s^{-i}) = \sigma(-x_s^i, -x_s^{-i})$. These properties are illustrated in Figure I.

The salience function in specification (5) satisfies additional properties besides those of Definition 1. For instance, it is symmetric, namely, $\sigma(x_s^1, x_s^2) = \sigma(x_s^2, x_s^1)$, which is a natural property in the case of two lotteries but which is dropped with $N > 2$ lotteries. Although our main results rely only on ordering and diminishing sensitivity, we sometimes use the tractable functional form (5) to illustrate our model.

Consider the choice between $L_1(z)$ and $L_2(z)$ introduced in Section II. When the common consequence is $z = 2400$, the possible payoff states are $S = \{(2500, 2400), (0, 2400), (2400, 2400)\}$. We then have:

$$(6) \quad \sigma(0, 2400) > \sigma(2500, 2400) > \sigma(2400, 2400).$$

The inequalities follow from diminishing sensitivity and ordering, respectively, and can be easily verified for equation (5). The state in which the riskier lottery $L_1(2400)$ loses is the most salient




FIGURE I
Properties of a Saliency Function, Equation (5)

one (which causes risk aversion).⁵ A similar calculation shows that, when the common consequence is $z = 0$, the state $(2500, 0)$ in which the risky lottery $L_1(0)$ wins is the most salient one, which points to risk seeking. In short, changing the common consequence affects the salience of lottery payoffs, as described in Section II. Section V.A provides a full analysis of the Allais paradoxes.

III.A. Saliency, Decision Weights, and Risk Attitudes

Given a saliency function σ , for each lottery L_i the local thinker ranks the states and distorts their decision weights as follows.

DEFINITION 2. Given states $s, \tilde{s} \in S$, we say that for lottery L_i state s is more salient than \tilde{s} if $\sigma(x_s^i, x_{\tilde{s}}^{-i}) > \sigma(x_{\tilde{s}}^i, x_s^{-i})$. Let $k_s^i \in \{1, \dots, |S|\}$ be the saliency ranking of state s for L_i , with lower k_s^i indicating higher salience. All states with the same

⁵In this example, constructing the state space from the alternatives of choice is straightforward. Section III.B discusses how the state space S is constructed in more complex cases.

salience obtain the same ranking (and the ranking has no jumps). Then the local thinker transforms the odds $\frac{\pi_{\tilde{s}}}{\pi_s}$ of \tilde{s} relative to s into the odds $\frac{\pi_{\tilde{s}}^i}{\pi_s^i}$, given by:

$$(7) \quad \frac{\pi_{\tilde{s}}^i}{\pi_s^i} = \delta^{k_{\tilde{s}}^i - k_s^i} \cdot \frac{\pi_{\tilde{s}}}{\pi_s}$$

where $\delta \in (0, 1]$. By normalizing $\sum_s \pi_s^i = 1$ and defining $\omega_s^i = \frac{\delta^{k_s^i}}{(\sum_r \delta^{k_r^i} \cdot \pi_r)}$, the decision weight attached by the local thinker to a generic state s in the evaluation of L_i is:

$$(8) \quad \pi_s^i = \pi_s \cdot \omega_s^i.$$

The local thinker evaluates a lottery by inflating the relative weights attached to the lottery's most salient states. Parameter δ measures the extent to which salience distorts decision weights, capturing the degree of local thinking. When $\delta = 1$, the decision maker is a standard economic decision maker: his decision weights coincide with objective probabilities (i.e., $\omega_s^i = 1$). When $\delta < 1$, the decision maker is a local thinker, namely, he overweights the most salient states and underweights the least salient ones. Specifically, s is overweighted if and only if it is more salient than average ($\omega_s^i > 1$, or $\delta^{k_s^i} > \sum_r \delta^{k_r^i} \cdot \pi_r$). The case where $\delta \rightarrow 0$ describes the local thinker who focuses only on a lottery's most salient payoffs.

The critical property of Definition 2 is that the parameter δ does not depend on the objective state probabilities. This specification implies:

PROPOSITION 1. If the probability of state s is increased by $d\pi_s = h \cdot \pi_s$, where h is a positive constant, and the probabilities of other states are reduced while keeping their odds constant, that is, $d\pi_{\tilde{s}} = -\frac{\pi_{\tilde{s}}}{1 - \pi_s} h \cdot \pi_{\tilde{s}}$ for all $\tilde{s} \neq s$, then:

$$(9) \quad \frac{d\omega_s^i}{h} = -\frac{\pi_s}{1 - \pi_s} \cdot \omega_s^i \cdot (\omega_s^i - 1).$$

Proposition 1 (see the Appendix for proofs) states that an increase in a state's probability π_s reduces the distortion of the decision weight in that state by driving ω_s^i closer to 1. That is, low-probability states are subject to the strongest

distortions:⁶ they are overweighted if salient and underweighted otherwise. In contrast to KT's (1979; Tversky and Kahneman 1992) assumption, low-probability (high-rank) payoffs are not always overweighted in our model; they are only overweighted if they are salient, regardless of probability (and rank). In accordance with KT, however, the largest distortions of choice occur precisely when salient payoffs are relatively unlikely. This property plays a key role for explaining some important findings such as the common ratio Allais paradox in Section V.A.⁷

Given Definitions 1 and 2, the local thinker computes the value of lottery L_i as:

$$(10) \quad V^{LT}(L_i) = \sum_{s \in S} \pi_s^i v(x_s^i) = \sum_{s \in S} \pi_s \omega_s^i v(x_s^i).$$

Thus, L_i 's evaluation always lies between the value of its highest and lowest payoffs.

Because salience is defined on the state space S , one may wonder whether splitting states, or generally considering a different state space compatible with the lotteries' payoff distributions, may affect the local thinker's evaluation (10). We denote by X the set of distinct payoff combinations of L_1, L_2 occurring in S with positive probability, and by S_x the set of states in S where the lotteries yield the same payoff combination $x \in X$, formally $S_x \equiv \{s \in S \mid x_s = x\}$. Clearly, $S = \cup_{x \in X} S_x$. By Definition 1, all states s in S_x are equally salient for either lottery, and thus have the same value of ω_s^i , which for simplicity we denote ω_x^i . Using (8) we can rewrite $V^{LT}(L_i)$ in equation (10) as:

$$(11) \quad V^{LT}(L_i) = \sum_{x \in X} \left(\sum_{s \in S_x} \pi_s \right) \omega_x^i v(x_x^i),$$

where x_x^i denotes L_i 's payoff in x . Equation (11) says that the state space only influences evaluation through the total probability

⁶This follows from the normalization of the decision weights. Since the expected distortion is zero, $\sum_i \pi_i \omega_s^i = 1$, and since the distortion factor ω_s^i for state s does not depend directly on its probability, states with lower probabilities are relatively more distorted.

⁷Proposition 1 can also be stated in terms of payoffs: if lottery L_i yields payoff x_k with probability p_k , then increasing p_k while reducing the probabilities $p_{k'}$ of other payoffs $x_{k'}$ (keeping their odds constant) decreases the distortion of p_k if and only if x_k is more salient than average. That is, in a given choice context, the probabilities of unlikely payoffs are relatively more distorted (see the Appendix for details).

of each distinct payoff combination x , namely, $\pi_x = \sum_{s \in S_x} \pi_s$. This is because salience σ depends on payoffs, and not on the probabilities of different states. Hence, splitting a given probability π_x across different sets of states does not affect evaluation (or choice) in our model. There is therefore no loss in generality from viewing S as the “minimal” state space X identified by the set of distinct payoff combinations that occur with positive probability. In the remainder of the article, we keep the notation of equation (10), with the understanding that S is this “minimal” state space (and omit the reference to the underlying lotteries).

In a choice between two lotteries, equation (10) implies that—due to the symmetry of the salience function (i.e., $k_s^1 = k_s^2$ for all s)—the local thinker prefers L_1 to L_2 if and only if:

$$(12) \quad \sum_{s \in S} \delta^{k_s} \pi_s [v(x_s^1) - v(x_s^2)] > 0.$$

For $\delta = 1$, the local thinker’s decision weights coincide with the corresponding objective probabilities. For $\delta < 1$, local thinking favors L_1 when it pays more than L_2 in the more salient (and thus less discounted) states.

III.B. Discussion of Assumptions and Setup

Saliency and Decision Weights. In our model the choice context shapes decision makers’ perception of lotteries through the mechanism of payoff salience. The properties of the salience function seek to formalize features of human perception, which we believe—in line with Kahneman, Tversky, and others—to be relevant for choice under risk. The intensity with which we perceive a signal, such as a light source, increases in the signal’s magnitude but also depends on context (Kandel, Schwartz, and Jessell 1991). Analogously, in choice under risk the signals are the differences in lottery payoffs across states. Via the ordering property, the salience function σ captures the signal’s magnitude in a given state. The role of context is captured by diminishing sensitivity (and reflection): the intensity with which payoffs in a state are perceived increases as the state’s payoffs approach the status quo of zero, which is our measure of context.⁸

⁸As in Weber’s law of diminishing sensitivity, in which a change in luminosity is perceived less intensely if it occurs at a higher luminosity level, the local thinker perceives less intensely payoff differences occurring at high (absolute) payoff levels. Interestingly, visual perception and risk taking seem to be connected at a more

Consistent with psychology of attention, we assume that the decision maker evaluates lotteries by focusing on, and weighting more, their most salient states. The “local thinking” parameter $\frac{1}{\delta}$ captures the strength of the decision maker’s focus on salient states, proxying for his ability to pay attention to multiple aspects, cognitive load, or simply intelligence. Our assumption of rank-based discounting buys us analytical tractability, but our main results also hold if the distortion of the odds in (7) is a smooth increasing function of salience differences, for instance, $\delta[\sigma(x_s^i, x_s^{-i}) - \sigma(x_s^i, x_s^{-i})]$.⁹ One benefit of this alternative specification is that it would avoid discontinuities in valuation. However, discontinuities play no role in our analysis, so for simplicity we stick to ranking-based discounting.

The main substantive restriction embodied in our model is that the discounting function does not depend on a state’s probability, which implies that unlikely states are subject to the greatest distortions. This notion is also encoded in prospect theory’s weighing function, in which “highly unlikely events are either ignored or overweighted” (KT). Together with subadditivity, this feature, also present in early work on probability weighing (Edwards 1962; Fellner 1961), allows KT to account for risk-loving behavior and the Allais paradoxes. Quiggin’s (1982) rank-dependent expected utility and Tversky and Kahneman’s (1992) cumulative prospect theory (CPT) develop weighing functions in which the rank order of a lottery’s payoffs affects probability weighting.¹⁰

fundamental neurological level. McCoy and Platt (2005) show in a visual gambling task that when monkeys made risky choices neuronal activity increased in an area of the brain (CGp, the posterior cingulate cortex) linked to visual orienting and reward processing. Crucially, the activation of CGp was better predicted by the subjective salience of a risky option than by its actual value, leading the authors to hypothesize that “enhanced neuronal activity associated with risky rewards biases attention spatially, marking large payoffs as salient for guiding behavior” (p. 1226).

⁹A smooth specification would also address a concern with the current model that states with similar salience may obtain very different weights. This implies that (1) splitting states and slightly altering payoffs could have a large impact on choice, and (2) in choice problems with many states the (slightly) less salient states are effectively ignored. However, because none of our results is due to these effects, we stick to rank-based discounting for simplicity.

¹⁰Prelec (1998) axiomatizes a set of theories of choice based on probability weighting, which include CPT. For a recent attempt to estimate the probability weighting function, see Wu and Gonzalez (1996).

Our theory exhibits two sharp differences from these works. First, in our model the magnitude of payoffs, not only their rank, determines salience and probability weights: unlikely events are overweighted when they are associated with salient payoffs, but underweighted otherwise. As a consequence, the lottery upside may still be underweighted if the payoff associated with it is not sufficiently high. As we show in Section IV, this feature is crucial to explaining shifts in risk attitudes. Second, and more important, in our model decision weights depend on the choice context, namely, on the available alternatives as they are presented to the decision maker. In Section V we exploit this feature to shed light on the psychological forces behind the Allais paradoxes and preference reversals.

Our main results rely on ordering and diminishing sensitivity of σ , as well as on the comparatively larger distortion of low probabilities. We however sometimes illustrate the model by using the more restrictive salience function in equation (5), which offers a tractable case characterized by only two parameters (θ , δ). This allows us to look for ranges of θ and δ that are consistent with the observed choice patterns.

The State Space. Salience is a property of states of nature that depends on the lottery payoffs that occur in each state, as they are presented to the decision maker. The assumption that payoffs (rather than final wealth states) shape the perception of states is a form of narrow framing, consistent with the fact that payoffs are perceived as gains and losses relative to the status quo, as in prospect theory.

In our approach, the state space S and the states' objective probabilities are a given of the choice problem.¹¹ In the lab, specifying a state space for a choice problem is straightforward when the feasible payoff combinations—and their probabilities—are available, for instance, when lotteries are explicitly described as contingencies based on a randomizing device. For example, $L_1 \equiv (10, 0.5; 5, 0.5)$ and $L_2 \equiv (7, 0.5; 9, 0.5)$ give rise to four payoff combinations $\{(10, 7), (10, 9), (5, 7), (5, 9)\}$ if they are

¹¹ In particular, we do not address choice problems where outcome probabilities are ambiguous, such as the Ellsberg paradox. This is an important direction for future work.

played by flipping two separate coins, but only to two payoff combinations, e.g. $\{(10, 7), (5, 9)\}$, if they are contingent on the same coin flip. In our experiments, we nearly always describe the lotteries' correlation structure by specifying the state space. However, classic experiments such as the Allais paradoxes provide less information: they involve a choice between (standard) lotteries, and the state space is not explicitly described. In this case, we assume that our decision maker treats the lotteries as independent, which implies that the state space is the product space induced by the lotteries' marginal distributions over payoffs.¹² Intuitively, salience detects the starkest payoff differences among lotteries unless some of these differences are explicitly ruled out.

For the choice problems we study, the choice set and thus the state space are unambiguous. All our results are obtained by equating the choice set with the set of options the decision maker is actively considering (the *consideration set*). In real-world applications, however, the consideration set need not coincide with the choice set. In some situations, the decision maker may in fact consider only a subset of the available options. For example, he may discard universally dominated lotteries from his choice set before evaluating other, more attractive, lotteries. As another example, suppose that the payoffs of two lotteries are determined by the roll of the same dice. One lottery pays 1,2,3,4,5,6, according to the dice's face; the other lottery pays 2,3,4,5,6,1. The state in which the first lottery pays 6 and the second pays 1 may appear most salient to the decision maker, leading him to prefer the first lottery. Of course, a moment's thought would lead him to realize that the lotteries are just rearrangements of each other and recognize them as identical. In the following, we assume that before evaluating lotteries, the decision maker edits the choice set by discarding all but one of the lottery permutations (at random, thus preserving indifference between the permutations). Both forms of editing are plausibly related to salience itself: in these cases, before comparing payoffs, what is salient to the decision maker are the properties of permutation or dominance of certain lotteries.

¹² In Online Appendix 2 we provide experimental evidence consistent with this assumption, as well as details on the information given in our surveys.

In other situations, the consideration set may be larger than the actual choice set, for instance, when the choice problem itself reminds the decision maker of options that are not actually available for choice. In this case, options outside the choice set may influence salience and choice. This perspective provides insight into the endowment effect, as shown in Bordalo, Gennaioli, and Shleifer (2012a), and into puzzles in consumer choice such as decoy effects and context-dependent willingness to pay for goods, as shown in Bordalo, Gennaioli, and Shleifer (2012b). Here this issue arises only in our discussion of preference reversals, where we argue that when people evaluate a lottery in isolation, they might compare it to the alternative of having nothing (see Section V.B).

Endogenizing the consideration set is an important direction for future work. There is a large literature on this topic in marketing and a growing one in decision theory (e.g., Manzini and Mariotti 2007; Masatlioglu, Nakajima, and Ozbay 2012), but a consensus model has not yet emerged. In a similar spirit, the model could be generalized to take into account determinants of salience other than payoff values, such as prior experiences and details of presentation, or even color of font. These may matter in some situations but are not considered here.

Salience and Context-Dependent Choice. We are not the first to propose a model of context-dependent choice among lotteries. Rubinstein (1988), followed by Aizpurua, Nieto, and Uriarte (1990) and Leland (1994), builds a model of similarity-based preferences, in which decision makers simplify the choice among two lotteries by pruning the dimension (probability or payoff, if any) along which lotteries are similar. The working and predictions of our model are different from Rubinstein's, even though we share the idea that the common ratio Allais paradox (see Section V.A) is due to subjects' focus on lottery payoffs. Loomes (2010) proposes a model that is closely related to Rubinstein's, and presents evidence suggesting a role for probability comparisons across lotteries in choice. In regret theory (Loomes and Sugden 1982; Bell 1982; Fishburn 1982), the choice set directly affects the decision maker's utility via a regret/rejoice term added to a standard utility function. In our model, instead, context affects decisions by shaping the salience of payoffs and decision weights. Regret theory shares with our model the idea that states with

higher payoff differences have a disproportionate effect on choices. In that theory, decision makers get (dis)utility from comparing forgone outcomes. Regret theory can account for a certain type of context dependence, such as a role for correlations among lotteries; however, by adopting a traditional utility theory perspective, it cannot capture framing effects or violations of procedural invariance (Tversky, Slovic, and Kahneman 1990). Moreover, since regret theory does not feature diminishing sensitivity (as it excludes the notion of a reference point), it has a hard time accounting for standard patterns of risk-preferences, including risk-averse preferences for fair 50-50 gambles over gains and their reflection over losses.

Formal models of context-dependent choice (e.g., Fishburn 1982) may be criticized as not being falsifiable because too many choice patterns can be justified. We stress that our psychologically based assumptions of ordering and diminishing sensitivity place tight restrictions on the predictions of our model under any value (and salience) function. To give one example, both the ordering and the diminishing sensitivity property make strong predictions regarding the conditions for and the directionality of the Allais paradoxes. In particular, they imply that the independence axiom of expected utility theory should hold when the mixture lotteries are correlated (see Section V.A). To give another example, the distortion of decision weights in Definition 2 implies that pairwise choice among two or three outcome independent lotteries having the same support is transitive and that choice is consistent with first-order stochastic dominance when lotteries are independent (see Online Appendix 1). In future work, it may be useful to uncover the precise axioms consistent with Definitions 1 and 2.

IV. SALIENCE AND ATTITUDES TOWARD RISK

We first describe how salience affects the risk preferences of a local thinker with linear utility. To do so, consider the choice between a sure prospect $L_0 = (x, 1)$ and a mean preserving spread $L_1 = (x + g, \pi_g; x - l, 1 - \pi_g)$, with $g\pi_g = (1 - \pi_g)l$. All payoffs are positive. In this choice, there are two states: $s_g = (x + g, x)$, in which the lottery gains relative to the sure prospect, and $s_l = (x - l, x)$, in which the lottery loses.

Because L_1 is a mean-preserving spread of L_0 , equation (12) implies that for any $\delta < 1$, a local thinker with linear utility

chooses the lottery if and only if the gain state s_g is more salient than the loss state s_l , that is, when $\sigma(x + g, x) > \sigma(x - l, x)$. In this case, using the notation of Definition 2, the weight π_g^1 attached to the event of winning under the lottery is higher than the event's probability π_g . As a result, the local thinker perceives the expected value of L_1 to be above that of L_0 , and exhibits risk-seeking behavior, choosing L_1 over L_0 .

Using the fact that $g\pi_g = (1 - \pi_g)l$, the condition for s_g to be more salient than s_l can be written as:

$$(13) \quad \sigma\left(x + \frac{1 - \pi_g}{\pi_g} \cdot l, x\right) > \sigma(x - l, x).$$

The ordering property of salience has two implications. First, when the state s_g is very unlikely, it is also salient: at $\pi_g \approx 0$ the lottery's upside is very large, its salience is high, and equation (13) always holds. Second, the salience of s_g decreases in π_g : as the lottery wins with higher probability, its payoff gain g is lower and thus less salient. Thus, equation (13) is less likely to hold as π_g rises. The diminishing sensitivity property in turn implies that when the lottery gain is equal to the loss (i.e., $g = l$), the loss is salient. As a consequence, when $\pi_g = \frac{1}{2}$ the state s_g is less salient than s_l , so equation (13) is violated.

In sum, condition (13) identifies a probability threshold $\pi_g^* < \frac{1}{2}$ such that: for $\pi_g < \pi_g^*$ the lottery upside is salient, the local thinker overweights it and behaves in a risk-seeking way; for $\pi_g > \pi_g^*$ the lottery downside is salient, the local thinker overweights it and behaves in a risk-averse way; for $\pi_g = \pi_g^*$ states s_g and s_l are equally salient and the local thinker is risk-neutral. Remarkably, these properties of decision weights recover key features of prospect theory's inverse S-shaped probability weighting function: overweighting of low probabilities, and underweighting of high probabilities. Figure II shows the decision weight π_g^1 as a function of probability π_g . Low probabilities are overweighted because they are associated with salient upsides of long-shot lotteries. High probabilities are underweighted as they occur in lotteries with a small, nonsalient, upside.

Note however that in our model the weighting function is context-dependent. In contrast to prospect theory, overweighting depends not only on the probability of a state but also on the salience of its payoff in equation (13). In particular, overweighting is shaped by the average level of payoffs x . To see this, denote




FIGURE II
Context-Dependent Probability Weighting Function

by $r = v^{LT}(L_0) - v^{LT}(L_1)$ the “premium” required by the local thinker to be indifferent between the risky lottery L_1 and the sure prospect L_0 (r is positive when the local thinker is risk-averse). For a rational decision maker with linear utility, $r = 0$ regardless of the payoff level x . To see how the local thinker’s risk attitudes depend on x , consider the following definition.

DEFINITION 3. A salience function is convex if, for any state with positive payoffs (y, z) and any $x, \epsilon > 0$, the difference $\sigma(y + x, z + x) - \sigma(y + x + \epsilon, z + x + \epsilon)$ is a decreasing function of the payoff level x . A salience function is concave if this difference increases in x .

A salience function is convex if diminishing sensitivity becomes weaker as the payoff level x rises. The Appendix then proves:

LEMMA 1. If the salience function is convex, then $r = v^{LT}(L_0) - v^{LT}(L_1)$ weakly decreases with x . Conversely, if the salience function is concave then r weakly increases with x .

If convexity holds and diminishing sensitivity becomes weaker with x , then a higher payoff level weakly reduces r , increasing the valuation of the risky lottery L_1 relative to that of the safe lottery L_0 . In equation (13), this increases the threshold π_g^* , boosting risk seeking. If instead diminishing sensitivity becomes stronger with x , a higher payoff level leads to an increase in r , weakly decreasing L_1 's valuation relative to that of L_0 . In equation (13) this reduces the threshold π_g^* , hindering risk seeking.

The salience function of equation (5) satisfies convexity. Using this function, the condition (13) for s_g to be more salient than s_l becomes:

$$(14) \quad \left(x + \frac{\theta}{2}\right)(1 - 2\pi_g) > l \cdot (1 - \pi_g),$$

which is indeed more likely to hold for higher x (so long as $\pi_g < \frac{1}{2}$).

Equation (14) implies that, holding the lottery loss l constant, risk attitudes follow Figure III (where for convenience we set $\theta \cdot l \simeq 0$). As x rises, the threshold π_g^* below which the decision maker is risk seeking increases, so that risk-seeking behavior can occur even at relatively high probabilities π_g (though never for $\pi_g > \frac{1}{2}$).

We tested the predictions illustrated in Figure III by giving experimental subjects a series of binary choices between a mean preserving spread $L_1 = (x + g, \pi_g; x - l, 1 - \pi_g)$ and a sure prospect $L_2 = (x, 1)$. We set the downside of L_1 at $l = \$20$, yielding an upside g of $\$20 \cdot \frac{(1 - \pi_g)}{\pi_g}$. We varied x in $\{\$20, \$100, \$400, \$2100, \$10500\}$ and π_g in $\{0.01, 0.05, 0.2, 0.33, 0.4, 0.5, 0.67\}$. For each of these 35 choice problems, we collected at least 70 responses. On average, each subject made five choices, several of which held either π_g or x constant. The observed proportion of subjects choosing the lottery for every combination (x, π_g) is reported in Table I; for comparison with the predictions of Figure III, the results are shown in Figure IV.

The patterns are qualitatively consistent with the predictions of Figure III. First, and crucially, for any given expected value x , the proportion of risk takers falls as π_g increases and there is a large drop in risk taking as π_g crosses 0.5. This prediction is consistent with the probability weighting function depicted in Figure II. Second, for a given $\pi_g < 0.5$, the proportion

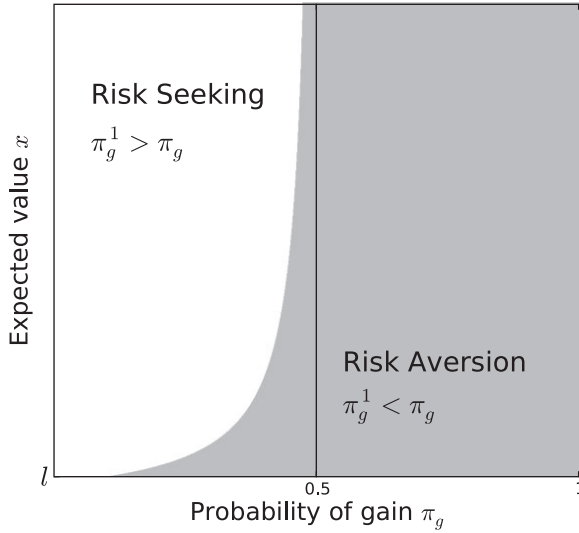


FIGURE III
Shifts in Risk Attitudes

TABLE I
PROPORTION OF RISK-SEEKING SUBJECTS

Expected value x	\$10500	0.83	0.65	0.50	0.48	0.46	0.33	0.23
	\$2100	0.83	0.65	0.48	0.43	0.48	0.38	0.21
	\$400	0.60	0.58	0.44	0.47	0.33	0.30	0.23
	\$100	0.58	0.54	0.40	0.32	0.22	0.30	0.13
	\$20	0.15	0.2	0.12	0.08	0.10	0.25	0.15
		0.01	0.05	0.2	0.33	0.4	0.5	0.67
				Probability of gain π_g				

of risk takers increases with the expected value x . The effect is statistically significant: at $\pi_g = 0.05$ a large majority of subjects (80%) are risk-averse when $x = \$20$, but as x increases to \$2100 a large majority (65%) becomes risk-seeking. This finding is consistent with the finer hypothesis, encoded in equation (5), that diminishing sensitivity may become weaker at higher payoff levels. The increase in x raises the proportion of risk takers from around 10% to 50% even for moderate probabilities in the range (0.2,0.4).




FIGURE IV
Proportion of Risk-Seeking Subjects

Although not a formal test of our theory, these patterns are broadly consistent with the predictions of our model.¹³ Online Appendix 2 describes additional experiments on long-shot lotteries whose results are also consistent with our model but inconsistent with prospect theory under standard calibrations of the value function. In that Appendix we show that, using the salience function in (5), the parameter values $\delta \sim 0.7$ and $\theta \sim 0.1$ are consistent with the evidence on risk preferences, as well as with risk preferences concerning long-shot lotteries. These values are not a formal calibration, but we employ them as a useful reference for discussing Allais paradoxes in the next section.

¹³ The weighing function of prospect theory and CPT can explain why risk seeking prevails at low π_g , but not the shift from risk aversion to risk seeking as x rises. To explain this finding, both theories need a concave value function characterized by strongly diminishing returns. In Online Appendix 2 we provide further support for these claims by showing that standard calibrations of prospect theory cannot explain our experimental findings. For example, the calibration in Tversky and Kahneman (1992) features the value function $v(x) = x^{0.88}$, which is insufficiently concave. Importantly, calibrations of the value function are notoriously unstable: using two other sets of choice data, Wu and Gonzalez (1996) estimate $v(x) = x^{0.5}$ and $v(x) = x^{0.37}$, respectively. The fact that calibration is so dependent on the choice context suggests that choice itself is context dependent.

V. LOCAL THINKING AND CONTEXT DEPENDENCE

V.A. *The Allais Paradoxes*

The “common consequence” Allais Paradox. Let us go back to the Allais paradox described in Section II. We now describe the precise conditions under which our model can explain it. Recall that subjects are asked to choose between the lotteries:

$$(15) \quad \begin{aligned} L_1(z) &= (2500, 0.33; \quad 0, 0.01; \quad z, 0.66), \\ L_2(z) &= (2400, 0.34; \quad z, 0.66) \end{aligned}$$

for different values of z . For $z = 2400$, most subjects are risk-averse, preferring $L_2(2400)$ to $L_1(2400)$, while for $z = 0$, most subjects are risk-seeking, preferring $L_1(0)$ to $L_2(0)$.

When $z = 2400$, the minimal state space is $S = \{(2500, 2400), (0, 2400), (2400, 2400)\}$. The most salient state is one where the risky lottery $L_1(2400)$ pays zero because, by ordering and diminishing sensitivity we have:

$$(16) \quad \sigma(0, 2400) > \sigma(2500, 2400) > \sigma(2400, 2400).$$

By equation (12), a local thinker then prefers the riskless lottery $L_2(2400)$ provided:

$$(17) \quad -(0.01) \cdot 2400 + \delta \cdot (0.33) \cdot 100 < 0,$$

which holds for $\delta < 0.73$. Although the risky lottery $L_1(2400)$ has a higher expected value, it is not chosen when the degree of local thinking is severe, because its downside of 0 is very salient.

Consider the choice between $L_1(0)$ and $L_2(0)$. Now both options are risky and, as discussed in Section III, the local thinker is assumed to see the lotteries as independent. The minimal state space now has four states of the world, that is, $S = \{(2500, 2400), (2500, 0), (0, 2400), (0, 0)\}$, whose salience ranking is:

$$(18) \quad \sigma(2500, 0) > \sigma(0, 2400) > \sigma(2500, 2400) > \sigma(0, 0).$$

The first inequality follows from ordering, and the second from diminishing sensitivity. By equation (12), a local thinker prefers the risky lottery $L_1(0)$ provided:

$$(19) \quad \begin{aligned} &(0.33) \cdot (0.66) \cdot 2500 - \delta \cdot (0.67) \cdot (0.34) \cdot 2400 \\ &+ \delta^2 \cdot (0.33) \cdot (0.34) \cdot 100 > 0 \end{aligned}$$

which holds for $\delta \geq 0$. Any local thinker with linear utility chooses the risky lottery $L_1(0)$ because its upside is very salient.

In sum, when $\delta < 0.73$ a local thinker exhibits the Allais paradox. This is true for any salience function satisfying ordering and diminishing sensitivity, and thus also for the parameterization $\delta = 0.7$, $\theta = 0.1$ obtained when using (5). It is worth spelling out the exact intuition for this result. When $z = 2400$, the lottery $L_2(2400)$ is safe, whereas the lottery $L_1(2400)$ has a salient downside of zero. The local thinker focuses on this downside, leading to risk aversion. When instead $z = 0$, the downside payoff of the safer lottery $L_2(0)$ is also 0. As a result, the lotteries' upsides are now crucial to determining salience. This induces the local thinker to overweight the larger upside of $L_1(0)$, triggering risk seeking. The salience of payoffs thus implies that when the same downside risk is added to the lotteries $L_1(2400)$ and $L_2(2400)$, the sure prospect $L_2(2400)$ is particularly hurt because the common downside payoff induces the decision maker to focus on the larger upside of the risky lottery, leading to risk-seeking behavior. This yields the "certainty effect" of prospect theory and CPT as a form of context dependence due to payoff salience.

This role of context dependence invites the following test. Suppose that subjects are presented the following correlated version of the lotteries $L_1(z)$ and $L_2(z)$ in equation (15):

Probability	0.01	0.33	0.66
payoff of $L_1(z)$	0	2500	z
payoff of $L_2(z)$	2400	2400	z

where the table specifies the possible joint payoff outcomes of the two lotteries and their respective probabilities. Correlation changes the state space *but not* a lottery's distribution over final outcomes, so it does not affect choice under either expected utility theory or prospect theory. Critically, this is not true for a local thinker: the context of this correlated version makes clear that the state in which both lotteries pay z is the least salient one, and also that it drops from evaluation in equation (12), so that the value of z should not affect the choice at all. This is due to the ordering property: states where the two lotteries yield the same payoff are the least salient ones and in fact cancel out in the local thinker's valuation (ordering leads to them being "edited out" by the local thinker). That is, in our model—but not in prospect

theory—the Allais paradox should not occur when $L_1(z)$ and $L_2(z)$ are presented in the correlated form equation (20).

We tested this prediction by presenting experimental subjects correlated formats of lotteries $L_1(z)$ and $L_2(z)$ for $z = 0$ and $z = 2400$. The observed choice pattern is the following:

	$L_1(2400)$	$L_2(2400)$
$L_1(0)$	7%	9%
$L_2(0)$	11%	73%

The vast majority of subjects do not reverse their preferences (80% of choices lie on the NW-SE diagonal), and most of them are risk-averse, which in our model is also consistent with the fact that (0, 2400) is the most salient state in the correlated choice problem (20). Among the few subjects reversing their preference, no clear pattern is detectable. This contrasts with the fact that our experimental subjects exhibit the Allais paradox when lotteries are presented in an uncorrelated form (see Online Appendix 2). Thus, when the lotteries pay the common consequence in the same state, choice is invariant to z and the Allais paradox disappears. Our model accounts for this fact because, as the common consequence z is made evident by correlation, it becomes nonsalient. As a result, subjects prune it and choose based on the remaining payoffs.¹⁴

This result captures the Savage (1972, p. 102) argument in defense of the normative character of the “sure thing principle,” and validates his thought experiment. Other experiments in the literature are consistent with our results. Conlisk (1989) examines a related variation of the Allais choice problem, in which

¹⁴ We tested the robustness of the correlation result by changing the choice problem in several ways: (1) we framed the correlations verbally (e.g., described how the throw of a common die determined both lotteries’ payoffs), (2) we repeated the experiment with uncertain real-world events, instead of lotteries, and (3) we varied the ordering of questions, the number of filler questions, and payoffs. As Online Appendix 2 shows, our results are robust to all these variations. We also ran an experiment where subjects were explicitly presented the lotteries of equation (15) with $z = 2400$ as uncorrelated, with a state space consisting of the four possible states. The choice pattern exhibited by subjects is: (1) very similar to the one exhibited when the state space is not explicitly presented, validating our basic assumption that a decision maker assumes the lotteries to be uncorrelated when this is not specified otherwise, and (2) very different from the choice pattern exhibited under correlation (with 35% of subjects changing their choice as predicted by our model, see Online Appendix 2).

each alternative is given in compound form involving two simple lotteries, with one of the simple lotteries yielding the common consequence z . Birnbaum and Schmidt (2010) present the Allais problem in split form, singling out the common consequence z in each lottery. In both cases, the Allais reversals subside. Our model also rationalizes the disappearance of the Allais paradox in Colinsk's second example, which uses nonboundary lotteries. See also Harrison (1994) for related work on the common consequence paradox.

The "common ratio" Allais Paradox. We now turn to the "common ratio" paradox, which occurs in the choice between lotteries:

$$(21) \quad L_1(\pi') = (6000, \pi'; 0, 1 - \pi'), \quad L_2(\pi) = (\alpha \cdot 6000, \pi; 0, 1 - \pi),$$

where $L_1(\pi')$ is riskier than $L_2(\pi)$ in the sense that it pays a larger positive amount ($\alpha < 1$) with a smaller probability ($\pi' < \pi$). By the independence axiom, an expected utility maximizer with utility function $v(\cdot)$ chooses the safer lottery $L_2(\pi)$ over $L_1(\pi')$ when:

$$(22) \quad v(\alpha \cdot 6000) \geq \frac{\pi'}{\pi} \cdot v(6000) + v(0) \left(1 - \frac{\pi'}{\pi}\right).$$

The choice should not vary with probabilities π, π' as long as $\frac{\pi'}{\pi}$ is kept constant. A stark case arises when $\frac{\pi'}{\pi} = \alpha$; now the two lotteries have the same expected value and a risk-averse expected utility maximizer always prefers the safer lottery $L_2(\pi)$ to $L_1(\pi')$ for any π . Parameter α identifies the "common ratio" between π' and π at different levels of π .

It is well known that, contrary to the expected utility theory, the choices of experimental subjects depend on the value of π : for fixed $\frac{\pi'}{\pi} = \alpha = 0.5$, when $\pi = 0.9$ subjects prefer the safer lottery $L_2(0.9) = (3000, 0.9; 0, 0.1)$ to $L_1(0.45) = (6000, 0.45; 0, 0.55)$. When instead $\pi = 0.002$, subjects prefer the riskier lottery $L_1(0.001) = (6000, 0.001; 0, 0.999)$ to $L_2(0.002) = (3000, 0.002; 0, 0.998)$. This shift toward risk seeking as the probability of winning falls has provided one of the main justifications for the introduction of the probability weighting function. In fact, KT account for this evidence by assuming that this function grows slower than linearly for small π ; hence, $\alpha\pi$ is overweighted relatively to π at low values of π , inducing the choice of $L_1(\pi')$ when $\pi = 0.002$.

Consider the choice between $L_1(\pi')$ and $L_2(\pi)$ in our model. For $\alpha = \frac{1}{2}$ there are four states of the world, $S = \{(6000, 3000), (0, 3000), (6000, 0), (0, 0)\}$. Once more, ordering and diminishing sensitivity suffice to imply that the salience ranking among states is

$$(23) \quad \sigma(6000, 0) > \sigma(0, 3000) > \sigma(6000, 3000) > \sigma(0, 0).$$

It is convenient to express the local thinker’s decision as a function of the transformed probabilities of the lottery outcomes (as opposed to those of states of the world).¹⁵ Denoting these transformed probabilities by $\hat{\pi}'$ and $\hat{\pi}$, we find that the local thinker evaluates the odds with which the riskier lottery $L_1(\pi')$ pays out relative to the safer one $L_2(\pi)$ as:

$$(24) \quad \frac{\hat{\pi}'}{\hat{\pi}} = \frac{\pi'}{\pi} \cdot \frac{(1 - \pi) + \pi\delta^2}{(1 - \pi')\delta + \pi'\delta^2}.$$

With a linear utility, the local thinker selects the safer lottery $L_2(\pi)$ if and only if $\frac{\hat{\pi}'}{\hat{\pi}} \leq \frac{1}{2}$. This implies that the local thinker chooses the safer lottery when:

$$(25) \quad \pi \geq \frac{2(1 - \delta)}{2 - \delta - \delta^2}.$$

As in the common ratio effect, the local thinker is risk-averse when π is sufficiently high and risk-seeking otherwise. In particular, for $\delta \in (0.22, 1)$, the local thinker switches from $L_2(0.9)$ to $L_1(0.001)$ just as experimental subjects do. This is true for any salience function satisfying ordering and diminishing sensitivity and thus also for the parameterization $\delta = 0.7, \theta = 0.1$ of equation (5).

The intuition for this result (see Proposition 1) is that salience exerts a particularly strong effect in low-probability states. The upside of the riskier lottery $L_1(\pi')$ is salient at every π , creating a force toward risk seeking. Crucially, however, this force is strong precisely when π is low. In this case, the greater salience of the risky lottery’s upside blurs the small probability difference $\pi - \pi' = (1 - \alpha)\pi$ between the two lotteries. When instead π is large, the decision maker realizes that the risky lottery is much

¹⁵ From any vector of state-specific decision weights $(\pi_s^i)_{s \in S}$, the decision weight $\pi^i(x)$ attached to lottery i ’s payoff x is equal to the sum of the decision weights of all states where lottery i pays payoff x . Formally, $\pi^i(x) = \sum_{s \in S_x^i} \pi_s^i$ where S_x^i is the set of states where i pays x .

more likely to pay nothing, inducing him to attach a large weight on the second most salient state (0, 3000). This is what drives the choice of the safe lottery $L_2(\pi)$.

Experimental evidence shows that this common ratio effect is also not robust to the introduction of correlation. KT asked subjects to choose between two lotteries of the type (23) in a two-stage game where in the first stage there is a 75% probability of the game ending without any winnings and a 25% change of going to stage two. In stage two, the lottery chosen at the outset is played out. The presence of the first stage is equivalent to reducing by 75% the winning probability for both lotteries, so in terms of final outcomes this setting is equivalent to the setting that leads to the common ratio effect. Crucially, KT document that in this formulation there is no violation of the independence axiom.

In explaining this behavior, KT informally argue that individuals “edit out” the correlated first-stage state where both lotteries pay zero. Our model yields this editing as a by-product of the low salience and cancellation of this state. Adding a correlated state where both lotteries pay 0 neither affects the salience ranking in equation (23) nor—more importantly—the odds ratios between states. As a result, the local thinker chooses as if he disregards the correlated state and its probability. This is what experimental subjects do.

In sum, our model explains the Allais paradoxes as the by-product of a specific form of context dependence working through the salience of lottery payoffs. Adding a common payoff to all lotteries or rescaling their probabilities changes risk preferences by changing the salience and the weighting of the lotteries’ upsides or downsides. These effects depend on how the lotteries are presented. Adding a common payoff or rescaling probabilities by introducing into the lotteries a correlated state does not affect choice: it is too enticing for subjects to disregard this state and abide by the independence axiom.

V.B. Preference Reversals

Context dependence in our model can also explain the phenomenon of preference reversal described by Lichtenstein and Slovic (1971) and confirmed by Grether and Plott (1979) and Tversky, Slovic, and Kahneman (1990). Subjects are asked to choose between a safer lottery L_π , which has a high probability of a low payoff, and a riskier lottery $L_\$$, which has a low

probability of a high payoff (we use conventional notation for the lotteries). Subjects may systematically choose the safer lottery L_π and yet state a higher minimum selling price for the riskier lottery $L_\$$. Preferences as revealed by choice are thus the opposite of preferences as revealed by pricing, leading to claims that choosing and pricing follow two fundamentally different principles. Neither prospect theory nor expected utility theory can rationalize preference reversals.

To study preference reversals in our model, consider how a local thinker prices a lottery. Given that in our model valuation is context-dependent, the concept of a minimum selling price can be interpreted in two distinct ways.¹⁶ Under the *valuation* approach, in jointly evaluating two lotteries $\{L_1, L_2\}$, the minimum selling price for either of them is the lottery's monetary valuation obtained by using the decision weights determined in $\{L_1, L_2\}$ according to Definitions 1 and 2. Formally, a local thinker with a value function $v(\cdot)$ prices L_1 at:

$$(26) \quad P_{\min}(L_1|L_2) = v^{-1} \left[\sum_{s \in S} \pi_s^1 v(x_s^1) \right],$$

where π_s^1 is the decision weights of state s for lottery L_1 in the context of its choice from the set $\{L_1, L_2\}$. Here salience is determined by the consideration set $\{L_1, L_2\}$, which does not coincide with the local thinker's choice set (since he is not choosing between the two lotteries, he is pricing them). With a linear value function, the price $P_{\min}(L_1|L_2)$ is the expected value of L_1 as perceived by the local thinker. If the local thinker is asked to price a lottery in isolation, this approach suggests that he evaluates it together with the alternative of not having the lottery, namely, having zero for sure, $L_0 \equiv (0, 1)$. We see this as a natural way to model the elicitation of minimum

¹⁶ Several authors have studied preference reversals by focusing on the details of the experimental procedure, in particular the incentive structures. Thus, Karni and Safra (1987) suggest that Grether and Plott's BDM (Becker, DeGroot, Marschak) mechanism should be interpreted as a choice between two-stage lotteries, and that preference reversals follow from violations of the independence axiom required to interpret elicited prices as certainty equivalents (see also Holt 1986). Segal (1988) suggests instead it results from violations of the reduction axiom. Evidently, the interpretation of experimental procedures is an important factor even in models which preserve transitivity and are not context-dependent.

selling prices in cases where—as in most preference reversal experiments—subjects must state this price (potentially under an incentive scheme).

Alternatively, under the *revealed preference* approach, the minimum selling price is found by revealed preference: the price of lottery L_1 is then the minimum amount of money c_1 such that adding the sure prospect $L_c = (c_1, 1)$ to the choice set makes the local thinker weakly prefer the sure prospect. In this approach, salience is determined by the decision maker’s actual choice set $\{L_1, L_c\}$. In this section we adopt the valuation approach to study reversals. In Online Appendix 1, we show that our model can also yield reversals using the revealed preference approach, but under more restricted circumstances than under the valuation approach. Consistent with the difference between the two approaches, experiments that *explicitly* implemented the revealed preference approach found significantly lower levels of reversals (Bostic, Herrnstein, and Luce 1990; Tversky, Slovic, and Kahneman 1990).

In the preference reversal experiments, subjects are first asked to price in isolation, and then to choose among, the following two independent lotteries:

$$L_\$ = \begin{cases} x, & \text{with prob. } \pi' \\ 0, & 1 - \pi' \end{cases}, L_\pi = \begin{cases} \alpha x, & \text{with prob. } \pi \\ 0, & 1 - \pi \end{cases} \tag{27}$$

where typically $\frac{\pi'}{\pi} = \alpha = \frac{1}{2}$, as in the common ratio experiments. We know from equation (25) that, with linear utility, the local thinker selects the safer lottery L_π when $\pi > \frac{2(1-\delta)}{2-\delta-\delta^2}$. In the literature, we typically have $\pi > \frac{3}{4}$, so this constraint holds for any $\delta \geq \frac{2}{3}$. Thus, when asked to choose, a local thinker having linear utility and $\delta = 0.7$ is risk-averse and prefers L_π to $L_\$$, just as most experimental subjects do.

In contrast, when the local thinker is asked to price the lotteries in isolation, he evaluates each lottery relative to $L_0 = (0, 1)$. In this comparison, each lottery’s upside is salient. As a consequence, since $\alpha = \frac{1}{2}$ the local thinker prices the lotteries as:

$$P(L_\pi|L_0) = \frac{x}{2} \cdot \frac{\pi}{\pi + (1 - \pi)\delta} \quad P(L_\$|L_0) = x \cdot \frac{\frac{\pi}{2}}{\frac{\pi}{2} + \left(1 - \frac{\pi}{2}\right)\delta} \tag{28}$$

For any $\delta < 1$, the local thinker prices L_\S higher than L_π in isolation, that is,

$$P(L_\S|L_0) > P(L_\pi|L_0) .$$

Both lotteries are priced above their expected value, but L_\S is more overpriced than L_π because it pays a higher gain with a smaller probability, and from Proposition 1 we know that lower probabilities are relatively more distorted.¹⁷

Thus, although in a choice context the local thinker prefers the safer lottery L_π , in isolation he prices the risky lottery L_\S higher, exhibiting a preference reversal. Crucially, this behavior is not due to the fact that choosing and pricing are different operations. In fact, in our model choosing and pricing are the same operation, as in standard economic theory. Preference reversals occur because, unlike in standard theory, evaluation in our model is context-dependent. Pricing and choosing differ because the underlying evaluations are performed in different contexts. One noteworthy feature of our model is that it generates preference reversals through violations of “procedural invariance,” defined by Tversky, Slovic, and Kahneman (1990) as situations in which a subject prices a lottery above its expected value, $P(L_1|L_0) > \mathbb{E}(x_s^1)$, and yet prefers the expected value to the lottery, $L_1 \prec (\mathbb{E}(x_s^1), 1)$. These authors show that the vast majority of observed reversals follow from the violations of procedural invariance, as predicted by our model. Regret theory can also generate preference reversals, using the revealed preference approach to determine certainty equivalents (Loomes and Sugden, 1983). These reversals are not due to violations of procedural invariance (in contrast to the evidence) but to intransitivity in choice.

One distinctive implication of our context-based explanation is that reversals between choice and pricing should only occur when pricing takes place in isolation but not if decision makers price lotteries in the choice context itself. We tested this hypothesis by giving subjects a choice between lotteries $L_\S = (16, 0.31)$;

¹⁷ These predictions are borne out by the literature as well as by our own experimental data. Tversky, Slovic, and Kahneman (1990) show that preference reversals follow from overpricing of L_\S in isolation, and that L_π is not underpriced. Our model predicts that decision makers price L_π close to its expected value because it offers an extremely high probability of winning, which is hardly distorted.

0, 0.69) and $L_\pi = (4, 0.97; 0, 0.03)$, which Tversky, Slovic, and Kahneman (1990) found to lead to a high rate of preference reversals. Subjects stated their certainty equivalents for the two lotteries, in isolation and in the context of choice.¹⁸ Our model then predicts that preference reversal should occur between choosing and pricing in isolation, but not between choosing and pricing in the choice context.¹⁹

Despite considerable variation in subjects' evaluations (which is a general feature of such elicitation, see Grether and Plott 1979; Bostic, Herrnstein, and Luce 1990; Tversky, Slovic, and Kahneman 1990), the results are consistent with our predictions. First, among the subjects who chose L_π over $L_\$$, the average (*avg*) price of L_π in isolation was lower than the average price of $L_\$$ in isolation:

$$\text{avg}[P(L_\pi|L_0)] = 4.6 < \text{avg}[P(L_\$|L_0)] = 5.2.$$

Thus, our subject pool exhibits the standard preference reversal between choice and average pricing in isolation.²⁰

Second, preference reversals subside when we compare choosing and pricing in the choice context. In fact, in this context the same subjects priced their chosen lottery L_π higher, on average, than the alternative risky lottery $L_\$$:

$$\text{avg}[P(L_\pi|L_\$)] = 4.3 > \text{avg}[P(L_\$|L_\pi)] = 4.1.$$

¹⁸ In our experimental design, each subject priced each lottery only once, and different lotteries were priced in different contexts. This design ensures that subjects do not deform their prices to be consistent with their choices; however, it also implies that preference reversals are not observed within-subject but only at the level of price distributions across subject groups (see Online Appendix 2 for more details).

¹⁹ This prediction does not change if we allow for the option $L_0 \equiv (0, 1)$ to be included in the choice set. See Online Appendix 1 for details on choice among more than two lotteries.

²⁰ This reversal holds not only with respect to average prices but also for the distribution of prices we observe. Assuming that subjects draw evaluations randomly from the price distributions, we estimate that around 54% of the subjects who choose L_π would exhibit the standard preference reversals (see Online Appendix 2). The average prices above imply that some subjects priced the safer lottery L_π above its highest payoff. Such overpricing can occur even in a laboratory setting and with incentive schemes (Grether and Plott 1979; Bostic, Herrnstein, and Luce 1990), perhaps due to misunderstanding of the pricing task. In Online Appendix 2 we consider truncations of the data that filters out such overpricing.

As predicted by our model, in the choice context the average price ranking is consistent with choice.²¹ One may object that this agreement is caused by the subjects' wish to be coherent when they price just after a choice. However, each subject priced only one of the lotteries in the choice context.²² It appears to be the act of comparing the lotteries that drives their evaluation during choice, and not (only) an adjustment of value subsequent to choice.

Another potential objection is that our experiments do not elicit true selling prices. It is well known that it is difficult to design price elicitation mechanisms for subjects who violate the independence axiom of expected utility theory. To avoid these problems, Cox and Epstein (1989) study preference reversals by only eliciting the *ranking* of selling prices across lotteries. In their experiments, Cox and Epstein directly compared lotteries to each other, so their procedure can be viewed as eliciting evaluations in the context of choice. They find some evidence of preference reversals, but crucially they show that these reversals are equally likely in both directions (from risk-averse choice to risk-seeking pricing, and from risk-seeking choice to risk averse pricing). Symmetric reversal patterns are typically attributed to arbitrary fluctuations in evaluation, see Bostic, Herrnstein, and Luce (1990) (although Cox and Epstein interpreted them as akin to a violation of procedural invariance). Thus we interpret Cox and Epstein's results as consistent with our prediction that systematic preference reversals subside when prices are elicited in a choice context.

These results suggest that choice and pricing may follow the same fundamental principle of context-dependent evaluation. Preferences based on choice could differ from those inferred

²¹ In our data, the distribution for $P(L_\pi|L_\$)$ does not dominate that for $P(L_\$|L_\pi)$. This is due to the fact that: (1) on average subjects attribute similar values to both lotteries in the choice context, and (2) there is substantial variability in choice (and thus in pricing), as about half the subjects chose each lottery. In Online Appendix 2 we look in a more detailed way at the manifestation and significance of fact (2) in light of Tversky, Slovic, and Kahneman's (1990) analysis of preference reversals.

²² We ran another version of the survey where we asked the subjects to price the lotteries under comparison but without having to choose between them. In line with the valuation approach, these subjects exhibited similar behavior on average, namely, pricing $L_\$$ higher than L_π in isolation, but similarly to L_π under comparison.

from pricing *in isolation* because they represent evaluations made in different contexts.

V.C. Reflection and Framing Effects

KT show that experimental subjects tend to shift from risk aversion to risk seeking as gains are reflected into losses. Our model yields these shifts in risk attitudes solely based on the salience of payoffs, without relying on the S-shaped value function of prospect theory. To see this, consider the choice between lottery $L_1 = (x_s^1, \pi_s)_{s \in S}$ and sure prospect $L_2 = (x, 1)$, both of which are defined over gains (i.e., $x_s^1, x > 0$) and have the same expected value $\mathbb{E}(x_s^1) = x$. For a local thinker with linear value function:

$$(29) \quad V^{LT}(L_1) = \sum_{s \in S} \pi_s \omega_s^1 x_s^1 = \mathbb{E}(x_s^1) + \text{cov}[\omega_s^1, x_s^1]$$

where $\text{cov}[\omega_s^1, x_s^1] = \sum_{s \in S} \pi_s [\omega_s^1 - 1][x_s^1 - x]$ (recall that $\mathbb{E}(\omega_s^1) = 1$). Thus, the local thinker is risk-averse, choosing L_2 over L_1 , when $\text{cov}[\omega_s^1, x_s^1] < 0$. If L_1 and L_2 are reflected into lotteries $L'_1 = (-x_s^1, \pi_s)_{s \in S}$ and $L'_2 = (-x, 1)$, property 3 in Definition 1 implies that the salience ranking among states does not change. As a result, the same decision maker is risk-seeking, choosing L'_1 over L'_2 when:

$$(30) \quad \text{cov}[\omega_s^1, -x_s^1] = -\text{cov}[\omega_s^1, x_s^1] > 0,$$

which is fulfilled if and only if the decision maker was originally risk-averse. Intuitively, a salient downside inducing risk aversion in the gain domain becomes a salient upside inducing risk seeking in the loss domain. Our model thus yields the fourfold pattern of risk preferences (reviewed below) without assuming a value function that is concave for gains and convex for losses. With the same logic our model can account for Tversky and Kahneman's (1981) famous framing experiments, even with a linear value function.

Consider the public health dilemma, which describes the outbreak of a disease that is expected to kill 600 people. When the choice between medical responses is framed in terms of lives saved (respectively lost), the local thinker specifies the payoffs as gains (respectively losses). In the lives saved frame, the most salient outcome is the one where nobody is saved, leading to a risk-averse choice, while in the lives lost frame, the most salient outcome is the one where nobody dies, triggering a risk-seeking choice.

Note, however, that in our model reflection of risk attitudes is a knife-edge property: it holds only if the decision maker's value function is linear. A concave value function $v(\cdot)$ in the loss domain would play against the reflection of salient payoffs, creating an intrinsic preference for a moderate and certain loss. The distinction between the salience of payoffs and the curvature of the value function can provide insight into findings that reflection of risk attitudes is only partial and decreases with payoff magnitude (Laury and Holt 2005).

VI. TAKING STOCK

We now take stock by summarizing the role of different assumptions in generating our results and by comparing our predictions to those of prospect theory. Denote by "Ordering" the ordering property, by "DS" the diminishing sensitivity property, and by "Odds" the property of Definition 2 that distortions do not depend on probability odds. Table II summarizes how our model and prospect theory account for Allais paradoxes and preference reversals.²³

Anomalies in our model can be driven by a change in the salience ranking of payoffs (such as the Allais common consequence paradox) or by the differential distortion of small probabilities (such as the Allais common ratio paradox, where the salience ranking does not change). The ordering property plays a crucial role throughout, determining the direction of the choice anomalies.

By providing insight into what drives the anomalies, the model also identifies circumstances where the anomalies disappear. Crucially, these follow from the same properties that cause the anomalies to begin with. Anomalies disappear when the choice problems are set up so that the representation of each lottery is stable for the decision maker across different treatments—for example, when the common consequence is made evident in the Allais paradox.

The situation with prospect theory is very different. Each Allais paradox is explained through a different assumption

²³Tables II and III list the main properties that drive each effect in either theory, but are not exhaustive. We use the 1979 version of prospect theory, but the cumulative version (Tversky and Kahneman 1992) is very similar for the purpose of this comparison.

TABLE II
TAKING STOCK OF ANOMALIES

	Saliency theory	Prospect theory
Allais common consequence	Ordering and DS	Subcertainty of $\pi(p)$
Allais common consequence (correlated)	Ordering	Editing (but not explicit)
Allais common ratio	Ordering, DS and odds	Subadditivity of $\pi(p)$
Allais common ratio (correlated)	Ordering	Isolation effect (editing)
Preference reversals	Ordering, DS and odds	No
Preference reversals (choice)	Ordering	No

Notes. p is probability, $\pi(\cdot)$ is probability weighting function.

about the probability weighting function, or on the editing process which is not formally modeled. Finally, prospect theory cannot account for preference reversals, since choice follows from context-independent evaluation.

Consider now another well-documented choice pattern: the fourfold pattern of risk preferences (Tversky and Kahneman 1992): risk aversion (RA) for gains of high probability, risk seeking (RS) for gains of low probability, and the reverse for losses. As shown in Section IV for gains, and in Section V.C for losses, our model reproduces this pattern based solely on the properties of saliency (including that diminishing sensitivity depends on the magnitude of payoff level, and not their sign, as encoded in the reflection property of Definition 1). We further predict that risk attitudes should depend on the payoff level x . In light of the experimental results of Section IV, we adopt the convexity property of Definition 3, whereby diminishing sensitivity gets weaker as payoff levels increase. These predictions are summarized in Table III for choices between a sure payoff and a mean-preserving spread.

In prospect theory, the main driver of risk attitudes is the curvature of the value function. As discussed in Section IV, different patterns of risk attitudes put different constraints on this function, which can be hard to reconcile. Our context-dependent account of risk preferences does not require any assumptions on the curvature of the value function. Instead, that role is taken by the diminishing sensitivity property of the saliency function.

TABLE III
TAKING STOCK OF RISK ATTITUDES

	Saliency theory	Prospect theory
<i>For gains</i>		
RA for high p	DS	Concave $v(\cdot)$ and subcertainty of $\pi(p)$
RS for low p , high x	Ordering	$v(\cdot)$ low curvature, $\pi(p) > p$ for small p
x -dependent switch to RS	Ordering, DS	Nonlinear $v(\cdot)$
RA for low p , low x	DS, convexity	$v(\cdot)$ very concave for low x
<i>For losses</i>		
RS for high p	DS	Convex $v(\cdot)$ and subcertainty of $\pi(p)$
RA for low p , high $ x $	Ordering	$v(\cdot)$ low curvature, $\pi(p) > p$ for small p
x -dependent switch to RA	Ordering, DS	Nonlinear $v(\cdot)$
RS for low p , low $ x $	DS, convexity	$v(\cdot)$ very convex for low $ x $

Notes. RA stands for risk aversion; RS stands for risk seeking p is probability; $\pi(p)$ is KT's probability weighting function.

Regret theory (Loomes and Sugden 1982) is also able to account for several patterns of choice under risk by assuming that states with higher payoff differences play a disproportionate role in choice. This intuition is closely related to ordering. To account for the Allais paradoxes, however, regret theory requires an additional convexity assumption which, as noted by the authors, lacks an independent psychological basis. More important, by focusing exclusively on payoff differences and neglecting diminishing sensitivity, regret theory can capture neither the dependence of risk attitudes on the payoff level x , nor framing effects and violations of procedural invariance, as previously noted.

In sum, we think that our model provides a parsimonious account of context dependent shifts in risk preference based on psychologically founded assumptions about the nature and the impact of the perceptual salience of lottery payoffs.

VII. CONCLUSION

Our article explores how cognitive limitations cause people to focus their attention on some but not all aspects of the world,

the phenomenon we call local thinking. We argue that salience, a concept well known to cognitive psychology, shapes this focus. In the case of choice under risk, this perspective can be implemented in a straightforward and parsimonious way by specifying that contrast between payoffs shapes their salience, and that people inflate the decision weights associated with salient payoffs. Basically, decision makers overweigh the upside of a risky choice when it is salient and thus behave in a risk-seeking way, and overweigh the downside when it is salient, and behave in a risk-averse way. This approach provides an intuitive and unified explanation of the instability of risk preferences, including the dramatic switches from risk-seeking to risk-averse behavior resulting from seemingly innocuous changes in the problem, as well as of some fundamental puzzles in choice under risk, such as the Allais paradox and preference reversals. It makes predictions for when these paradoxes will and will not occur, which we test and confirm experimentally.

Other aspects of salience have been used by economists to examine the consequences of people reacting to some pieces of data (salient ones) more strongly than to others. For example, Chetty, Looney, and Kroft (2009) show that shoppers are more responsive to sales taxes already included in posted prices than to sales taxes added at the register. Barber and Odean (2008) find that stock traders respond to “attention-grabbing” news. Perhaps most profoundly, Schelling (1960) has shown that people can solve coordination problems by focusing on salient equilibria based on their general knowledge, without any possibility for communication. Memory becomes a potential source of salient data. Our formal approach is consistent with this work, and stresses that in the specific context of choice under risk the relative magnitude of payoffs is itself a critical determinant of salience.

Our specification of contrast as a driver of salience could be useful for thinking about a variety of economic situations. For example, salience may affect consumer behavior: when considering which of different brands to buy, a consumer might focus on the attributes where the potential brands are most different, neglecting the others (Tversky and Simonson 1993; Bordalo 2011; Koszegi and Szeidl 2012). We use a version of this article’s model of salience to investigate consumer choices more broadly, as well as the endowment effect (Bordalo, Gennaioli, and Shleifer 2012a, 2012b). In many applications, the key idea of our approach is that

mental frames, rather than being fixed in the mind of the consumer, investor, or voter, are endogenous to the contrasting features of the alternatives of choice. This notion could perhaps provide a way to study how context shapes preferences in many social domains.

SUPPLEMENTARY MATERIAL

An Online Appendix for this article can be found at QJE online (qje.oxfordjournals.org)

APPENDIX

Proofs

PROPOSITION 1. If the probability of state s is increased by $d\pi_s = h\pi_s$ and the probabilities of other states are reduced while keeping their odds constant, that is, $d\pi_{\tilde{s}} = -\frac{\pi_s}{1-\pi_s}h\pi_{\tilde{s}}$ for all $\tilde{s} \neq s$, then for every lottery L_i :

$$\frac{d\omega_s^i}{h} = -\frac{\pi_s}{1-\pi_s} \cdot \omega_s^i \cdot (\omega_s^i - 1).$$

Proof. By definition,

$$\omega_s^i = \frac{\delta^{k_s^i-1}}{\sum_r \delta^{k_r^i-1} \cdot \pi_r}.$$

Therefore,

$$d\omega_s^i = -\frac{\omega_s^i}{\sum_r \delta^{k_r^i-1} \cdot \pi_r} \sum_r \delta^{k_r^i-1} \cdot d\pi_r.$$

Replacing $d\pi_s = h\pi_s$ and $d\pi_r = -\frac{\pi_s}{1-\pi_s}h\pi_r$ (for $r \neq s$) leads to

$$d\omega_s^i = -\frac{\omega_s^i}{\sum_r \delta^{k_r^i-1} \cdot \pi_r} \left[-\frac{h\pi_s}{1-\pi_s} \sum_{r \neq s} \delta^{k_r^i-1} \cdot \pi_r + h\delta^{k_s^i-1}\pi_s \right].$$

Thus

$$\frac{d\omega_s^i}{h} = -\omega_s^i \frac{1}{\sum_r \delta^{k_r^i-1} \cdot \pi_r} \left[-\frac{\pi_s}{1-\pi_s} \sum_{r \neq s} \delta^{k_r^i-1} \cdot \pi_r + \delta^{k_s^i-1}\pi_s \right].$$

The parentheses on the right-hand side can be rearranged to yield

$$\frac{\pi_s}{1 - \pi_s} \left[\delta^{k_s^i - 1} (1 - \pi_s) - \sum_{r \neq s} \delta^{k_r^i - 1} \cdot \pi_r \right] = \frac{\pi_s}{1 - \pi_s} \left[\delta^{k_s^i - 1} - \sum_r \delta^{k_r^i - 1} \cdot \pi_r \right],$$

where the sum is now over all states r . Inserting this term back into the equation above we get the result:

$$\frac{d\omega_s^i}{h} = -\omega_s^i \frac{\pi_s}{1 - \pi_s} (\omega_s^i - 1).$$

■

Proposition 1 has the following corollary: let lottery L yield payoff x_m with total probability p_m , with $\sum_n p_n = 1$. Let \mathbf{m} be the set of states where L pays x_m , and denote elements of \mathbf{m} by $s_{\mathbf{m}j}$, where $j = 1, \dots, |\mathbf{m}|$. Then $p_m = \sum_{s_{\mathbf{m}j} \in \mathbf{m}} \pi_{s_{\mathbf{m}j}} = p_m \sum_{s_{\mathbf{m}j} \in \mathbf{m}} \tilde{\pi}_{s_{\mathbf{m}j}}$, where we write $\tilde{\pi}_{s_{\mathbf{m}j}} = \pi_{s_{\mathbf{m}j}}/p_m$ (if L is being compared to another lottery L' and both lotteries are independent, then $\tilde{\pi}_{s_{\mathbf{m}j}}$ is just the probability that L' gives payoff y_j where the set \mathbf{m} is equal to the state (x_m, y_j)). Denote the salience distortion of p_m by

$$\omega_m = \frac{\sum_{s_{\mathbf{m}j} \in \mathbf{m}} \tilde{\pi}_{s_{\mathbf{m}j}} \delta^{k_{s_{\mathbf{m}j}}} - 1}{\sum_n \sum_{s_{\mathbf{n}j} \in \mathbf{n}} \tilde{\pi}_{s_{\mathbf{n}j}} \delta^{k_{s_{\mathbf{n}j}}} - 1} \cdot p_n}.$$

COROLLARY 1. If the probability p_m of payoff x_m is increased by $dp_m = hp_m$ and the probabilities of other states are reduced while keeping their odds constant, that is, $dp_{\tilde{m}} = -\frac{p_m}{1-p_m} hp_{\tilde{m}}$ for all $\tilde{m} \neq m$, then:

$$\frac{d\omega_m}{h} = -\frac{p_m}{1 - p_m} \cdot \omega_m \cdot (\omega_m - 1).$$

The proof of Corollary 1 is parallel to that of Proposition 1. In Lemma 1, $L_0 = (x, 1)$ and $L_1 = (x + g, \pi; x - l, 1 - \pi)$.

LEMMA 1. If the salience function is convex, then $r = v^{LT}(L_0) - v^{LT}(L_1)$ weakly decreases in x . Conversely, if the salience function is concave then r weakly increases in x .

Proof. First note that, due to linear utility, the premium r is independent of x for a given salience ranking. In fact, for any salience k_g^1, k_l^1 of L_1 's payoffs $x + g, x - l$ we have

$$\begin{aligned} v^{LT}(L_0) - v^{LT}(L_1) &= x - \frac{1}{\pi\delta^{k_g^1} + (1-\pi)\delta^{k_l^1}} \left[\pi\delta^{k_g^1}(x+g) + (1-\pi)\delta^{k_l^1}(x-l) \right] \\ &= \frac{1}{\pi\delta^{k_g^1} + (1-\pi)\delta^{k_l^1}} \left[l(1-\pi)\delta^{k_l^1} - g\pi\delta^{k_g^1} \right]. \end{aligned}$$

Second, note that if $\sigma(x+g, x) - \sigma(x, x-l)$ increases in x , then the upside of the risky lottery L_1 becomes weakly more salient as x increases. In particular, if L_1 's upside goes from being nonsalient (for low x) to being salient (for high x), the shift in r is negative:

$$\frac{l(1-\pi) - g\pi\delta}{\pi\delta + (1-\pi)} - \frac{l(1-\pi)\delta - g\pi}{\pi + (1-\pi)\delta} \propto -\pi(1-\pi)(1-\delta^2)[g+l].$$

which proves the claim when the salience function is convex (note that this shift goes to 0 as δ approaches 1). The concave case is analogous. ■

ROYAL HOLLOWAY, UNIVERSITY OF LONDON
 CREI (CENTRE DE RECERCA EN ECONOMIA INTERNACIONAL),
 UNIVERSITAT POMPEU FABRA AND BARCELONA GRADUATE
 SCHOOL OF ECONOMICS
 HARVARD UNIVERSITY

REFERENCES

- Aizpurua, Jose, Jorge Nieto, and Jose Uriarte, "Choice Procedure Consistent with Similarity Relations," *Theory and Decision*, 29 (1990), 235–254.
- Allais, Maurice, "Le Comportement de l'Homme Rationnel devant le Risque: Critique des Postulats et Axiomes de l'Ecole Americaine," *Econometrica*, 21 (1953), 503–546.
- Barber, Brad, and Terrence Odean, "All that Glitters: The Effect of Attention and News on the Buying Behavior of Individual and Institutional Investors," *Review of Financial Studies*, 21 (2008), 785–818.
- Bell, David, "Regret in Decision Making under Uncertainty," *Operations Research*, 30 (1982), 961–981.
- Birnbaum, Michael, and Ulrich Schmidt, "Allais Paradoxes Can Be Reversed by Presenting Choices in Canonical Split Form" (Mimeo, Kiel University, 2010).
- Bordalo, Pedro, "Choice Set Effects and Salience Theory" (Mimeo, Royal Holloway, University of London, 2011).
- Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer, "Salience and Consumer Choice" (NBER Working Paper 17947, 2012a).
- , "Salience in Experimental Tests of the Endowment Effect," *American Economic Review*, 102 (2012b), 47–52.

- Bostic, Raphael, Richard J. Herrnstein, and R. Duncan Luce, "The Effect on the Preference-Reversal Phenomenon of Using Choice Indifferences," *Journal of Economic Behavior & Organization*, 13 (1990), 193–212.
- Camerer, Colin "Choice under Risk and Uncertainty," in *Handbook of Experimental Economics*, ed. Kagel, John, and Roth, Alvin (Princeton, N.J.: Princeton University Press, 1995).
- Chetty, Raj, Adam Looney, and Kory Kroft, "Salience and Taxation: Theory and Evidence," *American Economic Review*, 99 (2009), 1145–1177.
- Conlisk, John, "Three Variants on the Allais Example," *American Economic Review*, 79 (1989), 392–407.
- Cox, James C., and Seth Epstein, "Preference Reversals without the Independence Axiom," *American Economic Review*, 79 (1989), 408–426.
- Edwards, Ward, "Subjective Probabilities Inferred from Decisions," *Psychological Review*, 69 (1962), 109–135.
- Fellner, William, "Distortion of Subjective Probabilities as a Reaction to Uncertainty," *Quarterly Journal of Economics*, 75 (1961), 670–689.
- Fishburn, Peter, "Nontransitive Measurable Utility," *Journal Mathematical Psychology*, 26 (1982), 31–67.
- Friedman, Milton, and Leonard J. Savage, "The Utility Analysis of Choices involving Risk," *Journal of Political Economy*, 56 (1948), 279–304.
- Gabaix, Xavier, "A Sparsity-Based Model of Bounded Rationality" (NBER working paper No. 16911, 2011).
- Gennaioli, Nicola, and Andrei Shleifer, "What Comes to Mind," *Quarterly Journal of Economics*, 125 (2010), 1399–1433.
- Grether, David, and Charles Plott, "Economic Theory of Choice and the Preference Reversal Phenomenon," *American Economic Review*, 69 (1979), 623–638.
- Harrison, Glenn, "Expected Utility Theory and the Experiments," *Empirical Economics*, 19 (1994), 223–253.
- Holt, Charles, "Preference Reversals and the Independence Axiom," *American Economic Review*, 76 (1986), 508–515.
- Kahneman, Daniel, "Maps of Bounded Rationality: Psychology for Behavioral Economics," *American Economic Review*, 93 (2003), 1449–1476.
- , *Thinking, Fast and Slow*. (New York: Farrar, Strauss, Giroux, 2011).
- Kahneman, Daniel, and Amos Tversky, "Prospect Theory: an Analysis of Decision under Risk," *Econometrica*, 47 (1979), 263–292.
- Kandel, Eric, James Schwartz, and Thomas Jessell, *Principles of Neural Science*. 4th ed. (Norwalk, CT: Appleton & Lange, 1991).
- Karni, Edi, and Zvi Safra, "'Preference Reversal' and the Observability of Preferences by Experimental Methods," *Econometrica*, 55 (1987), 675–685.
- Koszegi, Botond, and Adam Szeidl, "A Model of Focusing in Economic Choice" (Berkeley, Mimeo: University of California, 2012).
- Koszegi, Botond, and Matthew Rabin, "A Model of Reference-Dependent Preferences," *Quarterly Journal of Economics*, 121 (2006), 1133–1165.
- , "Reference-Dependent Risk Attitudes," *American Economic Review*, 97 (2007), 1047–1073.
- Laury, Susan, and Holt Charles, "Further Reflections on Prospect Theory" (Andrew Young School of Policy Studies Research Paper Series No. 06-11, 2005).
- Leland, Jonathan, "Generalized Similarity Judgments: An Alternative Explanation for Choice Anomalies," *Journal of Risk and Uncertainty*, 9 (1994), 151–172.
- Lichtenstein, Sarah, and Paul Slovic, "Reversals of Preference between Bids and Choices in Gambling Decisions," *Journal of Experimental Psychology*, 89 (1971), 46–55.
- Loomes, Graham, "Modeling Choice and Valuation in Decision Experiments," *Psychological Review*, 117 (2010), 902–924.
- Loomes, Graham, and Robert Sugden, "Regret Theory: An Alternative Theory of Rational Choice under Uncertainty," *Economic Journal*, 92 (1982), 805–824.
- , "A Rationale for Preference Reversal," *American Economic Review*, 73 (1983), 428–432.

- Manzini, Paola, and Marco Mariotti, "Sequentially Rationalizable Choice," *American Economic Review*, 97 (2007), 1824–1839.
- Masatlioglu, Yusufcan, Daisuke Nakajima, and Erkut Ozbay, "Revealed Attention," *American Economic Review* forthcoming (2012).
- McCoy, Allison, and Michael Platt, "Risk-Sensitive Neurons in Macaque Posterior Cingulate Cortex," *Nature Neuroscience*, 8 (2005), 1220–1227.
- Mullainathan, Sendhil, "Thinking through Categories" (Mimeo, Harvard University, 2002).
- Prelec, Drazen, "The Probability Weighting Function," *Econometrica*, 66 (1998), 497–527.
- Quiggin, John, "A Theory of Anticipated Utility," *Journal of Economic Behavior & Organization*, 3 (1982), 323–343.
- Rubinstein, Ariel, "Similarity and Decision-Making under Risk," *Journal of Economic Theory*, 46 (1988), 145–153.
- Savage, Leonard, *Foundations of Statistics*, 2 ed. (New York: Dover, 1972).
- Schelling, Thomas, *The Strategy of Conflict* (Cambridge, MA: Harvard University Press, 1960).
- Schwartzstein, Joshua, "Selective Attention and Learning" (Mimeo, Harvard University, 2009).
- Segal, Uzi, "Does the Preference Reversal Phenomenon Necessarily Contradict the Independence Axiom," *American Economic Review*, 78 (1988), 233–236.
- Taylor, Shelley, and Suzanne Thompson, "Stalking the Elusive Vividness Effect," *Psychological Review*, 89 (1982), 155–181.
- Tversky, Amos, and Daniel Kahneman, "The Framing of Decisions and the Psychology of Choice," *Science*, 211 (1981), 453–458.
- , "Advances in Prospect Theory: Cumulative Representation of Uncertainty," *Journal of Risk and Uncertainty*, 5 (1992), 297–323.
- Tversky, Amos, and Itamar Simonson, "Context-Dependent Preferences," *Management Science*, 39 (1993), 1179–1189.
- Tversky, Amos, Paul Slovic, and Daniel Kahneman, "The Causes of Preference Reversals," *American Economic Review*, 80 (1990), 204–217.
- Wu, George, and Richard Gonzalez, "Curvature of the Probability Weighting Function," *Management Science*, 42 (1996), 1676–1690

This page intentionally left blank