Online Appendix – Not for Publication

Part 1. Extensions

1.A Mixed Lotteries

In this section, we apply our model to mixed lotteries, those involving both positive and negative payoffs. To this end, we come back to the KT (1979) piecewise linear value function exhibiting loss aversion, for loss aversion provides an intuitive explanation for risk aversion with respect to small mixed bets. Using the salience function of Equation (5) in the text, for which $\sigma(x, y) = \sigma(-x, -y)$ for all x, y, all risk aversion for lotteries symmetric around zero is due to loss aversion. For non-symmetric lotteries, salience and loss aversion interact to determine risk preferences. To see this, consider Samuelson's wager, namely the choice between the lotteries:

$$L_S = \begin{cases} \$200, & 0.5 \\ -\$100, & 0.5 \end{cases}, \qquad L_0 = (\$0,\ 1)\ .$$

In this choice, many subjects decline L_S even though it has a positive and substantial expected value. With a symmetric salience function, we have that $\sigma(200, 0) > \sigma(100, 0) = \sigma(-100, 0)$, implying that in this choice the local thinker focuses on the lottery gain.

Consider now what happens under the following piecewise linear value function:

$$v(x) = \begin{cases} x, & \text{if } x > 0 \\ \lambda x, & \text{if } x < 0 \end{cases},$$

where $\lambda > 1$ captures loss aversion. Now the local thinker rejects L_S provided:

$$200 \cdot \frac{1}{1+\delta} - 100\lambda \cdot \frac{\delta}{1+\delta} < 0.$$

The decision maker rejects L_S when his dislike for losses more than compensates for his focus on the lottery gain, i.e. $\lambda > 2/\delta$.[1] In lotteries whose negative downside is larger in magnitude than the positive upside, salience and loss aversion go in the same direction in triggering risk aversion.

Although our approach can be easily integrated with standard loss aversion, we wish to stress that salience may itself provide one interpretation of the idea that "losses loom larger than gains" (KT 1979) where, independently of loss aversion in the value function, states with negative payoffs are ceteris paribus more salient than states with positive payoffs. The ranking of positive and negative states is in fact left unspecified by Definition 1. One could therefore add an additional property:

4) *Loss salience: for every state s with payoffs* $\mathbf{x}_s = (x_s^i)_{i=1,2}$ *such that* $x_s^1 + x_s^2 > 0$ *we* have that

$$\sigma(-x_s^1, -x_s^2) > \sigma(x_s^1, x_s^2).$$

This condition relaxes the symmetry around zero of the salience function of Equation (5), postulating that departures from zero are more salient in the negative than in the positive direction. In this specification, local thinking can itself be a force towards risk aversion for mixed lotteries, complementing loss aversion. In particular, if losses are sufficiently more salient than gains, one can account for Samuelson's wager based on salience alone (and linear utility): if $\sigma(-100, 0) > \sigma(200, 0)$, a local thinker with linear utility rejects Samuelson's bet as long as $200 \cdot \frac{\delta}{1+\delta} - 100 \cdot \frac{1}{1+\delta} < 0$, or $\delta < 1/2$. A specification where risk aversion for mixed lotteries arises via the salience of lottery payoffs may give distinctive implications from standard loss aversion, but we do not investigate this possibility here.

1.B Choice Among Many Lotteries

We now extend our model to a general choice among $N \geq 2$ of lotteries, which is particularly useful for economic applications. Before doing so, note that preferences over $N \geq 2$

[1] The role of loss aversion can also be gauged by considering the choice between two symmetric lotteries with zero expected value, $L_1 = (-x, 0.5;\ x, 0.5)$ and $L_2 = (-y, 0.5;\ y, 0.5)$, with $x > y$. Since (5) is symmetric, the states $(-x, y)$ and $(x, -y)$ have salience rank 1, whereas states $(-x, -y)$ and (x, y) have salience rank 2, so that L_1 is evaluated at $x(1-\lambda)/2$, and L_1 is evaluated at $y(1-\lambda)/2$. This implies that for any degree of loss aversion $\lambda > 1$, the Local Thinker prefers the safer lottery L_2.

lotteries cannot be inferred from pairwise comparisons because salience changes across comparisons and intransitivities can arise (see online Appendix 1.C).

To model choice from an arbitrary set of alternatives $\aleph = \{L_1, \ldots, L_N\}$ defined over a state space S (as in Section 3), we first generalize the notion of payoff salience. Let $x_s = (x_s^1, \ldots, x_s^N)$ be the vector of payoffs delivered in a generic state s, and denote by $x_s^{-i} = \{x_s^j\}_{i \neq j}$ the vector of payoffs excluding x_s^i. The salience of state s for lottery L_i is then captured by a function $\widehat{\sigma}(x_s^i, x_s^{-i})$ which contrasts L_i's payoff x_s^i in s with all other payoffs $x_{\mathbf{s}}^{-i}$ in the same state. Let $x_s^{-i} + \epsilon$ denote the vector with elements $\{x_s^j + \epsilon\}_{j \neq i}$. In line with Definition 1, we impose the following properties:

Definition 3 *Given a state space S and a choice set $\aleph$, the salience of state s for lottery L_i is given by a continuous and bounded function $\widehat{\sigma}(x_s^i, \mathbf{x}_s^{-i})$ that satisfies three conditions:*
1) *Ordering: if $x_s^i = \max \mathbf{x}_s$, then for any $\epsilon, \epsilon' \geq 0$ (with at least one strict inequality):*

$$\widehat{\sigma}\left(x_s^i + \epsilon, \mathbf{x}_s^{-i} - \epsilon'\right) > \widehat{\sigma}\left(x_s^i, \mathbf{x}_s^{-i}\right).$$

If $x_s^i = \min \mathbf{x}_s$, then for any $\epsilon, \epsilon' \geq 0$ (with at least one strict inequality):

$$\widehat{\sigma}\left(x_s^i - \epsilon, \mathbf{x}_s^{-i} + \epsilon'\right) > \widehat{\sigma}\left(x_s^i, \mathbf{x}_s^{-i}\right).$$

2) *Diminishing sensitivity: if $x_s^j > 0$ for all j, then for any $\epsilon > 0$,*

$$\widehat{\sigma}(x_s^i + \epsilon, \mathbf{x}_s^{-i} + \epsilon) < \widehat{\sigma}(x_s^i, \mathbf{x}_s^{-i})$$

3) *Reflection: for any two states $s, \widetilde{s} \in S$ such that $x_s^j, x_{\widetilde{s}}^j > 0$ for all j, we have*

$$\widehat{\sigma}(x_s^i, \mathbf{x}_s^{-i}) < \widehat{\sigma}(x_{\widetilde{s}}^i, \mathbf{x}_{\widetilde{s}}^{-i}) \quad \textit{if and only if} \quad \widehat{\sigma}(-x_s^i, -\mathbf{x}_s^{-i}) < \widehat{\sigma}(-x_{\widetilde{s}}^i, -\mathbf{x}_{\widetilde{s}}^{-i})$$

When $N > 2$, one can construct a salience function satisfying the above requirements by setting:

$$\widehat{\sigma}(x_s^i, \mathbf{x_s^{-i}}) \equiv \sigma(x_s^i, \mathbf{f}(\mathbf{x_s^{-i}})), \tag{1}$$

where $\sigma(.,.)$ is the salience function employed in the two lottery case of Section 3, and $f(x_s^{-i}) : R^{N-1} \to R$ is a function of the residual vector x_s^{-i}. Definitons 1 and 3 jointly impose some restrictions on the properties of $f(x_s^{-i})$. One intuitive specification which, together with Definition 1, satisfies Definition 3 is:

$$\mathbf{f}(\mathbf{x}_s^{-i}) = \frac{1}{N-1} \sum_{j \neq i} x_s^j. \tag{2}$$

That is, the salience of a lottery state s depends on the contrast between the lottery's payoff and the average payoff of the other lotteries in s. For $N = 2$, Equation (1) reduces to the salience function of Section 3.

Even though $\sigma(.,.)$ is symmetric, when there are more than two lotteries salience is in general not symmetric because the same state does not necessarily have the same salience for different lotteries. Consider for instance a state where lottery L_i's payoff x_s^i is very different from the payoffs of all the other lotteries in x_s^{-i}, but in turn the payoffs in x_s^{-i} are similar to each other. According to the salience function implied by (1) and (2), then, state s is very salient for L_i but not salient for the other lotteries. In contrast, a state s may be very salient for all lotteries if in that state half the lotteries have a very low payoff and the other half have a very high payoff.

Given a lottery specific salience ranking k_s^i based on the salience function $\widehat{\sigma}$, each state is assigned a decision weight π_s^i according to Equation (8) in the text, and a value $V^{LT}(L_i)$ is computed for each lottery L_i according to Equation (10).

One important new effect arising in the choice among $N > 2$ lotteries is that the preference ranking among any two lotteries depends on the remaining alternatives, potentially leading to violations of independence of irrelevant alternatives (IIA). By shaping payoff salience, the choice set is a source of context effects. A detailed analysis of these possibilities can be found in Bordalo (2011) and Bordalo, Gennaioli and Shleifer (2011).

1.C Intransitivity of pairwise preferences

Systematic intransitivities in choice under risk have been documented by several authors (Tversky 1969, Starmer and Sugden 1996, Luce 2000). Some models accounting for this phenomenon are Tversky's additive difference model and Regret Theory. Our model also yields intransitivities, and the structure imposed by salience allows to place restrictions on the circumstances in which intransitivities can or cannot occur.

To illustrate how intransitive preferences may arise in our model, consider the following three lotteries:

$$L_\pi = \begin{cases} \alpha x, & \pi \\ 0, & 1-\pi \end{cases}, \quad L_\$ = \begin{cases} x, & \alpha\pi \\ 0, & 1-\alpha\pi \end{cases}, \quad L_s = (y,1), \tag{3}$$

where $x, y > 0$ and $\alpha < 1$. Lotteries $L_\$$ and L_π are of the kind giving rise to the preference reversals of Section 5. In this case, a local thinker prefers the safer lottery L_π to $L_\$$ as long as π is large and δ is not too small. Suppose now that the sure prospect y is such that in the pairwise comparison with $L_\$$ the latter's gain is salient while in that with L_π the latter's loss is salient, i.e. $\sigma(x,y) > \sigma(0,y) > \sigma(\alpha x, y)$. It is then possible to find values (y, δ) such that choices are intransitive:[2]

$$L_\pi \succ L_\$, \quad L_\$ \succ L_s, \quad L_s \succ L_\pi.$$

Intransitivity arises because risk aversion in the direct comparison of L_π with $L_\$$ is reversed to risk seeking when the two lotteries are *indirectly* compared via their pairwise choice against the sure thing L_s. The intuition is as follows. In the direct comparison, $L_\pi \succ L_\$$ because lottery L_π pays off with much higher probability than $L_\$$. In the *indirect* comparison, $L_\pi \prec L_\$$ because the sure thing stresses the upside of the risky lottery and the downside of the safe lottery. This is "as if" in the direct comparison the decision maker chooses based on probabilities, while in the indirect comparison he chooses based on payoffs. This intuition is closely related to Tversky's (1969) account of intransitivities.

Although in our model intransitivities can arise in pairwise choices, choice is well defined

[2] One numerical example is $x = 100, \alpha = 1/10, \pi = 3/4, y = 4$ and $\delta = 0.75$.

within any given choice set. The choice set shapes the salience of each option and thus its valuation according to Definition 3. In particular, the local thinker's choice from $\{L_\pi, L_\$, L_s\}$ is well defined and can be found by applying Definition 3.[3]

1.C.1 Limits on Intransitivity

We now show that intransitivities never occur in choices among independent lotteries sharing the same support with two or three outcomes. The intuition is that for such choices, the state space – and the salience ranking of states – does not change across choices, only the probabilities of states do. The fact that a local thinker is transitive in such choices is consistent with the intuition that intransitivities require shifts in attention and salience from one choice to the next (Tversky, 1969).

Consider the space of payoffs $\{x, y, z\}$, where $x > y > z > 0$. A lottery is defined by the probability it yields each payoff, which we denote $L_a = \{a_x, a_y, a_z\}$. We now show:

Lemma 2 *A local thinker's preferences are pairwise transitive if $\delta = 0$ or $\delta = 1$. Furthermore, for any δ preferences are pairwise transitive within the set of independent lotteries sharing the same support consisting of up to 3 payoffs.*

Proof. Assuming all lotteries are independent (and all probabilities are positive), in a pairwise choice there are 9 states of the world. The salience ranking is univocal, up to the relative ranking of states (x, y) and (y, z). For now take the following ranking

$$\sigma(x, z) > \sigma(y, z) > \sigma(x, y) > \sigma(x, x) = \sigma(y, y) = \sigma(z, z).$$

Then $L_a \succ L_b$ iff

$$(x - z)[a_x b_z - a_z b_x] + \delta(y - z)[a_y b_z - a_z b_y] + \delta^2(x - y)[a_x b_y - a_y b_x] > 0 \tag{4}$$

We want to show that choice from such lotteries is transitive, i.e. $L_a \succ L_b$ and $L_b \succ L_c$ imply $L_a \succ L_c$.

[3] In particular, any pairwise intransitive choice pattern gives rise to a violation of independence of irrelevant alternatives when compared to choice from the full choice set.

Rewrite equation (4) in a more convenient way, by replacing $a_3 = 1 - a_1 - a_2$ and similarly for b_3. Then $L_a \succ L_b$ provided

$$(x - z)[a_x - b_x] + \delta(y - z)[a_y - b_y] > \Delta[a_x b_y - a_y b_x] \tag{5}$$

where $\Delta = (x - z) - \delta(y - z) - \delta^2(x - y)$. The left hand side captures the intuition that L_a is preferred to L_b when it has a sufficiently higher probability of a high payoff. In the limit $\delta \to 1$, Δ goes to zero and preferences among lotteries depend only on the linear terms in the left hand side

$$(x - z)a_x + (y - z)a_y > (x - z)b_x + (y - z)b_y \tag{6}$$

Since these conditions are linear, lotteries can be ranked by this weighted average of payoff differences and so preferences are transitive (and of course they coincide with the rational agent's preferences).

Intransitivities may arise only because of the quadratic term on the right hand side of (4), which comes from the fact that we are dealing with the product state space. The effect of local thinking on the preferences between lotteries cannot be easily characterized. To see that, note that Δ measures the average (perceived) difference between any two payoffs x, y and z. While for $\delta \to 1$ these differences cancel out, for $\delta < 1$ the largest payoff difference $(x - z)$ is overvalued, so in general Δ is positive and increasing as δ decreases. Then if e.g. $a_x > b_x$ and $a_y < b_y$ local thinking increases both sides of (5). Still, it is instructive to compare the fully rational case (6) to the case $\delta = 0$. Here $L_a \succ L_b$ when:

$$\frac{a_x}{a_z} > \frac{b_x}{b_z} \tag{7}$$

For the full local thinker, only the ratio of probabilities for the most salient payoffs matters. As a result, lotteries are ranked according to this ratio and we recover transitivity.

To examine the general case, suppose $L_a \succ L_b$ and $L_b \succ L_c$ for a given $\delta \in (0, 1)$. Then

(5) implies:

$$\begin{aligned}(x-z)a_x+\delta(y-z)a_y &> (x-z)b_x+\delta(y-z)b_y+\Delta[a_xb_y-a_yb_x]\\ &> (x-z)c_x+\delta(y-z)c_y+\Delta[a_xb_y-a_yb_x]+\Delta[b_xc_y-b_yc_x]\end{aligned}$$

Note that in deriving the second line, it is essential that the salience ranking does not change when we compare different lotteries. Here the fact that all lotteries have the same support plays a crucial role. Rearranging, we get that $L_a \succ L_b$ and $L_b \succ L_c$ together imply

$$(x-z)[a_x-c_x]+\delta(y-z)[a_y-c_y] > \Delta[a_xb_y-a_yb_x]+\Delta[b_xc_y-b_yc_x] \tag{8}$$

while $L_a \succ L_c$ means (note the same left hand side)

$$(x-z)[a_x-c_x]+\delta(y-z)[a_y-c_y] > \Delta[a_xc_y-a_yc_x]$$

Subtracting $\Delta[a_xc_y-a_yc_x]$ from both sides of (8) we get

$$\begin{aligned}(x-z)[a_x-c_x]+\delta(y-z)[a_y-c_y] &- \Delta[a_xc_y-a_yc_x]\\ &> \Delta[(a_xb_y-a_yb_x)+(b_xc_y-b_yc_x)-(a_xc_y-a_yc_x)]\\ &= \Delta[a_x(b_y-c_y)+c_x(a_y-b_y)+b_x(c_y-a_y)] \qquad (9)\end{aligned}$$

The left hand side of this equation (first line) is positive iff $L_a \succ L_c$. For simplicity write it as $\tilde{u}(L_a-L_b)$. We now show that the right hand side (third line) of this equation can also be written in terms of $\tilde{u}(L_a-L_b)$. We use again that $L_a \succ L_b$ and $L_b \succ L_c$, which put lower bounds on (a_y-b_y) and (b_y-c_y) respectively (according to (5)) to get:

$$\begin{aligned}a_x(b_y-c_y) &+ c_x(a_y-b_y)+b_x(c_y-a_y)\\ &> \frac{\Delta}{\delta(y-z)}b_x[a_xc_y-a_yc_x]-\frac{x-z}{\delta(y-z)}b_x[a_x-c_x]+b_x[c_y-a_y]\\ &= -\frac{b_x}{\delta(y-z)}\left[(x-z)[a_x-c_x]+\delta(y-z)[a_y-c_y]-\Delta[a_xc_y-a_yc_x]\right]\\ &= -\frac{b_x}{\delta(y-z)}\tilde{u}(L_a-L_c)\end{aligned}$$

Inserting this result in (9) we get

$$\tilde{u}(L_a - L_c) > -\frac{b_x \Delta}{\delta(y-z)} \tilde{u}(L_a - L_c)$$

which, because $y > z$ and $\Delta > 0$, implies that $\tilde{u}(L_a - L_c) > 0$, namely

$$L_a \succ L_c.$$

So we have transitivity. Moreover, we also have that $L_a \sim L_b$ and $L_b \sim L_c$ imply $L_a \sim L_c$.

■

1.D Preference Reversals under the Revealed Preference Approach

Section 5 in the text showed that under the "valuation approach" our model explains when preference reversals occur in choice and pricing among lotteries of the form

$$L_1 = (x, p;\ 0, 1-p) \qquad L_2 = (2x, p/2;\ 0, 1-p/2).$$

We now show that the same patterns can arise when selling prices are computed under the revealed preference approach. To do so, recall that a local thinker with linear utility chooses the safer lottery L_1 in the choice set $\{L_1, L_2\}$ if and only if:

$$p \geq \frac{2(1-\delta)}{2-\delta-\delta^2}. \tag{10}$$

Therefore, it is sufficient to show that the local thinker may state a higher "revealed preference price " for the riskier lottery L_2 than for the safer lottery L_1 when (10) holds.

1.D.1 Reversal and Pricing in Isolation

To define the "revealed preference price" for a lottery $L = (y, q; 0, 1-q)$, consider the choice set $\{L, L_P\}$, where $L_P = (P, 1)$ is a lottery promising the sure amount P. Define the revealed preference price of L as the minimum P such that a local thinker weakly prefers the sure payoff L_P to L. When $y = x$ and $q = p$ we have that $L = L_1$ and P is the price of L_1. When instead $y = 2x$ and $q = p/2$ we have that $L = L_2$ and P is the price of L_2.

With linear utility and a given salience ranking, the certainty equivalent of a lottery is its expected value computed using the decision weights implied by salience. If the upside of lottery L is salient, then the decision weight attached to the lottery's upside is $\frac{q}{q+(1-q)\delta}$, and so $P = y \cdot \frac{q}{q+(1-q)\delta}$. If instead the downside of lottery L is salient, the decision weight attached to the lottery's upside is $\frac{q\delta}{q\delta+(1-q)}$, and so $P = y \cdot \frac{q\delta}{q\delta+(1-q)}$.

Since lotteries L_1 and L_2 have the same mean, the certainty equivalent of the risky lottery L_2 is higher than that of the safe lottery L_1 if L_2 has a salient upside (even if L_1's upside is also salient). This is because L_2 pays its salient upside with a smaller probability, and we know that smaller probabilities are subject to greatest distortions. By the same argument, the certainty equivalent of the risky lottery L_2 is lower than that of the safe lottery L_1 if the downside of L_2 is salient. As a consequence, the necessary and sufficient condition for the certainty equivalent of the risky lottery L_2 to be higher than that of the safe lottery L_1 (and thus for preference reversals to arise when (10) holds), is that the upside of L_2 is salient.

The upside of the general lottery L is salient when the lottery is contrasted with its price $P = y \cdot \frac{q}{q+(1-q)\delta}$ provided:

$$\sigma\left(y, y \cdot \frac{q}{q+(1-q)\delta}\right) > \sigma\left(0, y \cdot \frac{q}{q+(1-q)\delta}\right). \tag{11}$$

Condition (11) imposes that, at the high revealed preference price, the lottery's upside is indeed salient.[4]

For preference reversals to occur in the context of lotteries L_1 and L_2, it must then be the case that the above condition is satisfied for L_2, namely when $y = 2x$ and $q = p/2$. That

[4]To fully characterize the revealed preference price P for L, note that the latter has a salient downside, and so $P = y\frac{\delta q}{q\delta+(1-q)}$ provided:

$$\sigma\left(y, y\frac{\delta q}{q\delta+(1-q)}\right) < \sigma\left(0, y\frac{\delta q}{q\delta+(1-q)}\right). \tag{12}$$

Conditions (11) and (12) are mutually exclusive, but for a given salience function there may be parameter values (y, q, δ) for which neither of them holds. This phenomenon is caused by jumps in the salience ranking, and disappears when probability distortions are a smooth function of salience. In this case, the minimum price P which is (weakly) preferred to the lottery is determined by $\sigma(y, P) = \sigma(0, P)$ (note that the certainty equivalent is not continuous, but is monotonic).

is, it must be that:

$$\sigma\left(2x, xp\frac{2}{p+(2-p)\delta}\right) > \sigma\left(0, xp\frac{2}{p+(2-p)\delta}\right). \tag{13}$$

The mechanism for preference reversals encoded in (13) is the same as the one creating these reversals in the valuation approach of Section 5. When considered in isolation, the upside of the riskier lottery becomes salient, and so the lottery is overvalued relative to L_1. The only difference now is that condition (13) imposes an additional restriction on the expected value xp, namely it cannot be too high.

We can also use (13) to characterize a sufficient condition for preference reversals to arise. Since the lottery's price P increases in the probability p, condition (13) is less likely to hold as p increases. As a result, preference reversals can only occur when (13) is satisfied at the lowest possible probability $p = \frac{2(1-\delta)}{2-\delta-\delta^2}$ consistent with (10), namely when:

$$\sigma\left(2x, x\frac{2}{1+\delta+\delta^2}\right) > \sigma\left(0, x\frac{2}{1+\delta+\delta^2}\right). \tag{14}$$

It is only when (14) holds that it is possible to find $p \geq \frac{2(1-\delta)}{2-\delta-\delta^2}$ such that preference reversals occur. Note that the condition is more likely to hold when δ is high. Using the salience function (5) in the text, the above condition becomes:

$$\frac{2x}{1+\delta+\delta^2} < \theta\left(\frac{1+\delta+\delta^2}{2} - 1\right).$$

The lotteries' payoff cannot be too large.

1.D.2 Preference Reversals in the Context of Choice

Consider now the case in which the certainty equivalents for the two lotteries are determined jointly, namely in the choice context. This is akin to presenting the local thinker with the choice set $\{L_1, L_2, L_{P_1}, L_{P_2}\}$, where $L_1 = (x, p; 0, 1-p)$, $L_2 = (2x, p/2; 0, 1-p/2)$ and $L_{P_i} = (P_i, 1)$, where P_i is the revealed preference price of lottery L_i. Now explicitly determining the prices is much more complicated because one needs to jointly determine two prices and two salience rankings. The key point, though, is that in equilbrium, the price of each lottery will be equal to the lottery's expected value calculated at the equilibrium salience ranking. This implies that when choosing among the lotteries the local thinker will value them at the resulting expected values. Accordingly, when pricing the lotteries the local thinker will state precisely the lotteries' perceived expected values. As a result, the local thinker's valuation of lotteries will be consistent with their pricing, and no preference reversals will occur.

1.E First Order Stochastic Dominance for Independent Lotteries

Let $\sigma : \mathbb{R} \times \mathbb{R} \to \mathbb{R}$ be a function satisfying ordering

$$[a, b] \subset [a', b'] \Rightarrow \sigma(a, b) < \sigma(a', b')$$

If a local thinker has a choice between lottery $L_1 = (x_i, p_i)$ (yields payoff x_i with probability p_i) and lottery $L_2 = (y_j, q_j)$, denote by L_i^{LT} the lotteries' representations derived from the salience function σ, i.e.

$$L_1 = (x_i, p_i^{LT}), \qquad L_2 = (y_i, q_i^{LT})$$

where decision weights are determined as in Definition 2. Set $i = 1, \ldots, N$ and $j = 1, \ldots, M$ such that $x_{i+1} > x_i$ and $y_{i+1} > y_i$.

Lemma 3 *Let $L_1 = (x_i, p_i)$ and $L2 = (y_j, q_j)$ be two independent lotteries over positive payoffs, whose supports intersect at most in one point. If lottery L_1 first order stochastically dominates (f.o.s.d.) lottery L_2 , then $L_1^{LT,\delta}$ f.o.s.d. $L_2^{LT,\delta}$ for any local thinker (meaning, a local thinker defined by any salience function satisfying ordering).*

Proof. L_1 f.o.s.d. L_2 if and only if

$$\sum_{j \geq k} q_j - \sum_{i \geq i_k} p_i \leq 0$$

for each $k = 1, ..., M$, where $i_k = \min\{i : x_i \geq y_k\}$. It is convenient to rewrite these stochastic dominance (SD) conditions in terms of sums over the state space

$$\sum_{j\geq k}\sum_{\tilde{i}} q_j p_{\tilde{i}} - \sum_{i\geq i_k}\sum_{\tilde{j}} p_i q_{\tilde{j}} \leq 0 \tag{15}$$

For a local thinker this becomes

$$\sum_{j\geq k} q_j \sum_{\tilde{i}} p_{\tilde{i}}\delta_{\tilde{i}j} \leq \sum_{i\geq i_k} p_i \sum_{\tilde{j}} q_{\tilde{j}}\delta_{i\tilde{j}} \tag{16}$$

But notice that the sums have states in common, which drop out of the expression. These are the states (i, j) with $i \geq i_k,\ j \geq k$. Thus the expression reduces to:

$$\sum_{j\geq k} q_j \sum_{\tilde{i}<i_k} p_{\tilde{i}}\delta_{\tilde{i}j} \leq \sum_{i\geq i_k} p_i \sum_{\tilde{j}<k} q_{\tilde{j}}\delta_{i\tilde{j}} \tag{17}$$

This is a friendlier expression than the full sum (16) because SD ensures that

$$\sum_{j\geq k} q_j \leq \sum_{i\geq i_k} p_i, \quad \text{and equivalently} \quad \sum_{\tilde{i}\leq i_k} p_{\tilde{i}} \leq \sum_{\tilde{j}<k} q_{\tilde{j}}.$$

We prove (17) by induction over the index k of outcomes of L_2. Starting from $k = M$, (17) becomes:

$$q_M \sum_{\tilde{i}<i_M} p_{\tilde{i}}\delta_{\tilde{i}M} \leq \sum_{i\geq i_M} p_i \sum_{\tilde{j}<M} q_{\tilde{j}}\delta_{i\tilde{j}} \tag{18}$$

Recall that S.D. implies

$$q_M \leq \sum_{i\geq i_M} p_i \tag{19}$$

so it is sufficient to show that the factor multiplying q_M has the upper bound

$$\sum_{\tilde{i}<i_M} p_{\tilde{i}}\delta_{\tilde{i}M} \leq \sum_{\tilde{j}<M} q_{\tilde{j}}\delta_{i^*\tilde{j}} \tag{20}$$

where i^* minimizes the coefficient of the right hand side of (18), that is $\sum_{\tilde{j}<k} q_{\tilde{j}}\delta_{i\tilde{j}}$. We can rewrite the left hand side of (20) as a sum over $\tilde{j}$, by grouping all the x_i payoffs that occur

between $y_{\tilde{j}}$ and $y_{\tilde{j}+1}$:

$$\sum_{\tilde{j}<i_M}\left(\sum_{\tilde{i}=i_{\tilde{j}}}^{i_{\tilde{j}}+1} p_{\tilde{i}}\delta_{\tilde{i}M}\right) \le \sum_{\tilde{j}<M} q_{\tilde{j}}\delta_{i^*\tilde{j}} \tag{21}$$

By SD, this holds when $\delta = 1$. To investigate the impact of the δs, recall that $i^* > M$ and $\tilde{j} < M$. Thus salience ordering guarantees that within any term $\tilde{j}$ in the sum, we have

$$\tilde{i} \in \left[i_{\tilde{j}}, i_{\tilde{j}+1}\right] \Rightarrow \delta_{\tilde{i}M} \le \delta_{i^*\tilde{j}}$$

since in fact $\left[x_{\tilde{i},y_M}\right] \subset \left[y_{\tilde{j},x_{i^*}}\right]$. Therefore, within each interval $\left[y_{\tilde{j}}, y_{\tilde{j}+1}\right]$, the left hand side is discounted more than the right hand side. Together with the stochastic dominance condition (19), this means that (18) holds. Moreover, this implies that

$$q_M^{LT,\delta} - \sum_{i\ge i_M} p_i^{LT,\delta} \le q_M - \sum_{i\ge i_M} p_i \le 0 \tag{22}$$

so that the difference between the cumulative probability distribution of L_1 and L_2 is amplified by the local thinker. Now suppose the following holds for some $k+1 < M$.

$$\sum_{j\ge k+1} q_j^{LT,\delta} - \sum_{i\ge i_k+1} p_i^{LT,\delta} \le \sum_{j\ge k+1} q_j - \sum_{i\ge i_k+1} p_i \tag{23}$$

Then to show it holds for k, we only need to show that the k-terms that are added to the sum above satisfy

$$q_k^{LT,\delta} - \sum_{i=i_k}^{i_{k+1}-1} p_i^{LT,\delta} \le q_k - \sum_{i=i_k}^{i_{k+1}-1} p_i \tag{24}$$

In other words, we need to show that the factor $q_k^{LT,\delta}/q_k$ has an upper bound which is smaller than the terms that enter in the sum over $p_i^{LT,\delta}$ (this is analogous to the discussion around (20) above). We want to show that (generalizing (20))

$$\sum_{\tilde{i}<i_k} p_{\tilde{i}}\delta_{\tilde{i}k} \le \min_i \left\{\sum_{\tilde{j}<k} q_{\tilde{j}}\delta_{i\tilde{j}}\right\} \tag{25}$$

Denote by i^* the index that minimizes the right hand side, as – as before – regroup the left hand side

$$\sum_{\tilde{j}<i_k}\left(\sum_{\tilde{i}=i_{\tilde{j}}}^{i_{\tilde{j}}+1} p_{\tilde{i}}\delta_{\tilde{i}k}\right) \leq \sum_{\tilde{j}<k} q_{\tilde{j}}\delta_{i^*\tilde{j}} \tag{26}$$

Again, for each $\tilde{j}$, we have $\left[x_{\tilde{i},y_k}\right] \subset \left[y_{\tilde{j},x_{i^*}}\right]$, so that $\delta_{\tilde{i}k} < \delta_{i^*\tilde{j}k}$. Together with SD, this implies that (26) holds. Therefore, we have shown that

$$\sum_{j\geq k} q_j^{LT,\delta} - \sum_{i>i_k+1} p_i^{LT,\delta} \leq \sum_{j\geq k} q_j - \sum_{i>i_k+1} p_i \tag{27}$$

for all k, thus concluding the proof of the proposition.

■

Part 2. Experimental Evidence and Calibrations

Our experimental results were obtained from surveys conducted on Amazon Mechanical Turk, an online marketplace service hosted by Amazon.com. MTurk allows *requesters* to post tasks that *workers* can complete in exchange for compensation. Typical tasks include data management (e.g. finding the best category for a product), content management (e.g. tagging contents with keywords), and consumer surveys. There is a growing usage of online surveys, and of MTurk in particular, as a tool for research in the social sciences (for an overview of its use in economics, see Horton, Rand and Zeckhauser, 2010). To test the performance of this tool as a means to study decision-making under risk, we ran several surveys consisting of well known choice problems, such as the Allais paradoxes, Samuelson's wager, framing effects and the Ellsberg paradoxes. In all such experiments we found agreement with the traditional lab experiment results.

We posted surveys consisting of several multiple choice questions, typically between four and six, which were either choice problems or elicitations of value (surveys are available upon request). Throughout the surveys, we presented choice problems in the traditional form used in (KT79). A representative example of a choice problem we used is:

Choose between:
Lottery A: \$2000 with 1% chance and \$0 with 99% chance,   Lottery B: \$20 for sure.

Independent lotteries were always presented side by side, as above, to prevent spurious interpretation of correlations. The sample size was typically between 75 and 120 subjects per survey. As requesters, we have no information about workers who complete our surveys, except for their worker ID. Using this information, we find that over 1100 different subjects participated in our surveys, of whom over 60% participated only once. We required two conditions for participation: i) that workers be living in the U.S. (so that subjects had as much as possible a similar understanding of survey questions) and ii) that workers' *reputation index* be above 0.96 out of 1 (see discussion below on incentives). We did not collect demographic information on our subjects. However, other surveys on MTurk workers demographics have found that, compared to the general U.S. population, U.S. MTurk workers are slightly more likely to be female (60%), have a similar age distribution but are somewhat younger on average (ages range from 16 to 60, and 45% are above 30 years old), have higher education level (about half have a bachelor degree or higher) and report an average household income level of $40,000. Even though this data is self reported, it indicates that the pool of workers is very representative of the internet-using population, and reasonably representative of the general population.

We used several approaches to ensure high quality data. Monetary incentives were not feasible due to the large volume of surveys and the range of lottery payoffs involved. Moreover, the evidence on the impact of monetary incentives suggests a modest quantitative correction to levels of risk aversion, but not a qualitative change of the subjects' risk appetite (at least in laboratory settings, see e.g. Grether and Plott 1979, and Holt and Laury, 2002). To understand the workers' motivation, note that they choose tasks in terms of their compensation and interest to them[5]. Once they choose a task, they have a strong incentive to perform, as it can affect their reputation index: requesters have the option to accept or reject a worker's task, and the reputation index captures the percentage of a worker's tasks which were accepted. We systematically discarded from the analysis surveys completed in a very short time, under 45 seconds, as these surveys were likely to have been answered without due attention. We included all other data from choice problems in the analysis. We occasionally introduced test questions, such as a choice between a two-outcome lottery and a sure prospect close to or lower than the lottery's downside. The rate of "wrong" anwers, where the sure prospect was chosen over the lottery, was always negligible. To check for consistency of preferences across subjects, we ran several surveys a few times in identical form. We found the results to be consistent, and data from identical surveys was pooled. To test for robustness, we repeated several surveys while changing some aspects of the presentation. For instance, we varied the ordering of questions, and also the order of presentation of the states of the world; we also varied numerical values, e.g. in the Allais paradoxes and in the evaluation problems. Preferences were largely robust to such manipulations, and we do not report these tests here.

In the next sections, we provide a detailed account of experiments that are important to

[5]Workers seemed to take a personal interest in our surveys, often providing feedback and justifying their choices.

assess our hypothesis of local thinking: those on the role of correlation among lotteries for risk preferences (Section 2.B), and those concerning the role of context dependent evaluation in preference reversals (Section 2.C). Section 2.D considers whether standard calibrations of Prospect Theory can account for the shifts in risk preferences documented in Section 4 of the paper.

2.B Robustness on the role of correlations in experiments

2.B.1 Correlation and the Allais paradoxes

To check the robustness of the results of Section 5.1, whereby the Allais paradox can be turned on and off depending on the correlation structure of the lotteries, we ran several variations of the experiment reported in the paper.

1) We checked that our results do not depend on the particular tabular representation of states of the world. To do so, we replaced the latter with verbal representations of world events such as: "Suppose you can choose between the following two lotteries, whose prizes depend on the draw of a card. Suppose the card deck is arranged such that with 1% chance you draw an ace, with 33% you draw a king and with 66% you draw a number. Lottery 1 yields \$2500 if you draw an ace, and \$0 otherwise, *etc.*" The experimental results were similar to those obtained by expressing lotteries in tabular form.

We also tested whether the results depended on precise knowledge of probabilities. To do so, we replaced risky outcomes for uncertain outcomes, for instance the following problem from (TK92):

"Suppose you can choose between two lotteries, which depend on the difference D between the closing value of the Dow-Jones stock market index today and its closing value tomorrow:

if D is	less than 30 pts	between 30 and 35 pts	more than 35 pts
L_1^z gives	0	z	7500
L_2^z gives	z	2500	2500

where z is the common consequence taking the values 0 and 2500. In both experiments, choice patterns were similar to those obtained in the original correlated version of the Allais paradox: subjects were risk averse in both problems, largely independent of the common consequence z. Indeed, preferences were:

	L_1^{2500}	L_2^{2500}
L_1^0	16%	11%
L_2^0	5%	68%

These results lend further support to our account that the state space shapes risk preferences, via the salience of the allowed states of the world.

2) To check whether the role of correlation in the common ratio paradox of Section 5.2 is due to special properties of compound lotteries, we ran a compound version of the common consequence paradox of Section 5.1, in the following pair of choice problems:

Problem 1: Suppose you are presented with a two-stage game. In the first stage, you have a 66% chance of getting \$0 and ending the game, and a 34% chance of going to the second stage. In the second stage, you will play one Lottery, $L_1 = (2500, .97; 0, .03)$ or $L_2 = (2400, 1)$.

Problem 1': identical to Problem 1, except in that the first stage offered a 66% chance of getting \$2400, instead of \$0. In both problems, subjects choose which lottery to play before the first stage of the game. The preferences were:

	L_1'	L_2'
L_1	6.7%	2.7%
L_2	6.7%	83.9%

Preferences were essentially the same in both Problems, as expected from local thinking and in accordance with the independence axiom. The choice patterns were similar to those observed for the one-stage problem with common consequence $z = \$2400$. As a result, the compounded lottery form did not add much to correlation, suggesting that the salience of the zero payoff in the space state generated by the lotteries may be responsible for choice patterns.

3) Finally, we tested our assumption that when lottery correlation is not described ex-

plicitly, subjects interpret the problem as a choice between two independent lotteries. To do so, we compare the implicitly uncorrelated choice problem,

Problem 2: $L_1 = (2500, .33; 0, .67)$ vs $L_2 = (2400, .34; 0, .66)$.

with a problem where the state space makes the absence of correlation explicit:

Problem 2':

	21.78%	22.78%	11.22%	44.22%
L_1'	2500	0	2500	0
L_2'	0	2400	2400	0

The choice patterns in the two problems, elicited from two different groups of subjects, were indistinguishable: the riskier lotteries L_1 and L_1' were chosen by 46% and 49% of subjects, respectively ($N = 75$ for each problem). The same subjects were then asked to choose between the following correlated version of the same problem:

Problem 2^c:

	1%	33%	66%
L_1^c	0	2500	0
L_2^c	2400	2400	0

Both groups of subjects shifted towards risk aversion, with large majorities (up to 78%) choosing L_2^c over L_1^c.

	L_1	L_2
L_1^c	.20	.1
L_2^c	.26	.44

	L_1'	L_2'
L_1^c	.14	.08
L_2^c	.35	.43

The mirroring behavior of the two groups of subjects supports our assumption that lotteries are interpreted as independent by default.

2.B.2 Further tests on the role of correlation on choice

We also tested the impact of correlations on choice problems that, unlike those of the Allais paradox, do not feature a common consequence state. Consider the following two choice problems:

Problem 3: $L_1 = (40, .67; 120, .33)$ vs $L_2 = (30, .5;\ 110, .5)$

Problem 3':

	1/6	1/3	1/2
L_1':	\$40	\$120	\$40
L_2':	\$110	\$110	\$30

Evidently, Problem 3' is simply a correlated version of Problem 3. The two lotteries have similar expected values, respectively \$66.4 and \$70. The joint preferences were:

	L_1'	L_2'
L_1	11%	28%
L_2	17%	44%

There is a significant shift towards lottery L_2', which wins in the most salient and least likely state. Notably, this occurs even though lottery A wins in 2 out of 3 states. Thus, correlation here affects choice in the way predicted by local thinking.

2.C Robustness on context dependent evaluation and preference reversals

The key implication of our model regarding Preference Reversals is that pricing in the context of choice is different from pricing in isolation because lotteries are being compared to different alternatives. Here we review this basic mechanism by: i) checking for context dependent evaluation in a setting where there are no preference reversals, and ii) providing a detailed description of our experiments on preference reversals.

We begin with point i). To check that context dependent evaluation is a feature of choice extending beyond preference reversals, we asked subjects to price $L_\$ = (600, .4;\ 0, .6)$ in isolation and in the context of a choice with $L_p = (400, 1)$. There are two differences with the preference reversals experiments: a) the expected value of $L_\$$ is much lower than that of L_p so we do not expect reversals of preference between pricing and choice, and b) L_p is a sure gain, which simplifies the pricing problem, so subjects only price $L_\$$. Crucially, since in the comparison with L_p the downside of $L_\$$ is salient, we expect that pricing in choice is lower than pricing in isolation.

To elicit pricing in isolation and in the choice context, we used wording similar to Grether and Plott (1974): "Suppose that you can choose between playing lottery $L_\$$ and getting a sure dollar amount. If this amount is very low, you might choose the lottery. As the dollar amount increases, you might prefer it to the lottery. For what dollar amount would you be indifferent between playing the lottery and getting the dollar amount?" Subjects chose from a menu of prices. We also tried to have subjects enter their own price; in this case the data is more noisy. The experiments confirm that subjects vastly prefer L_p to $L_\$$ and the prices satisfy:

$$P(L_\$|L_p) = 170 < P(L_\$^{iso}) = 230$$

Pricing in isolation is higher not only on average, but also probabilistically: the price distribution for $L_\$$ in isolation, $\mathcal{P}_{L_\$^{iso}}$ first-order stochastically dominates the price distribution for $L_\$$ in the context of choice, $\mathcal{P}_{L_\$|L_p}$, as shown in Figure 1. We ran another version of the survey where instead of the certainty equivalent we elicited the selling price, denoted by $P^{sell}(L_\$)$, which yielded similar results.

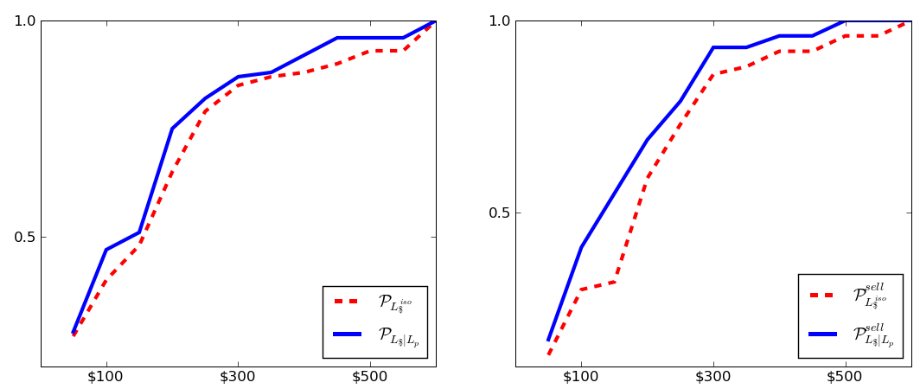

Figure 1: Context Dependent Evaluation

We now consider point ii). To test for preverence reversals, we used the following strategy. Subjects were presented with a choice between two-outcome lotteries L_p and $L_\$$, where L_p is a lottery with a high probability of a low payoff, and $L_\$$ pays a high payoff with low probability. The full state space of the choice problem was presented. Each subject was asked to choose between the lotteries and immediately afterwards to price one of the lotteries. This represents pricing in the choice context. After a few filler questions, that subject was asked to price the remaining lottery. This represents pricing in isolation. Which lottery was priced in which context (choice or isolation) was randomized across subjects. Our identifying assumption is that in the choice context the pricing of each lottery is based on its evaluation in comparison with the alternative lottery. Our account for preference reversals then predicts that they occur between choosing and pricing in isolation, but not between choosing and pricing in the choice context. To elicit pricing in either context, we used the Grether and Plott (1974) wording described above. As in the experiment described in point i), subjects chose from a menu of prices. We also tried to have subjects enter their own price, which yields noisier data. We considered the two lotteries

$$L_p = (4, .97; \;\; 0, .03), \quad \text{vs} \quad L_\$ = (16, .31; \;\; 0, .69)$$

which Tversky et al (1990) found to lead to a high rate of preference reversals. The experimental results allow us to make four points:

1) Standard preference reversals, where individuals choosing L_p over $L_\$$ price $L_\$$ higher, occur in terms of average prices in isolation. To see this, note that the mean prices $P(L_\$^{iso})$ and $P(L_p^{iso})$ for $L_\$$ and L_p, respectively, stated by individuals that had *chosen* L_p *over* $L_\$$ were equal to

$$P(L_p^{iso}) = 4.6 < P(L_\$^{iso}) = 5.2 \; .$$

Since $L_p \succ L_\$$ and $P(L_p^{iso}) < P(L_\$^{iso})$, Preference Reversal occurs with respect to mean prices in isolation.

We then tested whether preference reversals occur not only on average, but also probabilistically. Even though we do not directly replicate Tversky's within-subject test (since each subject priced only one lottery in each context), we checked from the distributions of prices in isolation the likelihood with which Preference Reversals would occur if individuals were drawing prices from these distributions. To do so, we denoted the price distribution for $L_\$$ in isolation by $\mathcal{P}_{L_\$^{iso}}$ and the price distribution for L_p in isolation by $\mathcal{P}_{L_p^{iso}}$. As a first step, we noted that $\mathcal{P}_{L_\$^{iso}}$ nearly first-order stochastically dominates $\mathcal{P}_{L_p^{iso}}$ as shown in Figure 2.

Figure 2 already gives a compelling visual effect of the potential prevalence of preference

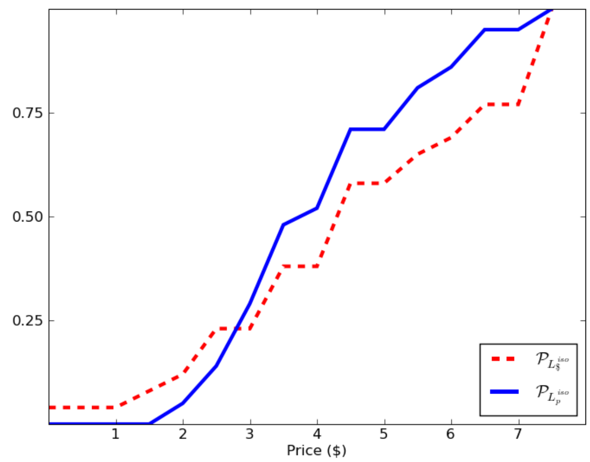

Figure 2: $P(L_\$^{iso})$ roughly first-order stochastically dominates $P(L_p^{iso})$.

reversals (remmeber that this is the price distribution for individual who *chose* L_p over $L_\$$). To quantify this effect, we denoted by $\mathcal{C}_{L_\$^{iso}}$ the cumulative pricing distribution for lottery $L_\$$ in isolation implied by the observed price distribution $\mathcal{P}_{L_\$^{iso}}$. Then the likelihood that a given subject exhibits a Preference Reversal is:

$$\text{Prob}\left[P(L_\$^{iso}) > P(L_p^{iso})\right] = \int \mathcal{P}_{L_p^{iso}}(p) \cdot \Big(1 - \mathcal{C}_{L_\$^{iso}}(p)\Big)\, dp.$$

The data imply that, conditional on $L_p \succ L_\$$, $\text{Prob}\big[P(L_\$^{iso}) > P(L_p^{iso})\big] = 0.52$.

Interestingly, the observed choice and pricing patterns allow us to predict the frequency of non-standard preference reversals, namely those whereby $L_\$$ is preferred to L_p and yet L_p is priced higher in isolation. These non-standard reversals are neither predicted by our model nor by any other theory of preference reversals (Tversky et al 1990, Loomes and Sugden 1983). However, studies on preference reversals have document that these non-standard patterns occur in the data, and experimenters have attributed them to arbitrary fluctuations in evaluation, especially if the two lotteries are evaluated similarly (Bostic et al 1990). Crucially, as stressed by (Tversky et al 1990) the preference reversal phenomenon is interesting precisely because in experiments the non-standard reversals appear to occur substantially less frequently than standard ones, suggesting that the riskier lottery tends to be systematically undervalued in the choice context and overpriced in isolation. If preference reversals were simply due to random fluctuations, the rates of standard and non-standard reversals would be the same on average. Using the above method, we find that in our experiments non-standard reversals occur with frequency around 25%, which is substantially lower than the frequency of standard preference reversals in our data.[6] Thus, also in our experiments reversals cannot be explained by random fluctuations in evaluation.

2) Standard preference reversals *do not* occur between choice and pricing in the context of choice, consistent with the key prediction of our model. Indeed, conditional on $L_p \succ L_\$$, the mean prices in the context of choice were

$$P(L_p|L_\$) = 4.3 > P(L_\$|L_p) = 4.1$$

Now the distribution for $P(L_p|L_\$)$ does not first order stochastically dominate that for $P(L_\$|L_p)$ because subjects attribute similar values to both lotteries in the choice context (indeed, about half the subjects chose each lottery), which magnifies the impact of noise in pricing in the chocie context. Interestingly, we estimate from the price distribution that using prices in the context of choice, around 45% of subjects would exhibit the standard preference reversal while around 49% incur in the non-standard preference reversal in the context of choice. There is no significant difference between these two rates, which suggests that any reversals in the context of choice are due to random fluctuations of evaluation.

To summarize points 1 and 2, our results show that the key feature of our experiments is that the price of the riskier lottery $L_\$$ substantially increases when performed in isolation relative to choice (see Figure 3), while the price of the safer lottery L_p does not change much in the two contexts. These features are both predicted by our model (the latter one is due to the fact that L_p wins with probability close to one, $p = 0.95$). The two prices are generally very close in the context of choice, although the precise ranking in this context is not clearly established by the data. These findings confirm that context-dependent evaluation provides

[6] The difference between these rates is less striking here than those found in the literature, which reports rates for standard preference of around 75% (conditional on choosing L_p), and rates for non-standard reversals of around 15% (conditional on choosing $L_\$$). This may be due to a potentially higher level of noise in our data, or more simply to errors introduced by our method of inferring preference reversals from independent pricing distributions.

an empirically valid account of preference reversals.

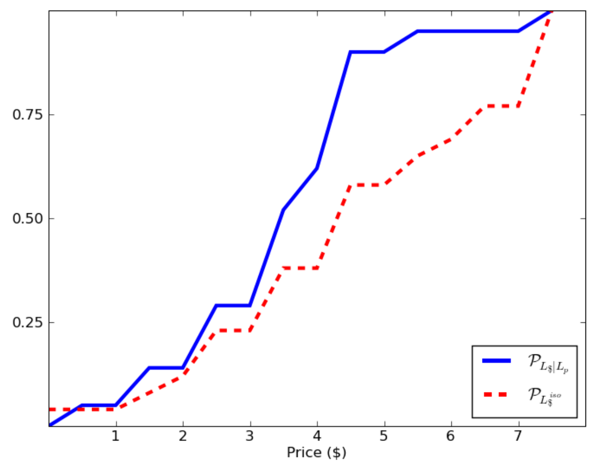

Figure 3: Effect of context on the evaluation of $L_\$$.

3) As is clear from average pricing, some subjects reported prices for L_p above its high payoff of \$4; since the range of prices in the menu extend much beyond the normal range of prices for L_p, this may induce some subjects to erroneously choose a higher price (this problem did not occur for $L_\$$, probably because the expected value is so much lower than the upper limit). Yet, using different menus for the two lotteries might have also induced relative mispricing (we conservatively chose mispricing of L_p because it goes against the standard preference reversals hypothesis). To get a more realistic distribution of prices, and to check for the robustness of the results, we also cleaned the data by truncating the most extreme overpricing.

To do so, we discarded surveys where stated prices exceeded the respective expected values by 50% or more, and we capped the stated price for L_p at \$4. Again, average prices support our prediction; conditional on choosing L_p over $L_\$$ the average prices are:

$$P(L_p^{iso}) = 3.88 < P(L_\$^{iso}) = 4.58$$

This result holds not only for the averages but also for the distributions of prices, in the sense that the distribution for $P(L_\$^{iso})$ essentially first order stochastically dominates the price for

$P(L_p^{iso})$, as before. Let us now compare choice and pricing in the choice context. Again conditional on $L_p \succ L_\$$, the subjects now price L_p on average just above $L_\$$:

$$P(L_p|L_\$) = 3.41 \sim P(L_\$|L_p) = 3.40$$

As before, the distributions do not suggest a significant difference between the two prices. The rate of standard and non-standard preference reversals are also very similar (around 30%). Thus, the truncated yield very similar results to the non-truncated data. This conclusion is robust to different specifications of the truncation (although lowering the threshold for the pricing of L_p affects $P(L_p)$ relatively more than $P(L_\$)$).

4) It is the comparison between the two lotteries, rather than the actual choice, that drives the large change in evaluation between choice and pricing in isolation. To test for this, in another version of the survey subjects were asked to price the lotteries under comparison, but without having to choose between them. We again observe, under the same truncation as in point 3, that $L_\$$ is priced higher in isolation, $P(L_p^{iso}) = 3.58 < P(L_\$^{iso}) = 4.34$. However, $L_\$$ is also priced higher under comparison, albeit much less so that in isolation: $P(L_p|L_\$) = 3.42 < P(L_\$|L_p) = 3.78$. The main difference between the two price distributions is again the large shift in evaluation of $L_\$$; to that extent, the results are still compatible with our predictions. Moreover, while the presentation of the lotteries was identical in the two versions of the survey, we expect that subjects who were asked to choose were more thorough in their examination of the state space. That may help explain why pricing under comparison was half-way between pricing under choice and pricing in isolation.

2.D Calibrations

2.D.1 Long-shot lotteries and bounds on θ

Recall the specification of the salience function in Equation (2) is:

$$\sigma(x, y) = \frac{|x - y|}{|x| + |y| + \theta}.$$

y	1	5	20	50	100	200
	.44	.50	.15	.06	.10	.10

Table 1: Proportion of risk-seeking subjects in longshot lotteries

The parameter θ plays an important role in determining the salience of states whose payoffs are close to zero; it can be interpreted as a cognitive limit to the resolution of payoff magnitude when a payoff approaches zero. This latter property of salience is crucial to determine a local thinker's risk attitudes with respect to longshot lotteries in Section 4.2: a local thinker takes the longshot lottery if and only if $y < \theta \cdot (1 - 2p)/2p$. We thus use the experimental results on longshot lotteries to derive constraints on θ. The following table reports the proportion of 100 subjects who chose the longshot lottery over its expected value as a function of $y \in \{1, 5, 20, 50, 100, 200\}$ when $p = 0.01$.

Thus, for $p = 0.01$, risk seeking decreases strongly between $y = \$5$ and $y = \$20$. This yields $\theta \geq 0.1$. We use the specification $\theta \sim 0.1$ to analyze the remaining results on risk preferences, see 2.C.3b.

2.D.2 Calibrations of Prospect Theory

We consider the ability of standard calibrations of Cumulative Prospect Theory (CPT) to account for the risk shifting experiments of Sections 4, where dramatic shifts in risk preferences occurred as the expected values of lotteries changed. To explore CPT's predictions in these experiments, we use the following standard functional forms for CPT's value function v and probability weighting function π (Tversky and Kahneman, 1992):

$$v(x) = \begin{cases} x^{\alpha} & \text{for } x \geq 0 \\ -\lambda(-x)^{\alpha} & \text{for } x < 0 \end{cases}, \qquad \pi(p) = \frac{p^{\gamma_{KT}}}{(p^{\gamma_{KT}} + (1-p)^{\gamma_{KT}})^{1/\gamma_{KT}}}$$

An agent characterized by $v(x)$ has constant relative risk aversion with coefficient $1 - \alpha > 0$ (implying decreasing absolute risk aversion). The parameter $\lambda > 1$ captures loss aversion. This functional form is standard in applications of Prospect Theory (e.g. Benartzi and Thaler, 1995). Importantly, the estimations of α from different sets of experiments yield very different values. In choices between two outcome lotteries and a sure prospect, Tversky and Kahneman calibrate $\alpha \sim 0.88$, $\lambda \sim 2.25$. In choices involving three-outcome lotteries Wu and Gonzalez (1996) get $\alpha \sim 0.5$ and $\alpha \sim 0.37$. The estimation of the probability weighting function is more stable, locating around $\gamma_{KT} \sim 0.61$ (TK1992). Wu and Gonzalez (1996) estimate $\gamma_{KT} \sim 0.71$, while Prelec's (1998) simplest representation of the weighting function is

$$\pi(p) = e^{-(-\ln p)^{\gamma_P}}$$

with $\gamma_P \sim 0.65$, which is numerically very similar to the original (TK92) calibration. These calibrations are useful references for generating predictions of CPT for our choice problems.

2.D.2a Long-shot lotteries

Consider again the choice between a sure prospect $S = \{y, 1\}$ and a long-shot lottery $L = \{x, p; 0, 1-p\}$ with the same expected value, $y = xp$. A CPT agent evaluates these lotteries as

$$V_{CPT}(L) = \pi(p) \cdot v(y/p), \quad \text{and} \quad V_{CPT}(S) = v(y)$$

If the value function is a power function, then L is preferred to S when $\pi(p)(y/p)^\alpha > y^\alpha$, that is, $\pi(p) > p^\alpha$. This condition is independent of the expected value y; thus, for a given level of probability, a CPT agent's either refuses any long-shot lottery, or takes every long-shot lottery, when that lottery L is compared with its expected value S. This is in contradiction with the experimental results (and also the intuition) whereby subjects prefer the lottery L as long as its expected value is small enough. CPT can account for these results by allowing a more general value function, featuring increasing relative risk aversion. However, several results (including the following) also need decreasing absolute risk aversion to be compatible with CPT, and constant relative risk aversion is a natural candidate for that.

2.D.2b Shifts in risk preferences

Consider again the risk shifting experiments of Section 4 where the lottery loss is kept constant at $l = 20$. Our experimental results, summarized in Table 2, are broadly consistent with the predictions of our model displayed in Figure 3.

\$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

Table 2: Proportion of Risk-Seeking subjects

Can CPT explain these patters through a combination of risk aversion and probability weighting? The answer is that it depends on the values of α and γ_{KT}. We simulated the choice behavior of a CPT agent using a single-agent stochastic choice model (Camerer and Ho 1994), which assumes that all subjects are described by the same calibrated model but their choices follow a stochastic process described by the logistic function $P(L_1 \succ L_2) = 1/(1 + e^{V(L_2)-V(L_1)})$. This process introduces no extra degrees of freedom and ensures that $P(L_1 \succ L_2) = 0.5$ if and only if $V(L_1) = V(L_2)$, and also $P(L_1 \succ L_2) = 1$ if and only if $V(L_1) >> V(L_2)$. The specification above fixes the value function's normalization, $v(1) = 1$. The quantitative CPT predictions thus generated for the baseline KT calibration ($\alpha = 0.88$ and $\gamma_{PT} = 0.61$) are presented in Table 3:

\$10500	1	0.99	0.85	0.52	0.39	0.28	0.21
\$2100	1	1	0.89	0.52	0.37	0.25	0.16
\$400	1	1	0.92	0.51	0.34	0.20	0.12
\$100	1	1	0.93	0.47	0.28	0.15	0.08
\$20	1	1	0.74	0.17	0.08	0.04	0.02
	0.01	0.05	0.2	0.33	0.4	0.5	0.67

Table 3: CPT with baseline calibration $\alpha = 0.88, \gamma_{KT} = 0.61$ (TK1992)

This table reveals two important facts:

1. CPT predicts a drop in risk seeking when $\pi(p) < p$, i.e. when $p > 0.35$, which corresponds broadly to what we observe. However, in its original intution, the probability weighting function was meant to overweight small probabilities (in KT, 1979 this seems to mean $p < 0.1$). As a result, from a conceptual standpoint our prediction of overweighting of probabilities (and thus of risk seeking) for moderate probabilities

is entirely original because it explains why agents can overweight even probabilities as high as $p \sim 0.35$.

2. When α is high (in Table 3, $\alpha = 0.88$), the value function is little curved and risk attitudes depend mainly on the weighting of probabilities. As a result, agents are risk seeking for any p such that $\pi(p) > p$ and risk aversion otherwise, largely independently of y. Thus in the baseline KT calibration $\alpha = 0.88$ for two outcome lotteries, CPT cannot reproduce the experimental patterns. Similarly, one can show that if α is small (e.g. $\alpha = 0.37$ as suggested by Wu and Gonzalez 1996), choices are driven mainly by risk preferences and agents are essentially risk neutral for any $y > 100$, independently of p. This reasoning suggests that an intermediate level of α may ensure that preferences depend both on payoff level x and on its variance (through p), as observed. Table 4 shows the predictions of CPT with the intermediate $\alpha = 0.6$.

\$10500	0.78	0.62	0.52	0.50	0.50	0.49	0.48
\$2100	0.89	0.70	0.54	0.50	0.49	0.48	0.47
\$400	0.94	0.80	0.57	0.49	0.48	0.45	0.44
\$100	0.93	0.82	0.57	0.47	0.44	0.41	0.38
\$20	0.32	0.22	0.13	0.11	0.10	0.10	0.11
	0.01	0.05	0.2	0.33	0.4	0.5	0.67

Table 4: CPT with $\alpha = 0.6, \gamma_{KT} = 0.61$

This calibration which is closer on average to the experimental results, even though differences remain; for instance, agents are not progressively more risk seeking as the expected value y increases; instead they are extremely risk seeking (for $y \geq 100$) and then progressively more risk neutral. To formally quantify this fit, a meaningful measure of distance between the two tables 2 and 4 would need to be developed. Still, our conclusion from this exercise is that the probability weighting function controls how risk attitudes change as p varies, while the value function controls how risk attitudes change as y varies. Since our results describe these two comparative statics, we find that by fine tuning both the probability weighting function and the value function, CPT can recover some of the features of our results.

We finally simulate our model by using the same stochastic choice model. With our

\$10500	0.79	0.79	0.75	0.71	0.69	0.34	0.39
\$2100	0.79	0.83	0.79	0.75	0.72	0.31	0.37
\$400	0.68	0.82	0.82	0.78	0.76	0.27	0.34
\$100	0.38	0.72	0.81	0.79	0.77	0.22	0.30
\$20	0.001	0.002	0.013	0.03	0.04	0.07	0.16
	0.01	0.05	0.2	0.33	0.4	0.5	0.67

Table 5: Local Thinking with $\delta = 0.7$, $\theta = 0.1$

specification of decision weights, a local thinker's evaluation is discontinuous at payoff levels where salience ranking changes. This presents a difficulty in the application of the stochastic choice model, since the latter is very sensitive to changes in the difference between the utilities, particularly if this difference moves away from zero. A straightforward application of the stochastic choice model would therefore lead to a very abrupt transition from risk aversion to risk seeking. To address this issue, we introduce a smoothing of the local thinker's evaluation of lotteries in the form of risk aversion in the value function (an alternative strategy would be to use a different stochastic choice model without this particular feature). We thus simulate our model with the value function $v(x) = x^{0.88}$, and fit (by least mean squares) the parameters θ and δ to the data. Table shows the results with the best fit $\delta = 0.7$, $\theta = 0.1$: The predictions are sensitive to δ, since it modulates large variations in the evaluation of the lotteries: for $\delta = 1$ the lotteries are evaluated at their expected value, and for $\delta \to 0$ they are evaluated at their most salient payoff. The specification $\delta = 0.7$ gives a good agreement with the experimental data. The value of θ is also important, but only for the lower left entry in the table: as for the long-shot lotteries, the experimental data for this experiment imply that when $y = 20$ the lotteries' downside is salient for any p. Setting $p = 0.01$, we again get $\theta \leq 0.4$. For $p > 0.01$ the constraint is eased and for $y \geq 100$, all payoffs are greater or equal to 80 so θ effectively does not count.

Having introduced risk aversion to generate predictions with logistic stochastic choice model, we also confirmed that the fit $\delta = 0.7$ and $\theta = 0.1$ is still consistent with the Allais common consequence and common ratio paradoxes; with risk aversion, the new constraints are $\delta \in (0.19, 0.82)$. Introducing this amount of risk aversion also has an impact on predictions about longshot lotteries. With $\alpha = 0.88$, even when the upside of the lottery is salient, a small majority of subjects are predicted to be risk averse; when the downside is salient, then the vast majority will be risk averse. This is close to the observed rates, see Table **??**.

2.D.2c Robustness of the Allais paradox:

Consider once more the common consequence paradox of Equation (13):

$$L_1^z = (2500,\ 0.33;\ 0,\ 0.01;\ z,\ 0.66), \quad \text{and} \quad L_2^z = (2400,\ 0.34;\ z,\ 0.66)$$

Salience Theory explains the shift toward risk seeking which occurs when z is reduced from 2400 to 0 as being triggered by the fact that when $z = 0$ the two lotteries have the same downside. This implies that the upside 2500 of the riskier lottery L_1^z becomes salient, enhancing risk taking. We have tested this intuition by considering correlated versions of these lotteries. Here we consider a different modification of the paradox which preserves lottery independence and allows for a comparison with CPT. The intuition for the Allais paradox in our model suggests that the shift toward risk seeking occurs only as z approaches 2400. In fact, at any $z > 0$ the risky lottery L_2^z still has a lower downside than L_1^z, which stifles risk seeking.

To test this possibility, we let subjects choose between L_1^z and L_2^z for z = \$0, \$5, \$25, \$100, \$200, and \$2400.

z	0	5	25	100	200
	.56	.71	.78	.80	.83

As predicted by our model, as soon as $z > 0$ as many as 70% of subjects prefer the safe lottery L_2^z even if z is as low as \$5. Interestingly, this boost in the preference for the safer lottery L_2^z as z increase from 0 to \$5 comes from subjects who had chosen L_1^z when $z = 0$. The pattern of responses for $z = 100$ already approaches that reported by KT for $z = 1000$.

Is this sharp reactions of subjects to $z = 5$ consistent with Prospect Theory? The Prospect Theory agent evaluates the difference between the prospects as:

$$V(L_2^z) - V(L_1^z) = v(z)[1 - \pi(.66) - \pi(.34)] - \pi(.33)v(2500) + \pi(.34)v(2400),$$

so that to prefer L_1^z when $z = 0$ (i.e. $L_2^0 - L_1^0 < 0$) subjects must have that $\pi(.34)v(2400) - \pi(.33)v(2500) < 0$. As z increases, since the weighting function is assumed to be subadditive [i.e. $1 > \pi(.66) + \pi(.34)$], the leftmost term in the above expression becomes larger and larger, eventually rendering the entire expression positive for $z = 2400$. One would however require a quite steep value function $v(\cdot)$ for the switch in the sign of the above expression to already occur at $z = 5$, especially because the extent of subadditivity of $\pi(\cdot)$ is typically small. Using the above calibrations for the value function v and the probability weighting function π, we get

$$V(L_2^z) - V(L_1^z) = 0.08 \cdot z^\alpha - 0.348 \cdot 2500^\alpha + 0.355 \cdot 2400^\alpha$$

Preferences are again modeled by the stochastic choice approach. Varying the risk aversion coefficient α has two effects: increasing α increases the impact of z on the demand for the safer option L_1^z as $z \to 2400$, but it also diminishes that impact when z is small. Instead, if α is small, increasing z from 0 to 5 significantly increases demand for L_1^z, but for large z the agent is indifferent between L_1^z and L_2^z. This trade-off is inconsistent with the experimental data, which shows both a significant increase in demand for L_1^z for $z = 5$ relative to $z = 0$, and also a nearly universal demand for L_1^z as $z \to 2400$. Both of these effects are predicted by our model. To see this, note that for any positive value of z in consideration (and $\theta \sim 0.1$), the salience ranking of states is

$$\sigma(0, 2400) > \sigma(0, z) > \sigma(2500, z) > \sigma(z, 2400) > \sigma(2500, 2400) > \sigma(z, z)$$

so that the most salient states favor L_2^z. Instead, for $z = 0$ the salience ranking collapses to the four less salient states:

$$\sigma(2500, 0) > \sigma(0, 2400) > \sigma(2500, 2400) > \sigma(0, 0)$$

which favors the riskier lottery.