

Appendix 1. Proofs

Proposition 1 *The local thinker's over-weighting ω_s has two properties:*

1) Change in a state's probability. *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, i.e. $d\pi_{\tilde{s}} = -\frac{\pi_s}{1-\pi_s}h\pi_{\tilde{s}}$ for all $\tilde{s} \neq s$, then:*

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

2) Change in a state's payoffs. *A change in the payoffs x_s, y_s of a state s which, by increasing (only) that state's salience $\sigma(x_s, y_s)$, improves its ranking from k_s to $k_s - 1$ increases this state's over-weighting ω_s .*

Proof. 1) By definition,

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

Therefore,

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

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

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

Thus

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

The parenthesis on the right hand side can be rearranged to yield

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

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

get the result:

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

2) Let k denote the original salience ranking, and \tilde{k} the modified salience ranking. We know that $\tilde{k}_s = k_s - 1$. Denote by ω_s the overweighting of π_s with respect to k , and by $\tilde{\omega}_s$ the overweighting with respect to \tilde{k} . Write

$$\omega_s - \tilde{\omega}_s = \frac{\delta^{k_s-1}}{\sum_{r \neq s} \delta^{k_r-1} \pi_r + \delta^{k_s-1} \pi_s} - \frac{\delta^{k_s-2}}{\sum_{r \neq s} \delta^{\tilde{k}_r-1} \pi_r + \delta^{k_s-2} \pi_s}$$

This is negative when

$$\delta \left(\sum_{r \neq s} \delta^{\tilde{k}_r-1} \pi_r + \delta^{k_s-2} \pi_s \right) < \sum_{r \neq s} \delta^{k_r-1} \pi_r + \delta^{k_s-1} \pi_s$$

which can be rearranged to yield

$$\sum_{r \neq s} \pi_r \left(\delta^{\tilde{k}_r} - \delta^{k_r-1} \right) < 0$$

By assumption, the rankings k and \tilde{k} differ only by a switch between s and another state t , namely $k_r = \tilde{k}_r$ for all $r \neq s, t$ and $k_s = \tilde{k}_t$ (which together imply that $k_t = \tilde{k}_s$). Thus $k_t = \tilde{k}_t - 1$. Then the above expression becomes

$$\sum_{r \neq s, t} \pi_r \delta^{k_r} (\delta - 1) + \pi_t \delta^{k_t-1} (\delta^2 - 1) < 0$$

and the result follows since both terms are negative. ■

Appendix 2. Experimental Evidence and Calibrations

Our experimental results were obtained from surveys conducted on Amazon Mechanical Turk (MTurk), 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 elicitation 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 them¹. 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" answers, 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

¹Workers 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 checks whether standard calibrations of Prospect Theory can account for the shifts in risk preferences documented in Section 4 of the paper, and quantifies the local thinking parameters (δ, θ) most consistent with this evidence.

2.B Robustness of 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 explicitly, subjects interpret the problem as a choice between two independent lotteries. To do so, we compare the implicitly uncorrelated choice problem,

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

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

		21.78%	22.78%	11.22%	44.22%
Problem 2':	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:

		1%	33%	66%
Problem 2 ^c :	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	L'_2
L_1^c		.20	.1	.14	.08
L_2^c		.26	.44	.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 $L_2 = (30, .5; 110, .5)$

		1/6	1/3	1/2
Problem 3':	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 of 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 $L_s = (600, .4; 0, .6)$ 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_s 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_s . Crucially, since in the comparison with L_p the downside of L_s 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 preference 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

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

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

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

Since $L_p \succ L_{\S}$ and $P(L_p^{iso}) < P(L_{\S}^{iso})$, 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_{\S} in isolation by $\mathcal{P}_{L_{\S}^{iso}}$ and the price distribution for L_p in isolation by $\mathcal{P}_{L_p^{iso}}$. As a first step, we noted that $\mathcal{P}_{L_{\S}^{iso}}$ 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: Cumulative distributions of $P(L_§^{iso})$ and $P(L_p^{iso})$.

reversals (remember 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:

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

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

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 documented 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.² 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 choice 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

²The 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 $L_p \succ L_\S$, the subjects now price L_p on average just above L_\S :

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

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 $P(L_\S)$).

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_\S is priced higher in isolation, $P(L_p^{iso}) = 3.58 < P(L_\S^{iso}) = 4.34$. However, L_\S is also priced higher under comparison, albeit much less so than in isolation: $P(L_p|L_\S) = 3.42 < P(L_\S|L_p) = 3.78$. The main difference between the two price distributions is again the large shift in evaluation of L_\S ; 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:

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

The parameter θ plays an important role in determining the salience of states whose payoffs

y	1	5	20	50	100	200
	.44	.45	.15	.06	.10	.10

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

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):

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

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 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 introspection) 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.

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

\$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

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 $V(L_1) \gg V(L_2)$. The specification above fixes the value function’s normalization, $v(1) = 1$. The quantitative CPT predictions 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 intuition, the probability weighting function was meant to overweight small probabilities (in KT, 1979 this seems to mean $p < 0.1$). From a conceptual standpoint, our prediction of overweighting (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 approximates reasonably the experimental results, even though differences remain; for instance, at low p agents are not progressively more risk seeking as the expected value y increases; instead they are extremely risk seeking (for $y = 100, 400$) and then progressively more risk neutral. Still, our conclusion from this exercise is that - from a formal standpoint - 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, by fine tuning both the probability weighting function and the value function, CPT can recover some of the features of our results.

Finally, we simulate our model and recover a quantification of the local thinking parameters (δ, θ) . One issue with this simulation exercise is that in our specification of decision weights, a local thinker's evaluation is discontinuous at payoff levels where salience ranking changes, particularly if the agent's utility is linear. This creates a difficulty in the application of the Camerer and Ho (1994) stochastic choice model, since the latter is very sensitive to changes in the difference between the utilities, particularly if this difference moves away

\$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 $\delta = 0.7$, $\theta = 0.1$

from zero. This might potentially lead to a very abrupt transition from risk aversion to risk seeking as y changes. To address this issue, we smooth the local thinker’s evaluation of lotteries by introducing mild concavity in the value function, which is an approach that is fully consistent with our model. We simulate our model using the standard TK (1992) calibration of the value function $v(x) = x^{0.88}$. Table 6 shows the results with the calibration $\delta = 0.7$, $\theta = 0.1$: 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 $\delta \rightarrow 0$ 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: 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 than or equal to 80 so θ effectively does not count.

Having introduced risk aversion to simulate the logistic stochastic choice model under local thinking, as a consistency check we confirmed that the value function $v(x) = x^{0.88}$ together with the fitted parameters $\delta = 0.7$ and $\theta = 0.1$ still predicts the Allais common consequence and common ratio paradoxes as well as the main features of subjects’ risk taking in longshot lotteries (all phenomena that in the paper we rationalized using linear utility). With risk aversion, the new constraints which guarantee the Allais paradoxes are $\delta \in (0.19, 0.82)$, which are verified at $\delta = 0.7$. With respect to longshot lotteries, when $\alpha = 0.88$ even if 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 1.