Regret and Ambiguity Aversion*

Daniel Krähmer† and Rebecca Stone‡

February 13, 2006

Abstract

The paper provides a behavioural explanation for ambiguity aversion based on the fear of regret. We propose a new model of regret in which the agent directly cares about the performance of his actual choice relative to the choice that would have been best ex post. An ambiguous alternative is represented as a compound lottery. The basic idea is that selecting a compound lottery reveals information, which alters the ex post assessment of what the best choice would have been, inducing regret. We provide sufficient conditions under which regret implies uncertainty aversion in the sense of quasi-concave preferences over compound lotteries.

Keywords: Ambiguity Aversion, Regret, Counterfactual Reasoning, Reference Dependence, Information Aversion, Hindsight Bias

JEL Classification: C72, D11, D81, D83

---

*Aknowledgements:

†Freie Universität Berlin, Institut für Wirtschaftstheorie, Boltzmannstr. 20, 14195 Berlin, Germany, ++49-(0)30-83855223, kraehmer@wiwiss.fu-berlin.de. Financial support through research grant SFB/TR 15 is gratefully acknowledged.

‡Department of Economics, University of Leicester, University Road, Leicester LE1 7RH, UK, ++44-(0)116-252-5368, rs181@le.ac.uk

1 Introduction

In this paper, we provide a behavioural explanation for ambiguity aversion based on an agent's fear of regret. Regret is the negative emotion that arises from an agent's perception that he should have chosen otherwise. Ambiguity aversion refers to an agent's distaste for making choices under conditions of uncertainty in which some relevant probabilities are unknown ex ante as opposed to conditions of risk in which all relevant probabilities are known in advance. Crucial for our argument is the observation that the resolution of an ambiguous lottery reveals information about the unknown probabilities of the events associated with that lottery. Thus, when an agent chooses such a lottery its resolution may alter his ex post assessment of what he optimally should have done at the moment of choice, leading him to regret his past choice. By contrast, the resolution of a risky prospect reveals no such information and so leaves the agent's original assessment of the wisdom of his past choices in tact.

To illustrate, consider the two-urn example presented by Ellsberg (1961). One of the urns, the "unambiguous urn", contains 50 black and 50 white balls. The other, the "ambiguous urn", contains 100 balls that are either black or white in unknown proportions. Suppose that the agent represents the ambiguous urn as compound lottery, where in the first stage, the composition of the urn is selected, and in the second stage a ball is drawn, conditional on first stage composition. Thus, there are 101 possible states of the ambiguous urn, each corresponding to one configuration of black and white balls. The agent has to select the urn from which a ball is to be drawn, knowing that he will receive £100 if a black ball is drawn from it and £0 if a white ball is drawn. For simplicity, suppose that the agent's prior beliefs assign equal probability to each of the possible states of the ambiguous urn. Then the two urns yield the same probability distribution over monetary outcomes. An agent is ambiguity averse if he strictly prefers the unambiguous urn.[1]

We model regret by assuming that the agent evaluates his past decisions in light of his ex post information. More precisely, after the ball has been drawn, the agent compares his actual payoff with the posterior expected payoff that would have been generated by the choice that is best, *conditional* on his ex post knowledge. This posterior expected payoff is the agent's reference point. We assume that beyond his material payoff, the agent obtains an emotional payoff, which is given by the difference between his actual payoff and this reference point.

When the agent chooses the ambiguous urn, the outcome reveals something about its true composition. Accordingly, he revises his assessment of its expected payoff. If he wins the bet, he infers that *his actual choice was the best choice* since his posterior assigns greater probability

[1]In this case, due to symmetry, he also strictly prefers the known urn if he is offered a bet that wins £100 if a white ball is drawn and £0 otherwise, which is inconsistent with expected utility maximisation.

to states in which there are more black balls than white balls. In this case he rejoices as he compares his payoff of £100 to the posterior mean of the ambiguous urn, which, given his posterior, is strictly greater than £50 and less than £100. If instead he loses the bet, his posterior puts greater probability weight on states in which there are more white balls than black balls. Accordingly, he infers that *the unambigous urn would have been the best choice*. In this case, he experiences regret as he compares his actual payoff of £0 to the expected payoff of the known urn of £50.

When the agent chooses the unambiguous urn instead, no outcome leads the agent to revise his assessment of either urn's expected payoff. Thus, irrespective of winning or losing the bet, the agent's actual choice remains optimal from an ex post perspective and in each case he compares his actual payoff to the posterior mean of the known urn, which is £50. Thus, he regrets and rejoices in equal measure, and in expected terms regret and rejoicing exactly offset each other.

By contrast, when he chooses the ambiguous urn, his regret in the event of a bad outcome *more than offsets* his rejoicing in the event of a good outcome because *he may change his assessment of what the optimal choice is* in the light of his ex post information, pushing up his reference point. Accordingly, his expected regrets exceed his expected rejoicing and he suffers an overall utility loss.

Notice that the fact that the agent assesses ex post what he should optimally have done induces a concavity in his regret function. This implies that from an ex ante point of view, the agent's expected regret looms larger than his expected rejoicing. In this sense, our model *endogenously* generates a form of loss-aversion. (We do not assume that the agent's regret function exhibits loss-aversion per se.)

In effect, our agent compares himself to a hypothetical counterfactual self who faces the same choice problem as he did but with the benefit of the knowledge that he has gained with hindsight. The expected payoff of this "hindsight self" enters negatively into the actual self's objective function and so the actual self benefits from ensuring that his hindsight self has less information. Since a draw from the unambiguous urn, in contrast to a draw from the ambiguous urn, does not provide the hindsight self with any information about the state of the world, choosing the unambiguous urn over the ambiguous urn shields the agent from regret.

In this paper, we generalise the example outlined above. In the spirit of Segal (1987), we model an ambiguous urn as a compound lottery whose first stage reflects the (second order) uncertainty with respect to the true (first order) distribution of the urn. Our two main propositions provide sufficient conditions such that our notion of regret leads to uncertainty aversion in the sense of Schmeidler (1989). An uncertainty averse agent prefers a convex mixture of two compound lotteries, its "component lotteries", over at least one them. Intuitively, regret and uncertainty aversion are linked because the possible first order distributions of a mixture of compound lotteries are less dispersed than those of its component lotteries and are therefore more similar to one another than those of at least one of them. This suggests that the information about the mixture's true distribution that is generated by a realisation of the mixture may be less conclusive than the information about the component's true distribution that is generated by a realisation of the component.

We identify two environments under which this intuition is formally true. In the first environment, we consider component lotteries whose first order distributions satisfy certain monotonicity conditions—in the sense of first order stochastic dominance.[2] In particular, these conditions hold automatically if there are only two outcomes such as "winning" or "losing" the bet as is the case in typical Ellsberg experiments. Thus, our notion of regret generates uncertainty aversion in all environments with two outcomes.

In the second environment, we impose certain symmetry conditions which essentially require that for each first order distribution (of a component lottery) that the agent considers possible, he also considers possible and equally likely the symmetric first order distribution in which outcomes are reversed. This seems to be a plausible assumption in situations in which the agent is completely ignorant about the composition of an urn: the principle of insufficient reason prescribes that the agent consider all possible urn compositions as equally likely. In this symmetric environment we can show not only that regret gives rise to uncertainty aversion, but also that uncertainty aversion is equivalent to second order risk aversion in the sense of an aversion to second order mean preserving spreads.

The model of regret we are proposing is very different from the original regret theory proposed by Loomes and Sugden (1982) and Bell (1982) in at least two respects. First, whereas in regret theory, the reference point is exogenously determined, in our model the reference point emerges from our explicit characterisation of the agent's counterfactual thought processes. Second, regret in our model is a response to the realisation of (second order) uncertainty about the (first order) risk the agent faces, whereas in regret theory regret is a response to the realisation of uncertainty about final outcomes. Given these significant differences between our account and regret theory, we are untroubled by recent experimental work that calls into question the empirical significance of the latter.[3]

Before we present the formal analysis in sections 3 to 5, we offer some motivation for our

[2] Similar conditions are provided in Halevy and Feltkamp (2005). We discuss the relation of their to our work in detail below.

[3] This work finds regret not primarily responsible for producing many of the empirical violations of expected utility theory such as violations of transitivity and monotonicity that inspired the formulation of regret theory. See, for instance, Harless (1992) and Starmer and Sugden (1993 and 1998).

approach by reviewing psychological evidence that supports our proposed model of preferences and points to a relationship between regret and ambiguity aversion. We postpone the discussion of how our approach relates to the existing literature on ambiguity aversion to section 6 and conclude in section 7.

2 Psychological motivation

Our model of the agent's preferences is motivated by two psychological observations. The first is the observation that a person's evaluation of an outcome is often determined by its comparison with salient *reference points* as well as by its intrinsic characteristics.[4] Specifically, people appear to care directly about the comparison of outcomes of their choices to pertinent counterfactuals such as what could, should, or might have happened. For example, a person experiences disappointment when an outcome falls short of the outcome he was expecting, while he experiences regret when he perceives that he would have received a higher payoff had he chosen differently.[5]

Second, psychological evidence indicates that reference points are often constructed not only ex ante but also *ex post*, once the outcomes of the events to which they pertain have been realised. This is the central claim of Kahneman and Miller's (1986) norm theory. Norm theory was developed to supplement the standard view that outcomes tend to be evaluated with respect to reference points that are predetermined by prior processes of anticipation and expectation formation with a conception of evaluation as an inherently backward looking phenomenon that is directly influenced by the experience of the events themselves. According to Kahneman and Miller (1986), "events in the stream of experience ... are interpreted in a rich context of remembered and constructed representations of what it could have been, might have been or should have been" (p.136).

To illustrate, consider for a moment the evaluative process that gives rise to disappointment. A person experiences disappointment when he perceives that his prior expectations about an outcome associated with a chosen option have been confounded. Standard conceptions posit the existence of a reference point that depends on expectations of the outcome that were formed by the agent prior to outcomes being realised.[6] In the spirit of Norm Theory, an alternative conception would allow an agent's evaluation to be influenced by the ex post knowledge that he

[4] Reference dependence is a feature of Kahneman and Tversky's (1979) Prospect Theory. For a review of some of the evidence on reference dependence see, for example, Rabin (1998).

[5] For experimental evidence that indicates that disappointment and regret are emotions that influence choice behaviour see, for example, Mellers et al. (1999), Mellers (2000) and Mellers and McGraw (2001).

[6] See, for instance, Koszegi and Rabin (2005).

acquires as a result of experiencing the actual outcome. This could account for the psychological finding that, even holding his ex ante expectations constant, a person's experience of failure tends to be rendered more painful when he nearly succeeded than when success was a long way off.[7] If, as seems plausible, nearly succeeding leads the agent to assign greater likelihood to states in which he was more likely to have succeeded than does the experience of failing by a larger margin, then the former experience will elicit a higher reference point and hence greater disappointment than the latter.

Our model of regret shares with this model of disappointment the assumption that the agent's evaluation is informed by knowledge gained in hindsight. The classic regret theory that was developed by Loomes and Sugden (1982) and Bell (1982, 1983) also shares this assumption to some extent. In regret theory, an agent's utility depends on the ex post comparison of the outcome of his chosen option with those of unchosen alternatives. However, while the comparison is made ex post, the *identities* of the options in the agent's original choice set whose payoffs will determine the reference point are fixed in advance. By contrast, in our model, the identity of the option whose payoff will serve as the reference point is itself determined in the light of the agent's ex post information and, thus, may actually be the agent's chosen option: the relevant counterfactual is the payoff that the agent judges with hindsight that he *should* have achieved, not merely the payoff that he perceives he *would* have gotten had he made a different choice.

In this sense, our agent is unduly harsh on his past self, for he judges himself in the light of posterior knowledge that was inaccessible at the moment of decision. However, psychological evidence indicates that when judging past decisions in hindsight, people consistently exaggerate what could have been anticipated in foresight, a tendency known as the hindsight bias (Fischoff 1982). Like our regretful agent, people tend to overestimate the prior inevitability of those outcomes and so curse themselves (and others) for failing to anticipate them.

Our theory requires that the agent decomposes the overall uncertainty to which he is exposed into two components: the uncertainty about the true distribution and the risk about the outcome, conditional on the true distribution. He blames himself only for his failure to know what he now knows about the true distribution, but not for failing to predict the actual outcome per se. Of course, whether or not he does this ultimately depends upon his psychological

[7] Tversky and Kahneman (1982) ask subjects to consider the following example:

Mr Crane and Mr Tees were scheduled to leave the airport on different flights, at the same time. They traveled from town in the same limousine, were caught in a traffic jam, and arrived at the airport 30 minutes after the scheduled departure time of their flights. Mr Crane is told that his flight left on time. Mr Tees is told that his flight was delayed, and just left five minutes ago. Who is more upset?

96% of subjects stated that Mr Tees would be more upset.

makeup rather than the world. In principle it may depend upon the way the agent perceives his choice problem and consequently upon certain contextual variables.

Our agent can be viewed as making a cognitive mistake rather than suffering from a motivational bias. If the agent could have known with foresight what he knows with hindsight then blaming himself for making an incorrect decision would seem like an appropriate emotional response. It is sensible to blame oneself for a failure that was forseeable and, hence, potentially avoidable, but not for a failure that was inherently unforseeable. Of course, by assumption, our agent could not have known with foresight what he knows only with hindsight, and it is this which renders his propensity to blame himself irrational. Thus, his emotional response to an outcome can be viewed as a rational reaction to the mistaken belief that he knew then what he now knows only with the benefit of hindsight. Indeed, if he was not prone to a hindsight bias then his behaviour would be indistiguishable from that of an expected utility maximiser: his reference point would always be given by the ex ante expected payoff of his chosen option and so he would regret and rejoice in equal measure.

The idea that people distinguish between uncertainty that arises from ignorance and the uncertainty that arises from intrinsic randomness has received backing in the psychology literature. Frisch and Baron (1988) conceive of ambiguity as precisely "the subjective experience of missing information relevant for a prediction", and, in support of this conception, Brun and Teigen (1990) find that subjects prefer guessing the outcome of an uncertain event before it has occurred to guessing it after it has occured but before they know it. Notice that while an event that has not yet occurred may be regarded as unknowable in principle, an event that has already occurred may be regarded as in principle knowable, in which case any failure to know the outcome is attributable to ignorance rather than inherent randomness. Thus, if subjects distinguish between different types of uncertainty in this way, then in the terms of our model subjects would make an analogous distinction between the postdiction and prediction scenario. Our model would then predict that they would experience greater regret when engaging in postdiction than prediction, leading them to prefer the latter over the former. In fact, Brun and Teigen's subjects commonly cite as a reason for their preference for prediction that wrong postdictions are much more embarassing than wrong predictions.

Psychologists have conducted experiments in which subjects' behaviour is influenced by the amount of feedback they expect to obtain. Specifically, subjects become more willing to choose a riskier gamble if they will learn the outcome of the gamble regardless of the choice they make, while such feedback on the safer alternative is provided only in the event that they choose it.[8] This suggests that people make choices to minimise their exposure to information about

[8] For a review of this evidence, see Zeelenberg (1999).

the outcomes of unchosen alternatives, and supports a version of regret theory that was first suggested by Bell (1983) according to which people are averse to such feedback.[9] As previously discussed, our agent also displays information aversion. However, our agent is averse to a different kind of information. Whereas Bell's agent is averse to feedback about outcomes, our agent is averse to feedback about the true probability distributions associated with the options in his choice set. This is unsurprising since Bell's theory is a theory of how regret arises under conditions of risk. Hence, the distribution of any lottery is determined from the outset. By contrast, ours is a theory of a different kind of regret that arises under conditions of uncertainty and rationalises Ellsberg type behaviour as the manifestation of the agent's aversion to feedback about the state of the ambiguous urn.

Finally, Ritov and Baron (2000) provide some evidence which points to a link between ambiguity aversion and regret. In their experiments, they find that the omission bias is intensified when subjects choose under conditions of uncertainty rather than mere risk. An ommission bias is a tendency to prefer inaction to action even if the consequences of inaction are worse than the consequences of action. Psychological evidence suggests that this may be due to the tendency of commissions to elicit greater regret than ommissions, at least in the short-run.[10] If, as our model suggests, regrets may be intensified under conditions of ambiguity, this would help to explain why the ommission bias is also amplified under these conditions.

3 The model

In order to fix ideas, we begin with an example to which we will make reference throughout the exposition of the general model.

Example An agent can draw a ball from one of two urns, α and γ, which contain black and white balls. If he draws a black ball he receives \$100, while he receives nothing if he chooses a white ball. There is uncertainty about an urn's composition each of which can be "good" or "bad" with equal probability. If the composition is "good", the fraction of black balls in urn α is $\frac{1}{2}+\eta^{\alpha}$ where $\eta^{\alpha}\in\left[\frac{1}{2},1\right]$, and an agent who chooses it faces the lottery $\Lambda_G^{\alpha}=\left(100,\frac{1}{2}+\eta^{\alpha};0,\frac{1}{2}-\eta^{\alpha}\right)$. If the composition is "bad", the fraction of black balls is $\frac{1}{2}-\eta^{\alpha}$ and the associated lottery is $\Lambda_B^{\alpha}=\left(100,\frac{1}{2}-\eta^{\alpha};0,\frac{1}{2}+\eta^{\alpha}\right)$. Similarly, the lotteries associated with the "good" and "bad" compositions of urn γ are, respectively, $\Lambda_G^{\gamma}=\left(100,\frac{1}{2}+\eta^{\gamma};0,\frac{1}{2}-\eta^{\gamma}\right)$ and $\Lambda_B^{\gamma}=\left(100,\frac{1}{2}-\eta^{\gamma};0,\frac{1}{2}+\eta^{\gamma}\right)$ where $\eta^{\gamma}\in\left[\frac{1}{2},1\right]$. We assume that $\eta^{\alpha}>\eta^{\gamma}$ which intuitively

[9] For a detailed discussion of feedback effects in a Bell-type framework, see Krähmer and Stone (2005).

[10] According to Gilovich and Medvec (1994), commissions tend to elicit greater regret in the short run while ommissions elicit greater regret in the long run.

means that urn α is the more ambiguous urn. In the extreme case in which $\eta^{\gamma} = \frac{1}{2}$, there is no uncertainty about the composition of urn γ. Figure 1 illustrates the compound lottery that describes urn α. $\square$

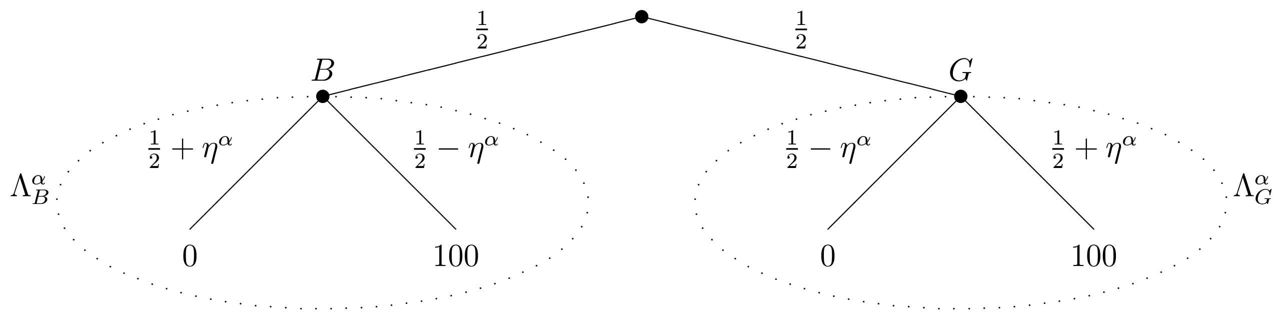

Figure 1: urn α

In the general environment, we represent an ambiguous urn as a general compound lottery. Thus, we describe urns by horse bets over roulette lotteries (Anscombe and Aumann, 1963). In our theory the agent will care directly about the information that his choice reveals about the composition of an urn. In general, this includes information that a draw from one urn might reveal about a different urn. In Ellsberg type problems, however, urns are typically independent. Therefore, we abstract from cross informational effects and model urns as independent.

There is a finite number of actions (urns) $d \in D$. For each $d \in D$, there is a finite marginal state space $\Omega^d = \left\{1, ..., \overline{\omega}^d\right\}, \overline{\omega}^d \in \mathbb{N}$. The product state space is given by $\widetilde{\Omega} = \Pi_{d \in D} \Omega^d$. For $\overline{x} \in \mathbb{N}$, let $\mathcal{X} = \{1, ..., \overline{x}\}$ be a finite outcome space. Let L be the set of all (roulette) lotteries with values in $\mathcal{X}$. Urn d is represented by the compound lottery $\Lambda^d : \widetilde{\Omega} \to L$, which gives the roulette lottery $\Lambda^d(\widetilde{\omega}) \in L$ in state $\widetilde{\omega} \in \widetilde{\Omega}$. We assume that Λ^d depends only on marginal states:

$$\Lambda^d\left(\omega^d, \omega^{-d}\right) = \Lambda^d\left(\omega^d, \theta^{-d}\right) \quad \forall \omega^{-d}, \theta^{-d} \in \Pi_{d' \in D \backslash d} \Omega^{d'}. \tag{1}$$

Hence, we also write $\Lambda^d\left(\omega^d\right)$ for $\Lambda^d\left(\omega^d, \omega^{-d}\right)$. The set of all probability distributions on $\mathcal{X}$ is denoted by $\Delta(\mathcal{X})$. We use the convention that elements in $\Delta(\mathcal{X})$ are *column* vectors.[11] Let

$$h^d_{\omega^d} = \left(h^d_{1\omega^d}, ..., h^d_{\overline{x}\omega^d}\right) \in \Delta(\mathcal{X})$$

be the distribution of $\Lambda^d\left(\omega^d\right)$. Thus, $h^d_{x\omega^d}$ is the probability that action d generates outcome x when the marginal state is ω^d. We summarise all the possible distributions associated with action d in the $\overline{x} \times \overline{\omega}^d$ Markov matrix[12]

$$h^d = \left(h^d_1, ..., h^d_{\overline{\omega}^d}\right).$$

[11] For space reasons, we write columns vectors as row vectors.

[12] Throughout, we use the convention that the *columns* of Markov matrices sum up to one.

Example (continued) Each urn has two marginal states, a "bad" state (state 1) and a "good" state (state 2), and so the resulting product space has four states. There are two possible outcomes: the agent either receives \$0 (outcome 1) or \$100 (outcome 2). The roulette lotteries associated with the compound lottery Λ^{α} that describes urn α are

$$\Lambda^{\alpha}(1,1) = \Lambda^{\alpha}(1,2) = \Lambda^{\alpha}(1) = \Lambda^{\alpha}_{B},$$

$$\Lambda^{\alpha}(2,1) = \Lambda^{\alpha}(2,2) = \Lambda^{\alpha}(2) = \Lambda^{\alpha}_{G}.$$

So the Markov matrix associated with urn α is

$$h^{\alpha} = \begin{pmatrix} \frac{1}{2} - \eta^{\alpha} & \frac{1}{2} + \eta^{\alpha} \\ \frac{1}{2} + \eta^{\alpha} & \frac{1}{2} - \eta^{\alpha} \end{pmatrix}.$$

The description of urn γ is analogous. □

The agent holds a prior belief $\widetilde{\pi} \in \Delta\left(\widetilde{\Omega}\right)$ over states. This prior may be subjective or objective. Let $\pi^{d} \in \Delta\left(\Omega^{d}\right)$ be the marginal prior: $\pi^{d}\left(\omega^{d}\right) = \sum_{\omega^{-d}} \widetilde{\pi}\left(\omega^{d}, \omega^{-d}\right)$. We assume that states are independent:

$$\widetilde{\pi}\left(\widetilde{\omega}\right) = \Pi_{d \in D} \pi^{d}\left(\omega^{d}\right) \qquad \forall \widetilde{\omega} = \left(\omega^{d}\right)_{d \in D}. \tag{2}$$

We denote by

$$X^{d} \in L$$

the reduced lottery that results from first playing out the state $\widetilde{\omega}$ according to $\widetilde{\pi}$ and then playing out the lottery $\Lambda^{d}\left(\widetilde{\omega}\right)$. Because of (1) and (2), the family $\left(X^{d}\right)_{d \in D}$ is stochastically independent. Let

$$p^{d} \in \Delta\left(\mathcal{X}\right)$$

be the distribution of X^{d}.

Example (continued) In the example, the marginal priors are all equal to one half: $\pi^{\alpha}(1) = \pi^{\alpha}(2) = \frac{1}{2}$ and $\pi^{\gamma}(1) = \pi^{\gamma}(2) = \frac{1}{2}$. Independence in this case means that prior beliefs are given by $\widetilde{\pi}(1,1) = \widetilde{\pi}(1,2) = \widetilde{\pi}(2,1) = \widetilde{\pi}(2,2) = \frac{1}{4}$. For $d \in \{\alpha, \gamma\}$, the reduced lottery is $X^{d} = \left(100, \frac{1}{2}; 0, \frac{1}{2}\right)$, and its distribution is $p^{d} = \left(\frac{1}{2}; \frac{1}{2}\right)$. □

Finally, we remark that it is without loss of generality to assume that for all $d, d' \in D$: $\Omega^{d} = \Omega^{d'}$ and $\pi^{d} = \pi^{d'}$. This is because states can always be appropriately "split".[13] Thus, we drop action

[13] To illustrate the argument, suppose D contains two actions a, b with $\overline{\omega}^{a} \neq \overline{\omega}^{b}$ and $\pi^{a} \neq \pi^{b}$. Then define $\Omega = \Omega^{a} \times \Omega^{b}$ and $\pi\left(\omega^{a}, \omega^{b}\right) = \pi^{a}\left(\omega^{a}\right) \pi^{b}\left(\omega^{b}\right)$.

indices from $\Omega, \overline{\omega}$, and π.[14]

Preferences The agent cares both about his material payoff and about his ex post evaluation of the performance of his action relative to the performance of a *reference action.* In the spirit of norm theory (Kahnemann and Miller, 1986), we first suppose that the reference action is constructed *ex post*, that is, after the agent has observed the outcome of his choice. Second, we assume that the reference action is the *ideal* action that the agent thinks he should have chosen had he known with forsight what he knows with hindsight.

The ex post comparison between what the agent actually got and what he perceives he should have gotten induces an emotional reaction: if the actual payoff falls short of his reference action's payoff, he experiences regret and his utility falls; otherwise, he rejoices and his utility rises.

Formally, the agent perceives that in period 0, Nature determines the state $\widetilde{\omega} \in \widetilde{\Omega}$, *once and for all,* according to $\widetilde{\pi}$, in a move that the agent cannot observe. In period 1, the agent selects an urn d, the lottery $\Lambda^d(\widetilde{\omega})$ is played out, and the lottery X^d realises. If $X^d = x$, the *material* component of the agent's instantaneous utility is given by

$$\phi(x),$$

where $\phi : \mathcal{X} \to \mathbb{R}$ is the material payoff function. We also write ϕ_x for $\phi(x)$ and identify ϕ with the column vector $\phi = (\phi_1, ..., \phi_{\overline{x}}) \in \mathbb{R}^{\overline{x}}$. Without loss of generality, ϕ is increasing in x.[15]

To describe the *emotional* component of the agent's instantaneous utility, we define the reference action as the action that would maximise the agent's material payoff in a hypothetical (counterfactual) period 2 choice problem, in which the agent could select among the *same* urns in the light of his ex post knowledge about the state of the world. Let

$$\widehat{\Lambda}^d(\widetilde{\omega}) \in L$$

be a stochastically independent version of $\Lambda^d(\widetilde{\omega})$ and let

$$\widehat{X}^d \in L$$

be the lottery that results from Nature's choice of $\widetilde{\omega}$ *in period 0* and from a subsequent draw of $\widehat{\Lambda}^d(\widetilde{\omega})$ in period 2. That is, the agent perceives that in his counterfactual choice problem the

[14] It is common in the literature to refer to a mapping $\Lambda : \widetilde{\Omega} \to L$ as an *act*, and to postulate preferences over the set of all acts. We deviate from this convention, because our independence assumption would make it necessary to introduce a state space component for each act. For our purposes, it is sufficient to look at a finite number of fixed acts with a simple state space representation. We use the term action also to highlight that we are considering preferences over specific acts only.

[15] Otherwise, re-order outcomes appropriately.

composition of the urn is the same as in his actual choice problem, but that a new, conditionally independent outcome is drawn from the urn.

Conditional on $X^d = x$, the agent's expected material payoff from a hypothetical period 2 choice $\widehat{d}$ is denoted[16]

$$\widehat{m}^{\widehat{d}}_{(d,x)} = E\left[\phi\left(\widehat{X}^{\widehat{d}}\right)|X^d = x\right].$$

The agent's *reference action*, $\mu_{(d,x)}$, is then defined as a best hypothetical period 2 action:

$$\mu_{(d,x)} \in \arg\max_{\widehat{d}\in D}\left\{\widehat{m}^{\widehat{d}}_{(d,x)}\right\}.$$

We refer to the reference action's payoff as the agent's *reference point*, $r_{(d,x)}$:

$$r_{(d,x)} = \max_{\widehat{d}\in D}\left\{\widehat{m}^{\widehat{d}}_{(d,x)}\right\}.$$

With this, the emotional component of the agent's instantaneous utility (his regret) is defined as the difference between his actual material payoff and his reference point, and his *overall* instantaneous utility from action d and outcome x, $u^d(x)$, is a linear combination of his material utility and his regret:

$$u^d(x) = \phi_x + \rho\left(\phi_x - r_{(d,x)}\right),$$

where $\rho \geq 0$ measures the agent's regret concerns.[17]

Example (continued) Let $\phi = (\phi_1, \phi_2) = (\phi(0), \phi(100))$. Suppose that the agent has chosen urn α. He uses the information contained in the outcome to update his belief about the state of the urns using Bayes' rules. Since the outcome of urn α contains no information about the state of urn γ, his expected material payoff from switching to γ is just the ex ante expectation of γ, irrespective of the outcome $X^\alpha = x$:

$$\widehat{m}^{\gamma}_{(\alpha,x)} = E\left[\phi\left(\widehat{X}^{\gamma}\right)|X^\alpha = x\right] = \frac{1}{2}\phi(100) + \frac{1}{2}\phi(0).$$

In contrast, his expected material payoff from sticking with α does depend on the outcome. If the outcome is good ($x = 2$), his posterior assigns probability $\frac{1}{2} + \eta^\alpha$ to urn α being in the "good" state and accordingly the expected material payoff from choosing urn α again is

$$\widehat{m}^{\alpha}_{(\alpha,2)} = \left(\frac{1}{2} + \eta^\alpha\right)\phi(100) + \left(\frac{1}{2} - \eta^\alpha\right)\phi(0).$$

[16] If $p^d_x = 0$, we set $\widehat{m}^{\widehat{d}}_{(d,x)} = E\left[\phi\left(\widehat{X}^{\widehat{d}}\right)\right]$.

[17] Strictly, reference action, reference point, and utility depend all on the choice set D. Since we do not consider variations across choice sets, we suppress this dependency.

Since $\widehat{m}^{\alpha}_{(\alpha,2)} > \widehat{m}^{\gamma}_{(\alpha,2)}$, it follows that the agent would stick with his choice of urn α and so his reference point is given by $r_{(\alpha,100)} = \widehat{m}^{\alpha}_{(\alpha,2)}$. Thus, his instantaneous utility is

$$u^{\alpha}(2) = \phi(100) + \rho\left(\phi(100) - \widehat{m}^{\alpha}_{(\alpha,2)}\right).$$

Conversely, if the outcome is bad ($x = 1$), his posterior assigns probability $\frac{1}{2} - \eta^{\alpha}$ to urn α being in the "good" state, and accordingly the expected material payoff from choosing urn α again is

$$\widehat{m}^{\alpha}_{(\alpha,1)} = \left(\frac{1}{2} - \eta^{\alpha}\right)\phi(100) + \left(\frac{1}{2} + \eta^{\alpha}\right)\phi(0).$$

Since $\widehat{m}^{\alpha}_{(\alpha,1)} < \widehat{m}^{\gamma}_{(\alpha,1)}$, it follows that the agent would switch to urn γ and so his reference point is given by $r_{(\alpha,100)} = \widehat{m}^{\gamma}_{(\alpha,1)}$. Thus, his instantaneous utility is

$$u^{\alpha}(1) = \phi(0) + \rho\left(\phi(0) - \widehat{m}^{\gamma}_{(\alpha,1)}\right).$$

□

Behaviour We assume that the agent anticipates his emotional response and maximises his expected utility taking his regret concerns into account. Thus, the agent chooses d so as to maximise

$$\begin{aligned} U^{d} &= E\left[u\left(X^{d}\right)\right] \\ &= E\left[\phi\left(X^{d}\right)\right] + \rho\left(E\left[\phi\left(X^{d}\right)\right] - E\left[\max_{\widehat{d}\in D}\left\{\widehat{m}^{\widehat{d}}_{(X^{d},d)}\right\}\right]\right). \end{aligned}$$

Our independence assumptions imply that

$$\widehat{m}^{\widehat{d}}_{(d,x)} = \begin{cases} E\left[\phi\left(\widehat{X}^{d}\right)|X^{d} = x\right] & \text{if } \widehat{d} = d \\ \\ E\left[\phi\left(X^{\widehat{d}}\right)\right] & \text{if } \widehat{d} \neq d \end{cases}.$$

Define

$$m^{\widehat{d}} = E\left[\phi\left(X^{\widehat{d}}\right)\right]$$

to be the payoff the agent would receive if he switches action after observing the outcome of action d and

$$\widehat{m}^{d}_{x} = E\left[\phi\left(\widehat{X}^{d}\right)|X^{d} = x\right]$$

to be the payoff that the agent receives if he sticks with action d. Let $\widehat{M}^{d}$ be the real-valued lottery, which gives $\widehat{m}^{d}_{x}$ if $X^{d} = x$. Thus, we can write U^{d} as

$$U^{d} = m^{d} + \rho\left(m^{d} - E\left[\max_{\widehat{d}\in D}\left\{\widehat{M}^{d}, m^{\widehat{d}}\right\}\right]\right). \tag{3}$$

We focus on the case in which a standard expected utility maximiser ($\rho = 0$) is indifferent between all actions: $m^d = m^{\widehat{d}} = m$ for all $d, \widehat{d} \in D$. In this way, we isolate clearly the effects that are driven by the agent's regret concerns. Thus, (3) simplifies to

$$U^d = m + \rho E\left[m - \max\left\{m, \widehat{M}^d\right\}\right]. \tag{4}$$

Note that the agent's objective is concave in $\widehat{M}^d$.

Example (continued) The ex ante expected utility associated with each urn is given by

$$m^\alpha = m^\gamma = m = \frac{1}{2}\phi(100) + \frac{1}{2}\phi(0).$$

Using the calculations from above, the ex ante expected utility from choosing urn α is given by

$$\begin{aligned} U^\alpha &= \frac{1}{2}u^\alpha(2) + \frac{1}{2}u^\alpha(1) \\ &= m - \frac{1}{2}\rho\eta^\alpha(\phi(100) - \phi(0)). \end{aligned}$$

By a similar argument, the ex ante expected utility from choosing urn γ is given by

$$U^\gamma = m - \frac{1}{2}\rho\eta^\gamma(\phi(100) - \phi(0))$$

Since $\eta^\gamma < \eta^\alpha$, the agent prefers the less ambiguous urn γ. □

Remarks It is useful to think of the reference action as the action that is chosen by a counterfactual incarnation of the agent, henceforth, a *hindsight self*, who moves after the actual self and seeks to maximise the reference payoff $\widehat{m}^{\widehat{d}}_{(d,x)}$. Since the agent's objective function depends directly on his beliefs about this hindsight self's action, the agent's behaviour can be formally viewed as the equilibrium outcome of an intra-personal psychological game in the sense of Geanakoplos et al. (1989) where, in equilibrium, the agent correctly predicts his hindsight self's action. When the agent has chosen d, then the hindsight self's optimal strategy is to stick to d if $\widehat{m}^d_x > m$ and to switch action otherwise.

While our agent takes as a standard of reference what he optimally *should* have chosen, we could have alternatively assumed that he cares about what he *would* have chosen in the "counterfactual past", in which he was blessed with the wisdom of hindsight. In this case, it is appropriate to assume that the hindsight self cares about his regrets in the same way as the "actual self." Thus, the hindsight self's action depends itself on a (second order) reference action that would have been chosen by the "hindsight self's hindsight self." Extending this logic ad infinitum gives rise to an infinite intra-personal psychological game in which all the possible hindsight selves of the agent interact.[18]

Notice that we are agnostic about how the agent's prior π comes about. In the classic Ellsberg problem, the "horse"-component of the lottery reflects the agent's subjective uncertainty about the composition of the ambiguous urn. Our theory is also applicable if horse-bets represent objective risk as would be the case if, for example, the experimenter told the agent the prior probabilities of the possible urn compositions.

For our purposes, it is important only that the agent makes a fundamental conceptual distinction between his uncertainty about the composition of an urn and the uncertainty about the ball that will be drawn from an urn of a given composition. Crucially, the agent does not reduce the compound lottery to a simple probability distribution over final outcomes (even if the horse bets represent objective risk). Our specification of preferences implies that the agent blames himself for his relative ex ante ignorance about the urn's composition but not his ignorance about the outcome per se. We shall discuss the significance of this assumption after we have presented our results.

4 Independent mixtures and uncertainty aversion

In the spirit of Schmeidler (1989), we define uncertainty aversion to be a general preference for mixtures of urns over the worse of these urns. Schmeidler refers to a mixture as a *statewise* mixture, that is, action c is a mixture of actions a and b if in each state $\widetilde{\omega} \in \widetilde{\Omega}$, $\Lambda^c(\widetilde{\omega})$ is determined by playing out $\Lambda^a(\widetilde{\omega})$ with probability λ and $\Lambda^b(\widetilde{\omega})$ with probability $1-\lambda$ for some $\lambda \in [0,1]$. Intuitively, when c is a mixture of a and b, then it is less ambiguous than urn a and urn b. This is so, since the composition of urn c in each state is an average of the compositions of urn a and urn b. Therefore, the compositions of urn c across all possible states are more similar to one another than the compositions of urn a or urn b, and thus there is less uncertainty about the true composition of urn c than about the true composition of urn a or urn b.

If c is a statewise mixture of a and b, then the composition of urn c is perfectly correlated with those of urn a and urn b. Thus, a draw from c is informative about the composition of a and b, and the lottery X^c is not independent of X^a and X^b. As noted earlier, we want to abstract from such cross informational effects. Therefore, we introduce what we call an *independent mixture*.

[18] Given the similarities of the trade-offs the actual self faces in the finite and infinite game, it is likely that the two models deliver similar predictions. Indeed, in simple two-urn Ellsberg problems, we can show that strict preference for the known urn is an equilibrium of the infinite game. How this extends to the more general environments we consider in this paper, is an open question.

Definition 1 *Let $\lambda \in [0,1]$. Action c is an independent λ-mixture of actions a and b if:*
(i) $h^c_{x\omega} = \lambda h^a_{x\omega} + (1-\lambda) h^b_{x\omega}$ for all $x \in \mathcal{X}, \omega \in \Omega$.
(ii) The family $\left(X^a, X^b, X^c\right)$ is stochastically independent.

Appendix A shows that independent mixtures can always be constructed. An independent mixture preserves the averaging feature of a statewise mixture, but avoids the implication that urns are correlated. In fact, an independent mixture of a and b is equivalent to a statewise mixture of two urns, a' and b', where a' and b' are independent versions of a and b, which are perfectly correlated with each other; that is, a' is in marginal state ω if and only if b' is in marginal state ω for all $\omega \in \Omega$. Thus, choosing c amounts to choosing urn a' with probability λ and urn b' with probability $1-\lambda$.

We define uncertainty aversion with respect to independent mixtures.

Definition 2 *Let c be an independent λ-mixture of actions a and b. An agent is uncertainty averse with respect to actions a, b and c if for all $\lambda \in [0,1]$ it holds that*

$$U^c \geq \min\left\{U^a, U^b\right\}$$

and if the inequality is strict for at least one $\lambda \in (0,1)$.

Example (continued) Returning to our example, let action β be an urn with the same Markov matrix as α except with the columns reversed. (This is, of course, exactly the same urn with a different labelling of states given our independence assumption.) Thus,

$$h^\beta = \begin{pmatrix} 1+\eta^\alpha & 1-\eta^\alpha \\ 1-\eta^\alpha & 1+\eta^\alpha \end{pmatrix}$$

Notice that γ is an independent mixture of α and β since if $\lambda = \frac{1}{2}\left(1+\eta^\gamma/\eta^\alpha\right)$ then $h^\gamma = \lambda h^\alpha + (1-\lambda) h^\beta$. We have seen that the agent prefers γ to α in this example. Thus, the agent displays uncertainty aversion with respect to actions α, β and γ. We will generalise this result in section 5.2. □

5 Analysis and results

For the rest of the paper, we consider the fixed choice set $D = \{a, b, c\}$ where c is an independent λ-mixture of a and b. We provide sufficient conditions such that the regretful agent displays uncertainty aversion.

5.1 Uniform switching and statewise first order stochastic dominance

In this section, we show that uncertainty aversion arises when two conditions hold. The first condition is that the hindsight self switches action whenever his observation exceeds a uniform threshold that is independent of the agent's actual choice d. The second condition requires the first order distributions of urn a and urn b be ranked in terms of first order stochastic dominance. The conditions are formally stated in Proposition 1.[19] (The proof is in Appendix B.)

Proposition 1 *If all of the following three conditions hold, then the agent displays uncertainty aversion:*
(i) There is a threshold $x^* \in \mathcal{X}$ such that for all $d \in D : \widehat{m}_x^d > m$ if and only if $x > x^*$.
(ii) For all $\omega \in \Omega$, either h_ω^a first order stochastically dominates h_ω^b or vice versa.
(iii) There is an $\omega \in \Omega$ with $\pi_\omega > 0$ such that h_ω^a strictly first order stochastically dominates h_ω^b or vice versa.

Since the agent is indifferent between the urns in material terms, he seeks to minimise his hindsight self's payoff. Thus, he selects the urn which provides the least valuable information for the hindsight self at stage 2. In general, an urn's value of information is high if it is likely to induce an action in those states in which this action yields a high payoff.

In our context, upon observing outcome x from urn d, the hindsight self optimally sticks to urn d if the posterior expectation $\widehat{m}_x^d$ of urn d exceeds its prior expectation m. In this case, he obtains payoff $\widehat{m}_x^d$, while otherwise he chooses some other urn and obtains the same payoff (m) that he would obtain in the absence of information. We refer to an observation which induces the hindsight self to stick as "critical". Condition (i) guarantees that all urns have the same critical observations.

The value of information of urn d is therefore high if it combines, on average across states, a high likelihood of making a critical observation with a high payoff from sticking to this urn. Roughly speaking, such an urn has the property that it is good at identifying states in which this urn yields a high payoff.

In effect, by choosing urn d, the agent provides the hindsight self with a free trial draw from this urn before the latter makes his payoff-determining decision $\widehat{d}$ at stage 2. Recall that choosing a trial from urn c amounts to revealing a draw from urn a' with probability λ and a draw from urn b' with probability $1 - \lambda$, where a' and b' are perfectly correlated with each other

[19] Condition (iii) is needed only to generate a strict preference for some α.

and have the same outcome distributions h^a and h^b as have a and b. Likewise, selecting urn c in stage 2 amounts to selecting urn a' with probability λ and urn b' with probability $1-\lambda$.

Hence, if a trial draw from urn c is chosen, then conditional on the draw being critical and from urn a', the hindsight self selects urn a' with probability λ. Thus, with overall probability λ^2, the trial from c is, on average, as valuable as a trial draw from urn a. Likewise, with probability $(1-\lambda)^2$, the trial from c is as valuable as a trial draw from urn b. However, conditional on the critical draw being from urn a' (b'), he selects urn b' (a') with probability $1-\lambda$ (λ).

Mixing an observation from a' with the selection of b' and vice versa in this way could be more valuable than a trial draw from urn a and urn b, only if there are some states in which a critical observation from a' is more likely than a critical observation from b' and, at the same time, the payoff from b' in this state is larger than the payoff from a' in this state (and vice versa). Intuitively, urn a' is then good at identifying states in which urn b' gives a high payoff (and vice versa). But the latter is ruled out by condition (ii) which guarantees that in all states in which a' is more likely to generate a critical observation than b', the payoff from a' in this state is also larger than the payoff from b' in this state (and vice versa).

To see the significance of condition (ii) more clearly, it is illuminating to consider an example in which (ii) is violated. There are three equally likely (marginal) states and three outcomes. Let

$$h^a = \frac{1}{10}\begin{pmatrix} 5 & 0 & 3 \\ 5 & 5 & 0 \\ 0 & 5 & 7 \end{pmatrix}, \quad h^b = \frac{1}{10}\begin{pmatrix} 5 & 3 & 0 \\ 5 & 0 & 5 \\ 0 & 7 & 5 \end{pmatrix}, \quad h^c = \frac{1}{10}\begin{pmatrix} 5 & \frac{3}{2} & \frac{3}{2} \\ 5 & \frac{5}{2} & \frac{5}{2} \\ 0 & 6 & 6 \end{pmatrix}.$$

Thus, c is .5-micture of a and b. The utility function is $\phi = (-1, 2, 3)$. Thus, the unconditional mean is the same for each action: $m = 1.6$. The following table depicts the conditional means, $E\left[\phi\left(X^d\right)|\omega\right]$, conditional on the state:

	ω_1	ω_2	ω_3
a	.5	2.5	1.8
b	.5	1.8	2.5
c	.5	2.15	2.15

It is easy to see that only observation $x = 3$ is critical ($x^* = 2$).[20] The value of a trial observation from urn a and urn b is

$$v^a = v^b = \frac{1}{3}\left[h^a_{32} \times \left(E\left[\phi\left(X^a\right)|\omega_2\right] - m\right) + h^a_{33} \times \left(E\left[\phi\left(X^a\right)|\omega_3\right] - m\right)\right]$$
$$= \frac{1}{3}\left[.5 \times (2.5 - 1.6) + .7 \times (1.8 - 1.6)\right] = \frac{59}{300}.$$

Similarly, the value of a trial draw from urn c can be seen to be

$$v^c = \frac{66}{300}.$$

The driving force behind this result is that in state ω_2 urn a gives a higher payoff than urn b: $2.5 > 1.8$; but that b is more likely than a to generate a critical observation in this state: $h^a_{32} > h^b_{32}$ (and the reverse in state ω_3). Thus, in the event in which urn c mixes an observation from a' with selecting b', it combines the high likelihood of a critical observation from urn a' with the high payoff of urn b' in state ω_2 while urn a combines this same likelihood only with the smaller payoff from urn a in this state (and similarly for state ω_3).

Notice that conditions (i) and (ii) are automatically satisfied if the number of outcomes is $\overline{x} = 2$.[21] Thus, our result is completely general in a typical Ellsberg experiment, in which the only outcomes are winning or losing the bet.

The next lemma provides a sufficient condition for (i) in Proposition 1 in terms of the primitives h^a_ω and h^b_ω. It states monotonicity conditions, which guarantee that the posterior expectations $\widehat{m}^a_x$ and $\widehat{m}^b_x$ are increasing in x and "cut through" m at the same threshold x^*. (The proof is in Appendix B.)

Lemma 1 *The following conditions are sufficient for condition (i) in Proposition 1. For $d = a, b$:*
(i) There is a threshold $x^* \in \mathcal{X}$ such that $h^d_{x\omega}$ is strictly increasing in ω if $x > x^*$ and decreasing in ω if $x \leq x^*$.
(ii) For all $\omega > \theta$, the difference $h^d_{x\omega} - h^d_{x\theta}$ is increasing in x.

5.2 Strong Symmetry

In this section, we explore the relationship between regret and uncertainty aversion via a different route. We use the fact that the integrand

$$\zeta(z) = m - \max\{m, z\} \tag{5}$$

[20] We have for $x = 1$: $\widehat{m}^a_1 = \widehat{m}^b_1 = 79/80$, $\widehat{m}^c_1 = 179/160 \approx 1.119$; for $x = 2$: $\widehat{m}^a_2 = \widehat{m}^b_2 = 1.5$, $\widehat{m}^c_2 = 179/160$; and for $x = 3$: $\widehat{m}^a_3 = \widehat{m}^b_3 = 351/120 = 2.925$, $\widehat{m}^c_2 = 2.15$.

[21] If $\overline{x} = 2$, condition (iii) is violated in degenerate cases only.

in the objective function (4) is concave in z. Thus, action c is preferred to action a if $\widehat{M}^a$ is a mean preserving spread of $\widehat{M}^c$. Recall that $\widehat{m}^d_x$ is the realisation of $\widehat{M}^d$ if $\widehat{X}^d = x$. Denote by $\widehat{m}^d \in \mathbb{R}^{\overline{x}}$ the *row* vector

$$\widehat{m}^d = \left(\widehat{m}^d_1, ..., \widehat{m}^d_{\overline{x}}\right).$$

With this notation, the definition of a mean preserving spread is as follows.

Definition 3 *$\widehat{M}^a$ is a mean preserving spread (MPS) of $\widehat{M}^c$ if there is a $\overline{x} \times \overline{x}$ Markov matrix τ such that*
(i) $\widehat{m}^c = \widehat{m}^a \tau$,
(ii) $p^a = \tau p^c$.[22]

We show that (i) and (ii) hold under specific symmetry conditions. To state these, we begin with some definitions. Let n, k denote integers. Let δ^n be the $n \times n$ identity matrix. The matrix "slash"

$$\diagup^n = \begin{pmatrix} & & 1 \\ & \cdots & \\ 1 & & \end{pmatrix}$$

is the $n \times n$ matrix with 1's on the secondary diagonale and 0's elsewhere. When the dimension is clear from the context, we omit n. Multiplication of $\diagup$ from the right reverses the order of the columns of a matrix. Multiplication of $\diagup$ from the left reverses the order of the rows of a matrix. We call a matrix that is preserved under these two operations *strongly symmetric.*

Definition 4 *A $n \times k$ matrix h is strongly symmetric if*

$$h = \diagup^n h \diagup^k.$$

We consider matrices h^a and h^b that are strongly symmetric. In addition, we assume that h^b is the columnwise mirror image of h^a, that is, $h^b = h^a \diagup$. Finally, we require that the prior π is strongly symmetric. To avoid rather uninteresting case distinctions, we also assume that $p^a_x > 0$ for all $x \in \mathcal{X}$.[23] Then we can prove existence of a matrix τ with the desired properties. (The proof is in Appendix C).

[22] A more common definition is that there is a real-valued lottery $\widetilde{\varepsilon}$ such that (a) $\widehat{M}^a = \widehat{M}^c + \widetilde{\varepsilon}$, and (b) $E\left[\varepsilon|\widehat{M}^c\right] = 0$. To see that the definitions coincide, define for each x the lottery $\widetilde{\varepsilon}_x$ with realisations $\varepsilon_{yx} = \widehat{m}^a_y - \widehat{m}^c_x$ for $y \in \mathcal{X}$, and let τ_{yx} be the probability with which $\widetilde{\varepsilon}_x = \varepsilon_{yx}$. Let $\widetilde{\varepsilon}$ be the lottery that results from first playing out $\widehat{M}^c$ and then, conditional on $\widehat{M}^c = \widehat{m}^c_x$, playing out $\widetilde{\varepsilon}_x$. Then, (ii) implies (a), and (i) implies (b).

[23] This is actually without loss of generality. As apparent from the proof of Proposition 2, strong symmetry implies that $p^a = p^b = p^c$. Thus, we can delete zero-probability outcomes from the outcome space.

Proposition 2 *Let π and h^a be strongly symmetric. Let $h^b = h^a \diagup$. Then $\widehat{M}^a$ is an MPS of $\widehat{M}^c$. The transition matrix τ is given by*

$$\tau = \left[\lambda^2 + (1-\lambda)^2\right]\delta + 2\lambda\,(1-\lambda) \diagup$$

To show that $\widehat{M}^a$ is an MPS of $\widehat{M}^c$, we show, in fact, that the (second order) distribution of the posterior distribution of $\widehat{X}^a$, $P\left[\widehat{X}^a = \widehat{x} | X^a = x\right]$, induced by the observations from urn a is a (second order) MPS—with transition matrix τ—of the posterior distribution of $\widehat{X}^c$ that is induced by the observations from urn c. To grasp intuitively why this is, note that under strong symmetry, we have that $h^c = \lambda h^a + (1-\lambda)\, h^a \diagup = h^a \tau$. This means that the (second order) distribution over the (first order) distributions h^a_ω of urn a is a (second order) MPS—with transition matrix τ—of the (second order) distribution over the (first order) distributions h^c_ω of urn c. Thus, the possible (first order) distributions of urn a are more dispersed than those of urn c. As a consequence, observations from urn a permit sharper predictions about the next period outcomes of a compared to the predictions permitted by outcomes of c about $\widehat{X}^c$, leading to the MPS relation between the outcome posteriors of urn a and urn c.

Note also that given that urn a bears more second order risk than urn c in the sense of the mean preserving spread notion, the agent's preference for urn c is, in effect, an aversion to second order risk. In other words, under strong symmetry, uncertainty aversion and second order risk aversion are equivalent in our setup. It is beyond the scope of the current paper to investigate if our agent is second order risk averse under more general conditions.

Even though strong symmetry is, of course, a restrictive condition, it appears plausible in the classical two-action problem, in which there is one unambiguous choice and one ambiguous choice, about which nothing is known whatsoever. The principle of insufficient reason prescribes that the agent consider *any* first order distribution of the ambiguous choice as a possible state and place equal probability weight on each of them. This yields a representation of the ambiguous choice as a strongly symmetric matrix h such that the .5-mixture of h and $h\diagup$ corresponds to the unambiguous choice.

In general, $\widehat{M}^a$ being an MPS of $\widehat{M}^c$ does not guarantee a strict preference of c over a for some λ as is required in our definition of uncertainty aversion. Note first that the mean of $\widehat{M}^d$ equals m for $d = a, b$ and that the objective ζ in (5) has a kink in $z = m$. Thus, the preference of c over a is strict, if (a) $\widehat{M}^a$ is non-degenerate and if (b) $\widehat{M}^a$ is not identical to $\widehat{M}^c$. Condition (a) means that $\widehat{M}^a$ has a realisation different from its mean m, that is, there is an x such that $\widehat{m}^a_x \neq m$. Similarly, condition (b) means that there is an x such that $\widehat{m}^c_x \neq \widehat{m}^a_x$, or: $\widehat{m}^a \neq \widehat{m}^a \tau$. While both conditions are true generically[24], the following lemma provides an

[24] A slight perturbation of the entries of h^a changes $\widehat{m}^a_x$ and turns the equalities into inequalities. Similarly, since $\widehat{m}^a_x$ depends on ϕ, conditions (a) and (b) are true for generic ϕ.

analytical sufficient condition for (b). (The proof is in Appendix C).

Lemma 2 *Under the conditions of Proposition 2, let* $\widehat{M}^a$ *be non-degenerate. Let* $\overline{\omega} \geq \overline{x}$*, let* h^a *have rank* $\overline{x}$*, and let* $\pi_\omega > 0$ *for all* ω*. Then the preference of* c *over* a *is strict for all* $\lambda \in (0, 1)$*.*

5.3 Robustness

Our results rest critically on how the extensive form of the psychological game that represents the agent's decision problem is specified. If, contrary to what we have assumed, Nature selects *both* the composition and draw from each urn once and for all at the beginning of the game, then the counterfactual self learns the draw that he will receive from the urn chosen by the actual self. The agent would then compare his actual payoff to the maximum expected payoff he could obtain if he could choose again and obtain the exact same draw from the urn that he has chosen.

Making this alternative assumption eliminates any impact of ambiguity on the agent's decision making, since choosing any urn reveals the same type of outcome information about the urn to the hindsight self, regardless of its degree of uncertainty. So the only feature of an urn that is relevant from an ex ante perspective is an urn's probability distribution over final outcomes.[25]

As noted in Section 2, the agent may be driven to make the necessary distinction between the two types of uncertainty by a psychological tendency to distinguish between facts that are knowable in principle and facts that are inherently unpredictable. If so, then it is possible that whether or not the agent makes the necessary distinction may be a contingent matter that depends on the manner in which problems are presented. For example, if the agent is told that the ball from each of the urns has been selected in advance then he might perceive that the colour of the ball is now an event that is, in principle, knowable (since the experimenter, for example, could inspect each ball that has been drawn from the urn). In terms of the extensive form game, the agent would envisage that Nature determines both the composition and draw from each urn at the beginning of the game, and so his aversion to ambiguity would disappear.[26] Thus, if this psychological hypothesis is true, then our model predicts that ambiguity aversion will be subject to framing effects.

[25] Similarly, uncertainty aversion is eliminated by assuming that Nature independently selects the composition of the urn again after the acutal self's and before the hindsight self's move.

[26] See Brun and Teigen (1990) for evidence that is supportive of this prediction.

6 Relation to the literature

Theoretical accounts of the Ellsberg paradox relax the consequentialist premise that the decision maker is always indifferent between a compound lottery and the corresponding simple lottery that gives the same probability distribution over final outcomes. Contrary to subjective expected utility theory they assume that the decision maker does not reduce the uncertainty about the true probabilities that he faces to ordinary risk. Therefore, a fundamental distinction between first order and second order probabilities arises where second order probabilities capture a decision maker's uncertainty about the probabilities he faces and first order probabilities capture his exposure to ordinary risk. Uncertainty aversion arises when an agent dislikes uncertainty about these second order probabilities. Thus, a common aim of the various theoretical accounts of Ellsberg behaviour is to characterise preference specifications that exhibit uncertainty aversion and to understand the relationship between uncertainty aversion and (some notion of) second-order risk aversion.

We contribute to this literature by identifying a psychological reason (regret) why an agent might directly care about the uncertainty about the probabilities he faces and, in turn, display uncertainty aversion.

Insofar as we seek a behavioural foundation for uncertainty aversion, the approach closest to ours is Halevy and Feltkamp (2005), henceforth HF, who argue that uncertainty averse behaviour arises from the conjunction of an agent's ordinary risk aversion with his adherence to a rule of thumb that may be appropriate in many (but not all) ambiguous situations. Specifically, they argue that in many situations of ambiguity, ordinary first order risk aversion will lead an agent to avoid ambiguous gambles. This is because the distribution of overall payoffs that is generated by a bundle of identical ambiguous gambles (e.g. a series of draws with replacement from the unknown urn) is more risky than that which is generated by a bundle of equivalent unambiguous gambles (e.g. a series of draws with replacement from the known urn).[27] Thus, if a decision maker ordinarily faces bundled risks, then avoiding ambiguous gambles is an effective rule of thumb, albeit one which leads him to make errors in scenarios, such as the Ellsberg experiments, in which he faces a choice between gambles which are played only once.[28]

HF obtain uncertainty aversion under the FOSD ranking condition (ii) of Proposition 1. In

[27] Whereas successive realisations of the unambiguous gambles are independent, realisations of the ambiguous gambles are correlated with one another due to their common dependence on the unknown state. This correlation increases the overall risk to which the decision maker is exposed.

[28] Morris (1997) explains ambiguity aversion as arising from the application of a different rule of thumb according to which the decision maker perceives the choice of an ambiguous urn as a competitive bet on the urn's composition against a privately informed experimenter. Avoidance of ambiguous urns then arises from adverse selection in the style of no-trade results.

fact, the driving force behind our and HF's result is similar. When HF's agent prefers urn a over urn b, then he does so because the first and the second draw from urn a are more correlated than the first and the second draw from urn b. But this implies that the first draw from urn a is a better predictor for the second draw from urn a than is the first draw from urn b for the second draw from urn b. Thus, observing the first draw from urn a provides more useful information for a second (hypothetical) decision than does observing the first draw from urn b. This intuitively suggests that our information averse agent would also prefer urn b over urn a. Our results show that this intuition is formally true if in addition to (ii), condition (i) of Proposition 1 is imposed.

Like HF, our theory is consistent with the experimental results of Halevy (2005) which shows that there is a negative correlation between the tendency to reduce *objective* compound lotteries and ambiguity aversion: subjects who reduce such lotteries tend to be ambiguity neutral, while subjects who fail to reduce such lotteries tend to be ambiguity averse. In line with our theory which is agnostic about how the agent's second order probabilities are determined, this suggests that ambiguity aversion is associated with a more general tendency to fail to reduce compound lotteries rather than the more specific tendency to fail to perform such a reduction only when one set of probabilities is subjective.

Despite these apparent similarities between our and HF's approach, the implications are different in at least two respects. First, it is easy to construct examples that satisfy our strong symmetry conditions of Proposition 2, for which the HF framework is without bite. We present such an example in Appendix D.

Second, and more importantly, there are testable implications that allow us to experimentally discriminate between the two models. One implication of our approach that is absent from HF is that our agent's behaviour depends on the nature of the feedback he expects to receive on his actual as well as his forgone choice. In the formal analysis, we have focussed on the case in which the agent only learns the outcome of the urn he has chosen. If, instead, he observes draws from both urns irrespective of his choice, our model predicts that the agent no longer displays uncertainty aversion, because his actions have no differential impact on his hindsight knowledge. Thus, we would expect experimental subjects' tendency to select the unambiguous urn to be mitigated when they are told in advance that they will receive feedback on the composition of both urns regardless of the choice that they make.[29]

[29] Curley et al. (1986) conduct an experiment in which they indeed manipulate subject's expected feedback by revealing the contents of the ambiguous urn ex post. They claim to rule out regret as a source of ambiguity aversion, since the data rejects their hypothesis that this feedback manipulation *increases* ambiguity aversion. However, Curley et al's findings are actually consistent with our model, which predicts that such forced feedback will *reduce* ambiguity avoiding behaviour.

Other theoretical accounts of ambiguity aversion can be broadly classified into two groups according to how they treat (second-order) uncertainty. In the first group, uncertainty is not treated as a probability. Segal (1987) considers an agent who uses a non-linear probability weighting function to evaluate the second-order uncertainty to which he is exposed while he evaluates first order risk in the manner of an expected utility maximiser. He derives conditions on the probability weighting function that deliver ambiguity aversion and shows that they are similar to the conditions which generate risk aversion in anticipated utility theory.[30] On a related note, Schmeidler (1989) provides an axiomatic foundation of preferences that permit an anticipated utility representation. His agent evaluates uncertainty using a non-additive measure rather than probabilities.

The second group conceives of uncertainty in terms of second order probabilities. A number of papers provide axiomatic foundations for preference representations that yield uncertainty aversion. Klibanoff et al. (2005) derive a recursive expected utility representation of uncertainty aversion, where a compound lottery is evaluated by a two-stage procedure: first, the (ordinary) expected utility of each realisation p (which is a simple lottery) of the compound lottery is computed ($\int u(x)\, p(dx)$). Second, an expected utility over all possible of these (ordinary) expected utilities is taken with respect to a subjective probability measure μ defined on the set of all simple lotteries ($\int \phi\left(\int u(x)\, p(dx)\right) \mu(dp)$). Klibanoff et al. show that ambiguity aversion, an aversion to second order mean preserving spreads, and ambiguity attitude can be characterised in terms of the shape of ϕ in much the same way as in treatments of ordinary risk aversion. Ergin and Gul (2002), Ahn (2005) and Nau (2005) provide alternative axiomatic foundations of similar recursive expected utility formulations. The maximin expected utility approach of Gilboa and Schmeidler (1988) can be viewed as an extreme case of Klibanoff et al.'s formulation in which ϕ places all the weight on the simple lottery that generates the lowest expected utility according to u.

7 Conclusion

This paper represents an attempt to gain a better understanding of the psychological forces that drive ambiguity aversion. We have proposed a model of regret that can account for Ellsberg type behaviour in choice under uncertainty. We have argued that ambiguity aversion arises from an agent's aversion to discovering that his choices are supoptimal from an ex post perspective.

We view our paper as forming part of a broader research agenda that seeks to find emotional underpinnings of behavioural phenomena. In general, this is a worthwhile enterprise if the

[30] Segal (1990) derives his 1987 model from preferences over two-stage lotteries.

resulting psychological models generate novel predictions about the phenomena they seek to explain. Our model of regret meets this test by suggesting that a kind of information aversion accounts for ambiguity aversion under certain conditions. Accordingly, it identifies situations in which agents will be more or less likely to make ambiguity avoiding choices. According to our model, agents will display greater ambiguity aversion when conditions are such that choosing the more ambiguous option reveals more information about the (first order) distribution of the chosen option. In particular, as we have previously argued, uniquely among existing theories of ambiguity aversion our model predicts that an agent's uncertainty aversion will be mitigated when he is exposed to the same feedback about the (first order) distributions of options in his choice set regardless of the option that he chooses.

Appendix A: Construction of independent mixtures

Consider a setup with two actions a and b. That is, $\Omega^a = \Omega^b = \Omega$, $\widetilde{\Omega} = \Omega^a \times \Omega^b$, and equs. (1) and (2) hold. We construct an independent λ-mixture of a and b.

Let $\Omega^c = \Omega$ be a new state space component and let $\widetilde{\widetilde{\Omega}} = \widetilde{\Omega} \times \Omega^c$ be the new state space. We extend $\widetilde{\pi}$ and $\Lambda^d, d = a, b$ on $\widetilde{\widetilde{\Omega}}$ by defining

$$\widetilde{\widetilde{\pi}}\left(\omega^a, \omega^b, \omega^c\right) = \pi\left(\omega^a\right)\pi\left(\omega^a\right)\pi\left(\omega^c\right),$$
$$\widetilde{\Lambda}^d\left(\omega^a, \omega^b, \omega^c\right) = \Lambda^d\left(\omega^a, \omega^b\right).$$

We now define $\widetilde{\Lambda}^c : \widetilde{\widetilde{\Omega}} \to L$. We begin on the "diagonale". For $\omega^a = \omega^b = \omega^c$, let

$$\widetilde{\Lambda}^c\left(\omega^a, \omega^b, \omega^c\right) = \lambda \circ \widetilde{\Lambda}^a\left(\omega^a, \omega^b, \omega^c\right) \oplus (1-\lambda) \circ \widetilde{\Lambda}^b\left(\omega^a, \omega^b, \omega^c\right)$$

be the compound lottery that is obtained by playing out $\widetilde{\Lambda}^a\left(\omega^a, \omega^b, \omega^c\right)$ with probability λ and playing out $\widetilde{\Lambda}^b\left(\omega^a, \omega^b, \omega^c\right)$ with probability $1-\lambda$. For "off-diagonale" states, we define $\widetilde{\Lambda}^c$ as independent from ω^a and ω^b. That is, for $\omega^c \neq \omega^a, \omega^b$, let

$$\widetilde{\Lambda}^c\left(\omega^a, \omega^b, \omega^c\right) = \widetilde{\Lambda}^c\left(\omega^c, \omega^c, \omega^c\right).$$

Thus, the probability with which $\widetilde{\Lambda}^c\left(\omega^a, \omega^b, \omega^c\right) = x$ in state $\omega^c \in \Omega$ is

$$h^c_{x\omega^c} = \lambda h^a_{x\omega^c} + (1-\lambda)\, h^b_{x\omega^c}.$$

In addition, the family $\left(X^a, X^b, X^c\right)$ is stochastically independent by construction. Thus c is an independent λ-mixture of a and b. $\square$

Appendix B: Proofs of Section 5.1

Proof of Proposition 1 Without loss of generality, suppose $U^a \geq U^b$. We have to show that $U^c \geq U^b$ for all $\lambda \in [0,1]$ with strict inequality for one $\lambda \in (0,1)$. Note first that the material term m in (4) is the same for all actions. Thus, only differences in the regret term $V^d = E\left[m - \max\left\{m, \widehat{M}^d\right\}\right]$ matter. Thus, $V^a \geq V^b$, and we need to show that $V^c \geq V^b$ for all $\lambda \in [0,1]$ with strict inequality for one $\lambda \in (0,1)$. To this end, we first re-write V^d for all $d \in D$ by conditioning on observations $X^d = x$. Note that by (i), for all $d \in D$, we have $\widehat{M}^d > m$ if and only if $X^d > x^*$. Thus,

$$V^d = \sum_{x=1}^{x^*} P\left[X^d = x\right](m - m) + \sum_{x=x^*+1}^{\overline{x}} P\left[X^d = x\right]\left(m - E\left[\phi\left(\widehat{X}^d\right)|X^d = x\right]\right).$$

The first term on the right hand side dissapears. The second term depends only on marginal states ω^d. Thus, since ω^d is independent of $\omega^{\widehat{d}}$ for all $\widehat{d} \neq d$, we can sum over the conditional expectation, conditional on $\omega^d \in \Omega^d$. Hence, we obtain (by dropping the index d from ω^d)

$$\begin{aligned} V^d &= \sum_{\omega=1}^{\overline{\omega}} \pi_\omega \sum_{x=x^*+1}^{\overline{x}} P\left[X^d = x|\omega\right]\left(m - E\left[\phi\left(\widehat{X}^d\right)|X^d = x, \omega\right]\right) \\ &= \sum_{\omega=1}^{\overline{\omega}} \pi_\omega \sum_{x=x^*+1}^{\overline{x}} h^d_{x\omega}\left(m - E\left[\phi\left(\widehat{X}^d\right)|\omega\right]\right), \end{aligned}$$

where, in the second line, we have used that $P\left[X^d = x|\omega\right] = h^d_{x\omega}$, and the fact that $E\left[\phi\left(\widehat{X}^d\right)|X^d = x, \omega\right] = E\left[\phi\left(\widehat{X}^d\right)|\omega\right]$. Now, define $t^d_\omega = m - E\left[\phi\left(\widehat{X}^d\right)|\omega\right]$, and $s^d_\omega = \sum_{x=x^*+1}^{\overline{x}} h^d_{x\omega}$. Thus,

$$V^d = \sum_{\omega=1}^{\overline{\omega}} \pi_\omega s^d_\omega t^d_\omega.$$

Note that $t^d_\omega = \sum_{x=1}^{\overline{x}} h^d_{x\omega}(m - \phi_x)$. Therefore, since c is a λ-mixture of a and b:

$$s^c_\omega = \lambda s^a_\omega + (1-\lambda)s^b_\omega, \quad \text{and} \quad t^c_\omega = \lambda t^a_\omega + (1-\lambda)t^b_\omega.$$

Using this in V^c and multiplying out, we obtain:

$$\begin{aligned} V^c &= \sum_{\omega=1}^{\overline{\omega}} \pi_\omega\left(\lambda^2 s^a_\omega t^a_\omega + (1-\lambda)^2 s^b_\omega t^b_\omega + \lambda(1-\lambda)\left(s^a_\omega t^b_\omega + s^b_\omega t^a_\omega\right)\right) \\ &= \lambda^2 V^a + (1-\lambda)^2 V^b + \lambda(1-\lambda)\sum_{\omega=1}^{\overline{\omega}} \pi_\omega\left(s^a_\omega t^b_\omega + s^b_\omega t^a_\omega\right). \end{aligned} \tag{6}$$

We now show that $V^c \geq V^b$. Indeed, since $V^a \geq V^b$ by assumption, we can estimate the first two terms in (6) against $\left(\lambda^2 + (1-\lambda)^2\right)V^b$, and it follows that $V^c \geq V^b$ if the last term in (6) is larger than $V^a + V^b \geq 2V^b$; that is,

$$\sum_{\omega=1}^{\overline{\omega}} \pi_\omega \left(s_\omega^a t_\omega^b + s_\omega^b t_\omega^a \right) \geq \sum_{\omega=1}^{\overline{\omega}} \pi_\omega \left(s_\omega^a t_\omega^a + s_\omega^b t_\omega^b \right),$$

which, by subtracting the right from the left hand side, is equivalent to

$$\sum_{\omega=1}^{\overline{\omega}} \pi_\omega \left(s_\omega^a - s_\omega^b \right) \left(t_\omega^b - t_\omega^a \right) \geq 0. \tag{7}$$

Now suppose that h_ω^a first order stochastically dominates h_ω^b. Since ϕ_x is increasing in x, this implies that $t_\omega^a \leq t_\omega^b$. Observe that we can write $s_\omega^d = \sum_{x=1}^{\overline{x}} h_{x\omega}^d \psi_x$ with the increasing function $\psi_x = 1_{\{x>x^*\}}(x)$. Hence, it is also true that $s_\omega^a \geq s_\omega^b$. If h_ω^b first order stochastically dominates h_ω^a, the argument is reversed. Thus, with (ii), we have for all ω:

$$s_\omega^a \geq s_\omega^b \Longleftrightarrow t_\omega^a \leq t_\omega^b.$$

But this implies (7).

To complete the proof, it remains to show that $V^c > V^b$ for some $\lambda \in (0,1)$. But this follows from (iii), which ensures that at least one term under the sum in (7) is strictly positive. This establishes the claim. □

Proof of Lemma 1 We proceed in two steps. In Step 1, we show that (i) implies

$$E\left[\phi\left(\widehat{X}^d\right)|X^d = x\right] = \widehat{m}_x^d > m \Longleftrightarrow x > x^*. \tag{8}$$

In Step 2, we show that if (i) and (ii) hold for $d = a, b$, then they hold for $d = c$. Thus, steps 1 and 2 imply the first condition in Proposition 1.

As for Step 1. We suppress the action index d. We begin by defining the agent's posterior: for a fixed outcome x with $p_x > 0$, Bayes' rule yields that

$$\widehat{\pi}_{\omega x} = \frac{h_{x\omega}\pi_\omega}{p_x} \tag{9}$$

is the probability with which the state is ω given $X = x$. For $p_x = 0$, we set $\widehat{\pi}_{\omega x} = \pi_\omega$.[31] Let $\widehat{\pi}_x \in \Delta(\Omega)$ be given by the *column* vector

$$\widehat{\pi}_x = \left(\widehat{\pi}_{1x}, ..., \widehat{\pi}_{\overline{\omega}x}\right).$$

With this definition, we proceed in two steps. In step (A), we show that if $h_{x\omega}$ is strictly increasing in ω, then the posterior $\widehat{\pi}_x$ strictly first order stochastically dominates the prior π. In step (B), we show that (c) implies that the function $\Psi : \Omega \to \mathbb{R}$ given by

$$\Psi_\omega = E\left[\phi\left(\widehat{X}\right)|\omega\right]$$

[31] This is consistent with our convention that $\widehat{m}_x = m$ for $p_x = 0$.

is increasing in ω. Thus, it follows by the strict dominance relation established in (A) that

$$\sum_{\omega} \Psi_\omega \widehat{\pi}_{\omega x} > \sum_{\omega} \Psi_\omega \pi_\omega$$

if $h_{x\omega}$ is strictly increasing in ω. But note that by the law of iterated expectation, the left hand side is equal to $\widehat{m}_x^d$, and the right hand side is equal to m. If $h_{x\omega}$ is decreasing in ω, an identical argument applies with all inequalities being weak and reversed. Thus, (i) implies (8).

As for (A). The argument is the same as in Milgrom (1981), proof of Proposition 1.

As for (B). By definition,

$$\Psi_\omega = \sum_{x=1}^{\overline{x}} \phi_x h_{x\omega}.$$

Let $\omega > \theta$. We show that $\Psi_\omega - \Psi_\theta \geq 0$. To see this, observe first that since $h_{x\omega} - h_{x\theta}$ is increasing in x by assumption (ii), there is a unique $\widetilde{x} \in \mathcal{X}$ such that $h_{x\omega} - h_{x\theta} \leq 0$ if and only if $x \leq \widetilde{x}$. Hence, since ϕ_x is increasing in x, it follows for all x:

$$\phi_x \left(h_{x\omega} - h_{x\theta}\right) \geq \phi_{\widetilde{x}} \left(h_{x\omega} - h_{x\theta}\right).$$

Observe second that $\sum_{x=1}^{\overline{x}} \left(h_{x\omega} - h_{x\theta}\right) = 0$. Together, the two observations imply that

$$\Psi_\omega - \Psi_\theta \geq \phi_{\widetilde{x}} \sum_{x=1}^{\overline{x}} \left(h_{x\omega} - h_{x\theta}\right) = 0.$$

This completes Step 1.

As for Step 2. Note that since $h_{x\omega}^c$ is the convex combination of $h_{x\omega}^a$ and $h_{x\omega}^b$, monotonicity properties carry over to $h_{x\omega}^c$. This completes the proof. $\square$

Appendix C: Proofs of Section 5.2

To prove Proposition 2, we introduce matrix notation. We can write the distribution of X^d as

$$p^d = h^d \pi^d,$$

The expected material payoff of action d is

$$m^d = \phi^T p^d,$$

where ϕ^T is the transposed of ϕ. Recall the definition of posteriors from (9) and re-attach the action index d. We summarise all possible posterior beliefs in the $\overline{\omega} \times \overline{x}$ Markov matrix

$$\widehat{\pi}^d = \left(\widehat{\pi}_{\omega x}^d\right)_{\omega,x}. \tag{10}$$

Thus, conditional on having observed $X^d = x$, the (posterior) probability with which $\widehat{X}^d = \widehat{x}$ is given by

$$\widehat{p}^d_{\widehat{x}x} = \sum_{\omega \in \Omega} h^d_{\widehat{x}\omega} \widehat{\pi}^d_{\omega x}.$$

We denote the conditional distribution of $\widehat{X}^d$, conditional on $X^d = x$, by the vector $\widehat{p}^d_x \in \Delta(\mathcal{X})$ and summarise all possible such distributions in the $\overline{x} \times \overline{x}$ Markov matrix

$$\widehat{p}^d = \left(\widehat{p}^d_1, ..., \widehat{p}^d_{\overline{x}}\right) = h^d \widehat{\pi}^d.$$

With this, we can write

$$\widehat{m}^d_x = \phi^T \widehat{p}^d_x.$$

For $v \in \mathbb{R}^n$, define δ_v as the $n \times n$ matrix with diagonale elements $(\delta_v)_{ii} = v_i$ for $i = 1, ..., n$ and off-diagonale elements 0. Denote by $S_{n,k}$ the set of all strongly symmetric matrices. We require the following lemma (proven below).

Lemma 3 *Under the conditions of Proposition 2:*
(a) $p^a \in S_{\overline{x},1}$
(b) $/\widehat{\pi}^a = \widehat{\pi}^a /$
(c) $p^a = p^b = p^c$.
(d) $\widehat{\pi}^b = /\widehat{\pi}^a$.

Proof of Proposition 2 We begin by noting that condition (i) from Definition 3 is implied by the condition

$$\widehat{p}^c = \widehat{p}^a \tau \qquad \text{(i')}$$

To see this, recall that $\widehat{m}^d_x = \phi^T \widehat{p}^d_x$. Thus, $\widehat{m}^d = \phi^T \widehat{p}^d$, and (i') implies

$$\widehat{m}^c = \phi^T \widehat{p}^c = \phi^T \widehat{p}^a \tau = \widehat{m}^a \tau.$$

We now show (i') and condition (ii) from Definition 3.

As for (i'). In light of Lemma 3 (c), we define $p = p^a$. By equ. (9), we have for all $d \in D$: $\widehat{\pi}^d_{\omega x} = \left(h^d_{x\omega}\pi_\omega\right)/p_x$ if $p_x > 0$. Since by assumption $p_x > 0$ for all $x \in \mathcal{X}$, (c) implies that we can write $\widehat{\pi}^d$ as

$$\widehat{\pi}^d = \delta_\pi \left(h^d\right)^T \delta_p^{-1}.$$

Thus, with $h^c = \lambda h^a + (1 - \lambda) h^b$:

$$\widehat{\pi}^c = \delta_\pi \left(\lambda h^a + (1 - \lambda) h^b\right)^T \delta_p^{-1} = \lambda \widehat{\pi}^a + (1 - \lambda) \widehat{\pi}^b.$$

Using this in $\widehat{p}^c = h^c \widehat{\pi}^c$ and multiplying out, we obtain:

$$\widehat{p}^c = \lambda^2 h^a \widehat{\pi}^a + (1 - \lambda)^2 h^b \widehat{\pi}^b + \lambda (1 - \lambda) \left(h^a \widehat{\pi}^b + h^b \widehat{\pi}^a\right).$$

The claim follows now from two observations. First, note that $\diagup\diagup = \delta$. Hence, since $h^b = h^a \diagup$ by assumption, and since $\widehat{\pi}^b = \diagup \widehat{\pi}^a$ by (d), it follows: $h^b \widehat{\pi}^b = h^a \diagup\diagup \widehat{\pi}^a = h^a \widehat{\pi}^a$. Second, since $h^b = h^a \diagup$, since $\widehat{\pi}^b = \diagup \widehat{\pi}^a$, and since $\diagup \widehat{\pi}^a = \widehat{\pi}^a \diagup$ by (b), we can deduce that

$$h^a \widehat{\pi}^b + h^b \widehat{\pi}^a = h^a \diagup \widehat{\pi}^a + h^a \diagup \widehat{\pi}^a = 2 h^a \widehat{\pi}^a \diagup .$$

Therefore, with $\widehat{p}^a = h^a \widehat{\pi}^a$ we conclude that

$$\begin{aligned} \widehat{p}^c &= \left[\lambda^2 + (1-\lambda)^2\right] \widehat{p}^a + 2\lambda (1-\lambda) \widehat{p}^a \diagup \\ &= \widehat{p}^a \tau , \end{aligned}$$

where $\tau = \left[\lambda^2 + (1-\lambda)^2\right] \delta + 2\lambda (1-\lambda) \diagup$. Notice that τ is trivially Markov, and this establishes (i').

As for (ii). By Lemma 3 (a), $p^a \in S_{\overline{x},1}$. Thus $\diagup p^a = p^a$, and it follows that $\tau p^a = p^a$. By (c), $p^c = p^a$, thus $\tau p^c = p^a$, and this completes the proof. □

Proof of Lemma 3 As for (a). Note that $\diagup\diagup = \delta$. Thus, $\diagup h^a \pi^a \diagup = \diagup h^a \diagup\diagup \pi^a \diagup = h^a \pi^a$. The last step follows because h and π are strongly symmetric by assumption. But this means that $p^a = h^a \pi^a \in S_{\overline{x},1}$.

As for (b). We will use the obvious property: (A) $v \in S_{n,1} \Longleftrightarrow \delta_v = \diagup \delta_v \diagup$.
Now, recall from (9) that $\widehat{\pi}^a_{\omega x} = \left(h^a_{x\omega} \pi_\omega\right) / p^a_x$ if $p^a_x > 0$. Since by assumption $p^a_x > 0$ for all $x \in \mathcal{X}$, we can therefore write $\widehat{\pi}^a$ as

$$\widehat{\pi}^a = \delta_\pi \left(h^a\right)^T \delta^{-1}_{p^a} .$$

Since $p^a \in S_{\overline{x},1}$ by (a), (A) implies that $\delta^{-1}_{p^a} = \diagup \delta^{-1}_{p^a} \diagup$. Likewise, $\delta_\pi = \diagup \delta_\pi \diagup$, as $\pi \in S_{\overline{\omega},1}$. Hence,

$$\diagup \widehat{\pi}^a \diagup = \diagup \delta_\pi \left(h^a\right)^T \delta^{-1}_{p^a} \diagup = \diagup\diagup \delta_\pi \diagup \left(h^a\right)^T \diagup \delta^{-1}_{p^a} \diagup\diagup .$$

Now note that $\diagup^T = \diagup$. Thus, $\diagup \left(h^a\right)^T \diagup = \left(\diagup h^a \diagup\right)^T = \left(h^a\right)^T$, where the last equality is true since h^a is strongly symmetric. With this and $\diagup\diagup = \delta$, we deduce that $\diagup \widehat{\pi}^a \diagup = \delta_\pi \left(h^a\right)^T \delta^{-1}_{p^a} = \widehat{\pi}^a$. Thus, $\diagup \widehat{\pi}^a = \diagup\diagup \widehat{\pi}^a \diagup = \widehat{\pi}^a \diagup$, which is (b).

As for (c). Since $\pi \in S_{\overline{\omega},1}$ it follows that $\diagup \pi = \pi$. Thus, as $h^b = h^a \diagup$ by assumption:

$$p^b = h^b \pi = h^a \diagup \pi = h^a \pi = p^a .$$

Moreover, $p^c = h^c \pi = \lambda h^a \pi + (1-\lambda) h^b \pi = p^a$, and (c) is established.

As for (d). By (c), $\delta^{-1}_{p^b} = \delta^{-1}_{p^a}$ is well-defined and with $h^b = h^a \diagup$ we obtain $\widehat{\pi}^b = \delta_\pi \left(h^a \diagup\right)^T \delta^{-1}_{p^a}$. Note that $\diagup^T = \diagup$. Thus $\left(h^a \diagup\right)^T = \diagup \left(h^a\right)^T$, and with $\diagup \delta_\pi = \delta_\pi \diagup$, we arrive at

$$\widehat{\pi}^b = \diagup \delta_\pi \left(h^a\right)^T \delta^{-1}_{p^a} = \diagup \widehat{\pi}^a ,$$

which is (d). This completes the proof. □

Proof of Lemma 2 We have to show that $\widehat{m}^a \neq \widehat{m}^a\tau$. We drop the action index. With $\tau = \left[\lambda^2 + (1-\lambda)^2\right]\delta + 2\lambda(1-\lambda)/$, it follows that $\widehat{m} = \widehat{m}\tau$ if and only if $\widehat{m} - \widehat{m}/ = 0$ and $\lambda \in (0,1)$. Hence, we have to show that $\widehat{m} - \widehat{m}/ \neq 0$. Recall that $\widehat{m} = \phi^T\widehat{p} = \phi^T h\widehat{\pi}$. Thus, from Lemma 3 (b) and with $h/ = /h$, we obtain:

$$\widehat{m} - \widehat{m}/ = \phi^T h\widehat{\pi} - \phi^T / h\widehat{\pi} = \left(\phi^T - \phi^T/\right) h\widehat{\pi}.$$

Note that $\widehat{\pi} = \delta_\pi h^T \delta_p^{-1}$. Since $\pi_\omega > 0$ for all ω by assumption, δ_π has rank $\overline{\omega}$. Likewise, δ_p^{-1} has rank $\overline{x}$. Hence, since h has rank $\overline{x}$ by assumption, it follows that $\widehat{\pi}$ has rank $\min\{\overline{x}, \overline{\omega}\} = \overline{x}$, and consequently, the $\overline{x} \times \overline{x}$ matrix $h\widehat{\pi}$ has rank $\overline{x}$. Therefore, $\left(\phi^T - \phi^T/\right) h\widehat{\pi} = 0$ if and only if $\phi^T - \phi^T/ = 0$. But since ϕ is increasing by assumption, the latter cannot be true. Thus, $\widehat{m} - \widehat{m}/ \neq 0$, and the claim is shown. □

Appendix D: Example distinguishing our model from HF's model

The following example illustrates that there are settings, in which our model generates uncertainty aversion while HF's model does not. There are two equally likely (marginal) states and five outcomes. We set

$$h^a = \frac{1}{10}\begin{pmatrix} 1.0 & 1.5 \\ 3.0 & 0.5 \\ 4.0 & 4.0 \\ 0.5 & 3.0 \\ 1.5 & 1.0 \end{pmatrix}, \quad h^b = \frac{1}{10}\begin{pmatrix} 1.5 & 1.0 \\ 0.5 & 3.0 \\ 4.0 & 4.0 \\ 3.0 & 0.5 \\ 1.0 & 1.5 \end{pmatrix}, \quad h^c = \frac{1}{10}\begin{pmatrix} 1.25 & 1.25 \\ 1.75 & 1.75 \\ 4.00 & 4.00 \\ 1.75 & 1.75 \\ 1.25 & 1.25 \end{pmatrix}.$$

Actions a and b clearly satisfy the strong symmetry conditions of Proposition 2 and c is a .5-mixture of a and b. Thus, in a pairwise choice between c and either a or b our agent will certainly prefer urn c to either a or b. However, actions a and b do not satisfy HF's requirement that corresponding states of a and b be ranked by first order stochastic dominance.

Denote by Y^d the sum of outcomes when d is chosen twice and played out conditionally independently. (In the notation of our model: $Y^d = X^d + \widehat{X}^d$.) According to HF's model, in a choice between action a and action c, any agent with a concave utility function will prefer the mixture c to a if and only if Y^c second order stochastically dominates Y^a, which is equivalent to

$$\sum_{y=2}^{\widetilde{y}} P\left[Y^a \leq y\right] \leq \sum_{y=2}^{\widetilde{y}} P\left[Y^c \leq y\right] \quad \text{for all } \widetilde{y} \in \{2, ..., 10\}.$$

Now notice that $P[Y^a = 2] = (1/2) \times (1/10)^2 + (1/2) \times (15/100)^2 = 0.01625$ and $P[Y^c = 2] = (125/100)^2 = 0.015625$. Thus,

$$\sum_{y=2}^{2} P[Y^a \leq y] = 0.01625 > 0.015625 = \sum_{y=2}^{2} P[Y^c \leq y].$$

Similarly, we obtain that

$$\sum_{y=2}^{3} P[Y^a \leq y] = 0.07000 < 0.075000 = \sum_{y=2}^{3} P[Y^c \leq y].$$

Thus, Y^c and Y^a are not ranked by second order stochastic dominance, so it is not true that any risk-averse agent will necessarily prefer c to a. □

References

Ahn, D. S. (2005). "Ambiguity without state space." mimeo.

Anscombe, F. J. and R. J. Aumann (1963). "A definition of subjective probability." *Annals of Mathematical Statistics*, 34: 199-205.

Bell, D. E. (1982). "Regret in decision making under uncertainty." *Operations Research*, 30: 961-981.

Bell, D. E. (1983). "Risk premiums for decision regret." *Management Science*, 29: 1156-1166.

Brun, W. and K. H. Teigen (1990). "Prediction and postdiction preferences in guessing." *Journal of Behavioral Decision Making*, 3: 17-28.

Curley, S. P., F. Yates and R. A. Abrams (1986). "Psychological Sources of Ambiguity Avoidance." *Organizational Behavior and Human Decision Processes*, 38: 230-256.

Ellsberg, D. (1961) "Risk, ambiguity and the Savage axioms." *Quarterly Journal of Economics*, 75: 643-669.

Ergin, H. and F. Gul (2002). "A subjective theory of compound lotteries." mimeo.

Frisch, D., and J. Baron (1988). "Ambiguity and rationality." *Journal of Behavioral Decision Making*, 1: 149-157.

Fischoff, B. (1975). "Hindsight ≠ foresight: The effect of outcome knowledge on judgment." *Journal of Experimental Psychology*, 1: 288-299.

Geanakoplos, J., D. Pearce, and E. Stacchetti (1989). "Psychological games and sequential rationality." *Games and Economic Behavior*, 1: 60-79.

Gilboa, I. and D. Schmeidler. "Maxmin expected utility with a non-unique prior." *Journal of Mathematical Economics*, 18: 141-153.

Gilovich, T. and V. H. Medvec (1994). "The temporal pattern to the experience of regret." *Journal of Personality and Social Psychology*, 67(3): 357-365.

Halevy, Y. and V. Feltkamp (2005). "A Bayesian approach to uncertainty aversion." *Review of Economic Studies*, 72: 449-466.

Halevy, Y. (2005). "Ellsberg revisited: an experimental study." mimeo.

Harless, D. (1992). "Actions versus prospects: the effect of problem representation on regret." *American Economic Review*, 82: 634-649.

Kahneman, D. and D. T. Miller (1986). "Norm theory: Comparing reality to its alternatives." *Psychological Review*, 93: 136-153.

Kahneman, D. and A. Tversky (1979). "Prospect theory: An analysis of decision under risk." *Econometrica*, 47: 263-291.

Klibanoff, P., M. Marinacci, and S. Mukerji (2005). "A smooth model of decision making under ambiguity." *Econometrica*, 73(6): 1849-1892.

Koszegi, B, and M. Rabin (2005). "A model of reference dependent preferences." mimeo.

Krähmer, D. and R. Stone (2005. "Regret in sequential decision problems." mimeo

Loomes, G. and R. Sugden (1982). "Regret theory: An alternative approach to rational choice under uncertainty." *The Economic Journal*, 92: 805-824.

Mellers, B. (2000). "Choice and the relative pleasure of consequences." *Psychological Bulletin*, 126: 910-924.

Mellers, B., A. Schwartz and I. Ritov (1999). "Emotion-based choice." *Journal of Experimental Psychology: General*, 128: 332-345.

Mellers, B. and A. P. McGraw (2001). "Anticipated emotions as guides to choice." *Current Directions in Psychological Science*, 10(6): 210-214.

Milgrom, P. (1981). "Good news and bad news: Representation theorems and applications." *Bell Journal of Economics*, 13: 380-391.

Morris, S. (1997). "Risk, uncertainty, and hidden information." *Theory and Decision*, 42: 235-269.

Nau, R. F. (2005). "Uncertainty aversion with second-order utilities and probabilities." mimeo.

Rabin, M. (1998). "Psychology and Economics." *Journal of Economic Literature*, 36: 11-46.

Ritov, I. and J. Baron (2000). "Reluctance to vaccinate: omission bias and ambiguity." In C. R. Sunstein (ed.), *Behavioral Law and Economics*, (pp. 168-186). New York: Cambridge University Press.

Schmeidler, D. (1989). "Subjective probability and expected utility without additivity." *Econometrica*, 57: 571-587.

Segal, U. (1987). "The Ellsberg paradox and risk aversion: an anticipated utility approach." *International Economic Review*, 28: 175-202.

Segal, U. (1990). "Two-stage lotteries without the reduction axiom." *Econometrica*, 58: 349-377.

Starmer, C., and R. Sugden (1993). "Testing for juxtaposition and event-splitting effects." *Journal or Risk and Uncertainty*, 6: 235-254.

Starmer, C., and R. Sugden (1998). "Testing alternative explanations for cyclical choice." *Economica*, 65: 347-361.

Tversky, A., & Kahneman, D. (1982). "Causal schemas in judgments under uncertainty." In D. Kahneman, P. Slovic, and A. Tversky (Eds.), *Judgment under uncertainty: Heuristics and biases* (pp. 117-128). New York: Cambridge University Press.

Zeelenberg, M. (1999). "Anticipated regret, expected feedback, and behavioral decision making." *Journal of Behavioral Decision Making*, 12: 93-106.