Randomization Devices and the Elicitation of Ambiguity-Averse Preferences

Sophie Bade*

May 24, 2015

Abstract

In random incentive mechanisms agents choose from multiple problems and a randomization device selects a single problem to determine payment. Agents are assumed to act as if they faced each problem on its own. While this approach is valid when agents are expected utility maximizers, ambiguity-averse agents may use the randomization device to hedge and thereby contaminate the data.

Keywords: Random Incentive Mechanisms, Ambiguity Aversion, Experiments, Hedging. *JEL Classification Numbers*: C91,D81.

1 Introduction

Experimental economists often ask subjects to choose from several different problems simultaneously. One of these problems is then randomly drawn; the subject's choice from this problem determines the outcome of the experiment. The agent might, for example, be asked to report choices from six different sets of bets, with the experimenter then rolling a die to determine which of the six choices is payoff-relevant. Any experimental design which uses a randomization device to elicit choices from several problems is a **random incentive mechanism**.

*Royal Holloway College, Uuniversiy of London and Max Planck Institute for Research on Collective Goods, Bonn sophie.bade@rhul.ac.uk

If a subject's choice from each separate problem is identical to his choice from the same problem when it appears as part of a random incentive mechanism, the mechanism has many advantages over separate single choice experiments. First, large sets of data can be elicited with one payment. Second, the payments do not accrue; so the later choices in an experiment are not affected by the agent's earlier earnings or losses. Finally and most importantly, to check for behavioral regularities we must elicit choices from various problems; a single choice experiment carries no information about the consistency of an agent's behavior. But if an agent's behavior in a random incentive mechanism differs from his behavior in separate choice situations it is not clear how one should interpret the data generated by the mechanism.

Random incentive mechanisms have been used widely in the experimental literature on ambiguity aversion (Camerer and Weber [5], Halevy [10] and Ahn et al. [1]). However, there are no theoretical results on the incentives for ambiguity-averse agents to reveal their true preferences in such mechanisms. The present study argues that random incentive mechanisms stand on shaky ground when agents are ambiguity-averse. Consider a mechanism that is designed to elicit preferences over ambiguous acts. Let all acts that the agent can choose in the mechanism be conditioned on a set of possibly ambiguous events. If this set of events is independent of the randomization device then the agent can use the randomization device to hedge against the ambiguity associated with his choices. Preference reversals, where agents behave differently in random incentive mechanisms and in single choice experiments, will occur.

Example: an urn and a coin. There is an urn filled with 30 blue balls and 60 green and red balls in unknown proportion. We are interested in an agent's preferences over "urn-acts" $f = (f(B), f(G), f(R))$ where $f(B), f(G)$ and $f(R)$ denote the agent's utility-payoffs in the events B, G, and R that a blue, green, or red ball is drawn.[1] Let the agent choose among a "blue act" that delivers utility 5 when a blue ball is drawn from the urn, a "green act" that delivers utility 9 when a green ball is drawn and a "red act" which also delivers 9 when a red ball is drawn. Represent these acts as $blue\colon = (5, 0, 0)$,

[1]The assumption that acts map to lotteries over outcomes is more common. If we assume that an agent's preferences over objective lotteries has an expected utility representation on lotteries, we can derive acts f which directly map states to utilities from more basic acts g which map to lotteries over outcomes by letting $f(\omega) = u(g(\omega))$ for every state ω.

green: $= (0, 9, 0)$ and *red*: $= (0, 0, 9)$.

Assuming our agent believes that the probability of a blue ball is $\frac{1}{3}$ the preference *blue* $\succ$ *green* $\sim$ *red* is inconsistent with expected utility theory. If our agent was an expected utility maximizer he would have to believe that either R or G occurs with a probability of at least $\frac{1}{3}$. Consequently his preferred act among *red* and *green* would have to deliver an expected utility of at least $\frac{1}{3} \times 9$ whereas *blue* delivers only $\frac{1}{3} \times 5$. But an ambiguity-averse agent might well prefer the objective lottery *blue* to the acts *green* and *red* that leave winning probabilities uncertain. Would an ambiguity-averse agent reveal the preference *blue* $\succ$ *green* $\sim$ *red* in a random incentive mechanism?

To pose this question concretely, let the preference $\succsim$ over acts f be represented by a maxmin expected utility $U(f) = \min_{\pi \in C} \sum_{\Omega} f(\omega)\pi(\omega)$. Unlike an expected utility maximizer our agent holds a set of beliefs C on the state space Ω, not a single prior. He calculates an expected utility with respect to every prior in C and evaluates his overall utility as the lowest among these. To match our assumption that the agent believes a blue ball is drawn with probability $\frac{1}{3}$, let $\pi(B) = \frac{1}{3}$ hold for all $\pi \in C$. To reflect the agent's uncertainty about the probabilities of green and red balls let $\pi(G)$ either equal $\frac{1}{9}$ or $\frac{5}{9}$, implying that $\pi(R)$ also equals either $\frac{1}{9}$ or $\frac{5}{9}$.[2] So our agent evaluates any urn-act by either $\pi^{red} = (\frac{1}{3}, \frac{1}{9}, \frac{5}{9})$ or $\pi^{green} = (\frac{1}{3}, \frac{5}{9}, \frac{1}{9})$ where the components of these vectors denote the probabilities of the events B, G and R. Since our agent believes that a blue ball is drawn with probability $\frac{1}{3}$, his utility of *blue* is $\frac{1}{3} \times 5$. His utility of *green* is just $\pi^{red}(G)9 = 1$, since π^{red} is the most pessimistic prior in C to evaluate *green*. Similarly, the agent's utility of *red* is $\pi^{green}(R)9 = 1$. Our agent prefers *blue* to *green* and *red*.

Now let's construct a random incentive mechanism to elicit these preferences. First, let us ask our agent to choose one act each from $S_H\colon = \{blue, green\}$ and $S_T\colon = \{blue, red\}$. Let us then toss a fair coin to declare one of these two choices as payoff-relevant. If heads comes up, the agent is paid according to his choice from S_H, otherwise he is paid according to his choice from S_T.[3]

To model the agent's behavior we need to specify his preferences over acts that are not only conditioned on the events B, G and R but also on the events H and T that the coin comes up heads or tails. Define the state space Ω such that any state $\omega \in \Omega$ is the

[2] Since C is not convex, U deviates from the maximin expected utility model of Gilboa and Schmeidler [8]. The analysis goes through unchanged if we replace C with its convex hull.

[3] This example was inspired by the experimental setup in Ahn et al. [1].

intersection of a coin- and an urn-event. For example, the unique state ω at which the coin comes up heads and a blue ball is drawn is defined by $\{\omega\} = H \cap B$. Assuming that the agent assigns probability $\frac{1}{2}$ to heads and tails, let C consist of the priors defined by the following two matrices:

	B	G	R
H	$\frac{1}{6}$	$\frac{1}{18}$	$\frac{5}{18}$
T	$\frac{1}{6}$	$\frac{1}{18}$	$\frac{5}{18}$

	B	G	R
H	$\frac{1}{6}$	$\frac{5}{18}$	$\frac{1}{18}$
T	$\frac{1}{6}$	$\frac{5}{18}$	$\frac{1}{18}$

According to C, the agent evaluates any urn-act at π^{red} or at π^{green} and he assigns a probability of $\frac{1}{2}$ to either side of the coin. The coin and the urn are independent according to either prior in C; we have, for example, $\pi^{red}(H \cap R) = \pi^{red}(H)\pi^{red}(R)$.

Consider the choice of *green* from S_H and *red* from S_T in random incentive mechanism. For any prior π in C the expected utility of this plan is $\frac{1}{2}\pi(G)9+\frac{1}{2}\pi(R)9$. As $\pi(G)+\pi(R) = \frac{2}{3}$ holds for any $\pi \in C$, this expected utility equals $\frac{1}{2}\pi(G)9 + \frac{1}{2}(\frac{2}{3} - \pi(G))9 = 3$ for any prior $\pi \in C$. Consequently the maxmin expected utility of this plan also equals 3. On the other hand, choosing *blue* from both sets yields a utility of only $\min_{\pi \in C} \pi(B)5 = \frac{5}{3}$. The remaining two options (choosing *blue* from exactly one of the two sets) deliver a yet lower utility. In sum, there is a preference reversal. While the agent prefers *blue* to *green* and *red* it is optimal for him to choose *green* from S_H and *red* from S_T in the mechanism. □

The main result of the paper shows that the preceding example is no accident. Preference reversals must occur when agents are ambiguity-averse. To make this point I consider the two most popular models of ambiguity-averse preferences: the maxmin expected utility model of Gilboa and Schmeidler [8] and Klibanoff, Marinacci and Mukerji's [16] model of smooth ambiguity aversion. Fix a randomization device $\mathcal{D}$, defined by a set of "coin-events", and a set of - possibly ambiguous - events $\mathcal{A}$. Consider the set of random incentive mechanisms that present the agent with choice-sets of acts that are conditioned on events in $\mathcal{A}$ and that use the randomization device $\mathcal{D}$ to determine which of the agent's choices is operative for payment. Assume that $\mathcal{D}$ and $\mathcal{A}$ are independent and that the agent is strictly ambiguity-averse with respect to acts conditioned on $\mathcal{A}$. Theorem 1 shows that the agent's preference must exhibit a reversal in some mechanism in this set. Such reversals can be ruled out only if the agent is an expected utility maximizer with respect to the acts under study.

Schmeidler [20] suggests that "intuitively, uncertainty aversion means that 'smoothing' or averaging utility distributions makes the decision maker better off." This is exactly what happens here: our experimental subject uses the coin to smooth the two utility distributions associated with uncertain acts *green* and *red*. On its own, each of these acts delivers utility 9 in one uncertain event and 0 otherwise. But in the compound act according to which *green* is played if the coin comes up heads and *red* is played if tails, the event R is neither as advantageous as it is under *red* nor as unfavorable as it is under *green*. If the event R occurs the compound act delivers utility 9 with tails and 0 with heads and vice versa for G. The coin, therefore, averages the utilities delivered by the two urn-acts *green* and *red*. In the present case, this averaging or hedging is efficient enough to make the agent's choices in the random incentive mechanism differ from his choices in the two single choice experiments.

My arguments share some similarity with Karni and Safra's [13] and Holt's [11] analysis of preference reversals in random incentive mechanisms. These two studies show that some empirically well-documented reversals (Lichtenstein and Slovic [17] and Grether and Plott [9]) of preferences over lotteries are consistent with the assumption that agents have rank-dependent preferences. Similarly to the present paper, Karni and Safra [13] and Holt [11] argue that, without the assumption of expected utility preferences, an agent's behavior in the random incentive mechanism as a whole need not reflect the agent's behavior in single choice experiments. However Karni and Safra [13] and Holt [11] do not address the question whether random incentive mechanisms truthfully elicit the preferences of ambiguity-averse agents; their studies consider only the case of objective lotteries. To make sure that my results are driven by ambiguity aversion alone, I assume that preferences over objective lotteries have expected utility representations.

Bade [3] first raised the issue that an ambiguity-averse agents might use the randomization device of an experiment to hedge and thereby contaminate the data. Since then other authors have made related observations. Kuzmics [15] notes that it is impossible to detect the Ellsberg paradox experimentally if subjects can use coin tosses to smooth out all ambiguity in the experiment. Working in a more general setup Azrieli, Chambers and Healy [2] study the family of preferences under which random incentive mechanisms are incentive compatible. Oechssler and Roomets [18] point out that there is a potential hedging problem in experiments that use random incentives with ambiguity-averse agents. Noting the same problem, Baillon, Halevy and Li [4] identify conditions on the agent's preferences under which the incentive compatibility of random incentive mechanisms is restored. In Section 4 I argue that it is impossible to ascertain whether or not these conditions are met in any particular case.

The fact that some experimental studies on ambiguity aversion, such as Stahl [21], find inconclusive results can be viewed as empirical motivation for my study. If experimental subjects use the randomization device as a hedging device, the full extent of their ambiguity aversion will not be visible in the data. If other subjects do not hedge, the empirical picture might turn out very hard to analyze.

2 Definitions

2.1 Basics. The agent has a complete and transitive preference $\succsim$ over acts which are functions from a finite state space Ω to $\mathbb{R}$. Under the act f the agent obtains utility $f(\omega)$ in state ω. Anscombe-Aumann acts, in contrast, map a state space to lotteries over outcomes. If we assume an expected utility representation u over lotteries, then we can map any Anscombe-Aumann act g to an act $f : \Omega \to \mathbb{R}$ by letting $f(\omega) \colon = u(g(\omega))$ for all $\omega \in \Omega$. A **constant act** maps every state to the same utility level $x \in \mathbb{R}$. As a shorthand a constant act is also denoted x. The constant act $x_f \in \mathbb{R}$ that is indifferent to f is the **certainty equivalent** of f. For any pair of acts f, g and event $E \subset \Omega$, the **compound act** $f_E g$ delivers $f(\omega)$ if $\omega \in E$ and $g(\omega)$ otherwise. If f and g are constant acts then $f_E g$ is a **bet (on** E**)**. In the compound act $green_H red$ in the introductory example payoffs are determined by *green* in case of heads and by *red* in case of tails. The acts *blue*, *green* and *red* are all bets.

Fix some partition $\mathcal{P}$ of Ω. Then f is a $\mathcal{P}$-act if any two states in the same cell of $\mathcal{P}$ yield the same utility level. For any $\mathcal{P}$-act f and any cell E of $\mathcal{P}$ let $f(E) \colon = f(\omega)$ when $\omega \in E$. Any union of events in $\mathcal{P}$ is a $\mathcal{P}$-event, the complement of E is $\overline{E}$, and $\Delta\Omega$ is the set of probability distributions on Ω. For any $\pi \in \Delta\Omega$, the marginal distribution on $\mathcal{P}$ (the restriction of π to $\mathcal{P}$-events) is $\pi_{\mathcal{P}}$. The conditional distribution of $\mathcal{P}$-events given E is $\pi_{\mathcal{P}}(\cdot \mid E)$.

2.2 Random Incentive Mechanisms. The experimenter is interested in an agent's preference over the set of $\mathcal{A}$-acts, where some events in the partition $\mathcal{A}$ might be ambiguous. The experiment uses a **randomization device** $\mathcal{D}$. Formally $\mathcal{D} \colon = \{D_1, \ldots, D_n\}$ is another partition of the state space Ω into n cells. Both $\mathcal{A}$ and $\mathcal{D}$ are fixed throughout. In terms of the introductory example we have $\mathcal{D} = \{H, T\}$ and $\mathcal{A} = \{B, G, R\}$. Since the events in the partitions $\mathcal{A}$ and $\mathcal{D}$ are the only ones that matter to the present study, I assume that any state is the intersection of a $\mathcal{D}$-event and an $\mathcal{A}$-event: each $\omega \in \Omega$ is identified with a pair of events $A \in \mathcal{A}$ and $D_i \in \mathcal{D}$ such that $\{\omega\} = A \cap D_i$.

A **random incentive mechanism** S is an indexed collection $\{S_1, \ldots, S_n\}$ of sets S_i of $\mathcal{A}$ acts. The agent chooses one act from each S_i. Which of these acts is payoff relevant depends on the randomization device $\mathcal{D}$: if D_i is drawn, the agent is paid according to his choice from S_i. Any list of choices $(f[1], \ldots, f[n]) \in (S_1 \times \cdots \times S_n)$ defines an act $f : \Omega \to \mathbb{R}$ that is **available** in S via $f(\omega) = f[i](A)$ for $\{\omega\} = D_i \cap A$. Conversely, an act $f : \Omega \to \mathbb{R}$ is available in S if for each i there exists some $f' \in S_i$ such that $f(A \cap D_i) = f'(A)$ for all $A \in \mathcal{A}$. For any available f, define a set of n $\mathcal{A}$-acts $f[i]$ via $f[i](A) := f(A \cap D_i)$ for all $A \in \mathcal{A}$. The preference $\succsim$ **does not exhibit a preference reversal** in the random incentive mechanism S if the following equivalence holds for any available f^* in S[4]:

$$f^* \succsim f \text{ for all available } f \text{ in } S \Leftrightarrow$$
$$f^*[i] \succsim f[i] \text{ for all } f[i] \in S_i \text{ and all } i \in \{1, \ldots, n\}.$$

So $\succsim$ does not exhibit a preference reversal in S if the agent chooses the same $f^*[i]$ from the set S_i, whether he faces just that choice or has to choose from the entire list $\{S_1, \ldots, S_n\}$. Conversely, $\succsim$ exhibits a preference reversal in S if the agent's optimal choices within the mechanism differ from his optimal choices in the separate choice problems. If $\succsim$ does not exhibit a preference reversal in any random incentive mechanism (that uses the randomization device $\mathcal{D}$ to elicit preferences over $\mathcal{A}$-acts) then $\succsim$ is **transparent**.

2.3 Representations. A representation U of the preference $\succsim$ is a **MMEU representation** (maxmin expected utility representation, Gilboa and Schmeidler [8]) if there exists a convex and compact set of beliefs C on Ω such that $U(f) = \min_{\pi \in C} \sum_{\Omega} f(\omega)\pi(\omega)$ for all acts $f : \Omega \to \mathbb{R}$. A representation V is a **SAA representation** (smooth ambiguity aversion representation, Klibanoff, Marinacci and Mukerji [16]) if there exists a concave function $\phi : \mathbb{R} \to \mathbb{R}$ and a probability measure μ on the set of priors $\Delta\Omega$ such that

[4] To clarify, for $\Rightarrow$ the $\mathcal{A}$-acts $f^*[i]$ and $f[i]$ are derived from f^* and f, while for $\Leftarrow$ f^* and f are derived from the $\mathcal{A}$-acts $f^*[i]$ and $f[i]$.

$V(f) = \int_{\Delta\Omega} \phi\Big(\sum_\Omega f(\omega)\pi(\omega)\Big) d\mu(\pi)$ for all acts $f : \Omega \to \mathbb{R}$. Given that acts directly map states to utilities in this paper, a MMEU representation is defined through the set of beliefs C, and a SAA representation is defined through the prior over priors μ and the function ϕ. A preference that has either a MMEU or a SAA representation is **strictly ambiguity-averse** if it does not have an expected utility representation.

2.4 Independence. Given a probability measure π, two events E_1 and E_2 are **(stochastically) independent** if the probability of both events occurring equals the product of their probabilities, $\pi(E_1 \cap E_2) = \pi(E_1)\pi(E_2)$. Two partitions $\mathcal{P}_1$ and $\mathcal{P}_2$ are (stochastically) independent if any $\mathcal{P}_1$-event E_1 and a $\mathcal{P}_2$-event E_2 are independent.

Since the representations we consider involve multiple priors we cannot use the standard notion of independence to define the independence of $\mathcal{D}$ and $\mathcal{A}$. Instead we need a behavioral concept of independence, which I define following Klibanoff [14]. Fix two events E_1 and E_2 and a bet b, that delivers 1 if E_2 occurs and 0 otherwise. Consider the compound act $b_{E_1}x_b$ according to which the agent gets to play the bet b if E_1 occurs and receives x_b, the certainty equivalent of b, otherwise. If the agent is an expected utility maximizer with a prior π according to which E_1 and E_2 are independent then he must be indifferent between $b_{E_1}x_b$ and x_b. The reason is that the two acts differ only in the event E_1, and, due to the independence of E_1 and E_2, the preferred outcome of the bet b is just as likely under $\pi(\cdot \mid E_1)$ as it is under π.

Transferring this intuition to any preference $\succsim$, say that the events E_1 and E_2 are "independent" if $b_{E_i}x_b \sim b$ holds for all bets b on E_j and all $\{i, j\} = \{1, 2\}$. If the indifference $b_{E_1}x_b \sim b$ holds for a bet b on E_2 then the value that the agent assigns to b does not depend on E_1 occurring or not. So the agent cannot consider his preferred event under b, be it E_2 or $\overline{E}_2$, to be correlated with E_1. Generalizing this idea to the case of two partitions $\mathcal{P}_1$ and $\mathcal{P}_2$ we obtain:

Definition 1 *Two partitions $\mathcal{P}_1$ and $\mathcal{P}_2$ on Ω are* ***(behaviorally) independent*** *according to $\succsim$ if $f \sim f_E x_f$ holds for all pairs of a $\mathcal{P}_j$-act f and a $\mathcal{P}_i$-event E with $\{i, j\} = \{1, 2\}$.*

This definition applies the intuition developed above to any combination of a $\mathcal{P}_i$-event and a $\mathcal{P}_j$-act, not just bets. Klibanoff [14] shows that behavioral independence reduces to the classical definition when $\succsim$ has an expected utility representation.

3 An Impossibility Result

Independent randomization devices offer agents the opportunity to hedge. When the agent does hedge, his choices in the mechanism appear to be more ambiguity accepting than the choices he would make if he faced the problems separately. Any strictly ambiguity-averse preference exhibits a reversal in some random incentive mechanism with an independent randomization device. For the statement of the following theorem, recall that (1) $\succsim$ is a binary relation on acts $f : \Omega \to \mathbb{R}$, (2) each state $\omega \in \Omega$ is the intersection of a $\mathcal{D}$-event and an $\mathcal{A}$-event, (3) $\succsim$ is *transparent* if it does not exhibit a reversal in any random incentive mechanism that uses $\mathcal{D}$ to elicit preferences over $\mathcal{A}$-acts, and (4) $\succsim$, which either has a MMEU or a SAA representation, is *strictly ambiguity averse* if and only if it does not have an expected utility representation.

Theorem 1 *Let $\succsim$ either have a MMEU or a SAA representation and let $\mathcal{A}$ and $\mathcal{D}$ be independent according to $\succsim$. If $\succsim$ is transparent, then the restriction of $\succsim$ to $\mathcal{A}$-acts has an expected utility representation.*

If experiments use an independent randomization device to elicit preferences that can be represented by either of the two most prominent theories of ambiguity aversion, then Theorem 1 presents a dilemma: if the preferences under investigation (the preferences over $\mathcal{A}$-acts) are strictly ambiguity averse, then $\succsim$ must exhibit reversals in some experiments.

The proof of Theorem 1, which can be found in the Appendix, starts by translating transparency into a condition on preferences. Lemma 1 shows that $\succsim$ is transparent if and only if any act f is for any i indifferent to $f_{\overline{D}_i} x_{f[i]}$, which is identical to f except that the outcomes of f when some D_i obtains ($f[i]$) are replaced by $x_{f[i]}$, the certainty equivalent of $f[i]$. If the condition holds the agent prefers $f[i]_{D_i} g$ to $f'[i]_{D_i} g$ if and only if he prefers the $\mathcal{A}$-act $f[i]$ to the $\mathcal{A}$-act $f'[i]$: the act g which is paid if D_i does not occur is irrelevant for agents preference over $f[i]_{D_i} g$ and $f'[i]_{D_i} g$.

Now fix an agent's MMEU preferences over $\mathcal{A}$-acts and over $\mathcal{D}$-acts; that is fix two sets of beliefs $C^*_{\mathcal{A}}$ and $C^*_{\mathcal{D}}$ on $\mathcal{A}$ and $\mathcal{D}$, respectively. Consider all MMEU preferences on acts $f : \Omega \to \mathbb{R}$ that coincide with these fixed preferences on the sets of $\mathcal{A}$-acts and $\mathcal{D}$-acts. Lemma 2 shows that such a preference is transparent if and only if it is represented by a belief set C that consists of *all* priors π on Ω with the following two features. First $\pi_{\mathcal{D}}$, the distribution of the randomization device implied by π, is contained in $C^*_{\mathcal{D}}$. Second,

$\pi_{\mathcal{A}}(\cdot \mid D_i) \in C^*_{\mathcal{A}}$ holds for all $D_i \in \mathcal{D}$, so conditioning on any event of the randomization device, the distribution on $\mathcal{A}$ belongs to $C^*_{\mathcal{A}}$.[5] To evaluate an f, an agent with such a representation separately calculates the MMEU of each $\mathcal{A}$-act $f[i]$ using the set of beliefs $C^*_{\mathcal{A}}$. The MMEU of f is a weighed average of these MMEUs.

Reconsidering the introductory example of the urn and the coin, let the set of priors on the coin $C^*_{\mathcal{D}}$ contain only the prior where heads and tails are equally likely. Let the set of priors on the urn $C^*_{\mathcal{A}}$ consist of the two priors $\pi^{red} = (\frac{1}{3}, \frac{1}{9}, \frac{5}{9})$ and $\pi^{green} = (\frac{1}{3}, \frac{5}{9}, \frac{1}{9})$. Lemmas 1 and 2 imply that a transparent MMEU is defined by a set C' that consists of the following four priors:

	B	G	R
H	$\frac{1}{6}$	$\frac{1}{18}$	$\frac{5}{18}$
T	$\frac{1}{6}$	$\frac{1}{18}$	$\frac{5}{18}$

	B	G	R
H	$\frac{1}{6}$	$\frac{5}{18}$	$\frac{1}{18}$
T	$\frac{1}{6}$	$\frac{5}{18}$	$\frac{1}{18}$

	B	G	R
H	$\frac{1}{6}$	$\frac{1}{18}$	$\frac{5}{18}$
T	$\frac{1}{6}$	$\frac{5}{18}$	$\frac{1}{18}$

	B	G	R
H	$\frac{1}{6}$	$\frac{5}{18}$	$\frac{1}{18}$
T	$\frac{1}{6}$	$\frac{1}{18}$	$\frac{5}{18}$

The first two priors make up the set C in the introductory example: both satisfy $\pi_{\mathcal{D}}(H) = \frac{1}{2}$ and $\pi_{\mathcal{A}}(\cdot \mid H), \pi_{\mathcal{A}}(\cdot \mid T) \in \{\pi^{red}, \pi^{green}\}$. However, C does not fit Lemma 2's characterization of transparent MMEU-preferences as it misses some π with $\pi_{\mathcal{D}}(H) = \frac{1}{2}$ and $\pi_{\mathcal{A}}(\cdot \mid H), \pi_{\mathcal{A}}(\cdot \mid T) \in \{\pi^{red}, \pi^{green}\}$. According to the first two priors the agent evaluates the urn act in case of heads and the urn act in case of tails at the same posterior. Hedging occurs since the agent does not choose two different posteriors to pessimistically maximize the probability of the worst outcome of the two different bets. The hedging problem disappears under the MMEU defined by C' as the agent separately chooses either π^{red} or π^{green} to evaluate the urn-act in case of heads and the urn-act in case of tails.

Lemma 4 shows that $\mathcal{D}$ and $\mathcal{A}$ cannot be independent if $\succsim$ has a representation that is characterized by Lemma 2 and if $\succsim$ is strictly ambiguity averse with respect to $\mathcal{A}$-acts ($C_{\mathcal{A}}$ not a singleton). The two partitions $\mathcal{A}$ and $\mathcal{D}$ are not interchangeable in Lemma 2's characterization: The conditions, $\pi_{\mathcal{D}} \in C_{\mathcal{D}}$ and $\pi_{\mathcal{A}}(\cdot \mid D_i) \in C_{\mathcal{A}}$ for all $D_i \in \mathcal{D}$ treat $\mathcal{A}$ and $\mathcal{D}$ differently. This asymmetry clashes with the independence-requirement which treats the two partitions symmetrically.

The proof for the case of SAA-preferences follows the same pattern. Lemma 3 characterizes the set of transparent SAA preferences. According to any such SAA representa-

[5] The resulting set of beliefs is rectangular with respect to the filtration $\{D_1, \ldots, D_n\}$ in the sense of Epstein and Schneider [7]. This means that full Bayesian updating with respect to any D_i is dynamically consistent. What is not implied by dynamic consistency but by independence is that the set of posteriors on $\mathcal{A}$ when updating with respect to D_i is identical to the set of priors on $\mathcal{A}$.

tion, the agent is an expected utility maximizer with respect to the randomization device. Moreover the agent's unconditional prior over priors on $\mathcal{A}$ is identical to the agent's conditional prior over priors on $\mathcal{A}$ given any event D_i of the randomization device. Lemma 4 also concludes the proof of the SAA-part of Theorem 1: transparency in the context of ambiguity aversion requires an asymmetry between the agent's beliefs on $\mathcal{D}$ and on $\mathcal{A}$ that cannot coexist with the symmetric independence property.

4 Discussion

Theorem 1 is due to the clash among transparency, ambiguity aversion, and independence. To salvage random incentive mechanisms on might firstly impose a weaker notion of independence, secondly to drop independence altogether or thirdly weaken transparency.

Given that there is no agreed upon notion of stochastic independence for ambiguity-averse preferences one might consider replacing Klibanoff's [14] notion of independence by a weaker one. However, Klibanoff's [14] notion is the weakest in the literature. Gilboa and Schmeidler [8], for example, introduced a more restrictive notion of independence in their original article on MMEU representations: any two events that are independent according to that notion are also independent according to Klibanoff's [14]. Hence, Theorem 1 remains valid if independence is defined following Gilboa and Schmeidler [8].

Second, Lemmas 2 and 3 in the Appendix demonstrate that transparency and strict ambiguity aversion can be reconciled - if we replace independence by a no-hedging condition. To justify random incentive mechanisms Baillon, Halevy and Li [4] as well as Johnson et al. [12] refer to conditions on an agent's preferences under which a randomization device is not used for hedging. A new question arises: which randomization devices meet these conditions? Saito [19] convincingly argues that the suitability of a randomization device to hedge is a feature of the individual's preferences. He tells two equally appealing stories on the interaction between coin tosses and Ellsberg urns. According to one of these stories a coin toss serves to hedge away all uncertainty in the Ellsberg experiment, according to the other the coin does not mitigate the effects of uncertainty. So decision theory provides no guidance for the design of hedging-proof randomization devices. An empirical investigation whether an agent considers a randomization device well suited to hedge would require a large set of choice data. But to elicit such data we would have to set up a random incentive mechanism within which the agent does not hedge. Seemingly no progress has been made.

The third suggestion is to weaken transparency. Consider a mechanism in the coin-urn example in which the agent only gets to choose among "G-acts", whose payoff in the event G is at least as high as their payoff in the event R. Given that the better payoff is associated with the same ambiguous event for all acts that can be chosen, there is no hedging opportunity in such a mechanism. If we assume that any such mechanism is transparent we can truthfully elicit the agent's preferences over G-acts. Since there is no a priori difference between green and red balls in the description of the urn we might assume that the agent is indifferent between any f and h with $f(G) = h(R)$, $f(R) = h(G)$ and $f(B) = h(B)$. Combining the weakened transparency assumption with the symmetry assumption we can identify the agent's preference over any two acts. This approach indeed works for the agent with the MMEU preferences defined in the introduction. But for a preference over a more complex set of acts it might not be so easy to identify a subset of mechanisms where the preference does not exhibit reversals and to find an assumption that allows the researcher to derive the remainder of the preference from the elicited choice data. Consider payoff plans that are conditioned on the size and student evaluation scores of a particular class. Without any obvious "symmetry" between sizes and scores, it is hard to come up with an assumption that would relate preferences over acts conditioned on scores to preferences over acts conditioned on class sizes.

Finally, let me conjecture that other classes of ambiguity averse preferences suffer from the same hedging-in-experiments problem. Specifically, I believe that the conclusion of Theorem 1 holds for any uncertainty averse preference as defined by Cerreia-Vioglio, Maccheroni, Marinacci and Montrucchio [6].

5 Appendix

For any partition $\mathcal{P}$ of Ω let $\sigma_{\mathcal{P}}$ denote the algebra generated by $\mathcal{P}$. So an event E is a $\mathcal{P}$-event if and only if $E \in \sigma_{\mathcal{P}}$ and f is a $\mathcal{P}$-act if and only if f is $\sigma_{\mathcal{A}}$-measurable. Let $\succsim^{\mathcal{P}}$ denote the restriction of $\succsim$ to the set of $\mathcal{P}$-acts. For any set $C \subset \Delta\Omega$ and any partition $\mathcal{P}$ of Ω define $C_{\mathcal{P}}$ as the set marginal distributions on $\pi_{\mathcal{P}}$ of all priors in the set C, formally $C_{\mathcal{P}} = \{\pi_{\mathcal{P}} \mid \pi \in C\}$. Let $\Sigma_{\mathcal{A}}$ be the algebra on $\Delta\Omega$ generated by the partition of $\Delta\Omega$ into sets $\{\pi \mid \pi_{\mathcal{A}} = \pi^*_{\mathcal{A}}\}$. For any prior μ over priors $\Delta\Omega$, let $\mu_{\Sigma_{\mathcal{A}}}$ be the marginal distribution with respect to the algebra $\Sigma_{\mathcal{A}}$. For any set S let $co(S)$ be the convex hull of the set S, so $co(S)$ is the smallest convex set that contains S.

An event E is **ambiguous** according to a MMEU-representation associated with the set of beliefs C if $\{\pi(E) \mid \pi \in C\}$ is not a singleton. If $\succsim$ has a SAA representation associated with μ, ϕ, then E is ambiguous if ϕ is strictly concave and $\{\pi(E) \mid \pi \in supp(\mu)\}$ is neither a singleton nor a subset of $\{0, 1\}$. The preference $\succsim^{\mathcal{A}}$ is strictly ambiguity averse if and only if there exists an ambiguous $\mathcal{A}$-event. If no $\mathcal{A}$-event is ambiguous and if $\succsim^{\mathcal{A}}$ either has a MMEU- or an SAA-representation, then $\succsim^{\mathcal{A}}$ has an expected utility representation. Represent any $f : \Omega \to \mathbb{R}$ as $f : = (f[1], \ldots, f[n])$ with the understanding that $f[i]$ is a $\mathcal{A}$-act defined by $f[i](A) = f(\omega)$ for $\{\omega\} = A \cap D_i$ and write $f \in S$ if f is available in S.

The randomization device $\mathcal{D}$ is **isolated** from $\mathcal{A}$ if $f \sim x_{D_i} f$ holds for $x \in \mathbb{R}$, act f and $D_i \in \mathcal{D}$ if and only if $x \sim f[i]$. So if $\mathcal{D}$ is isolated from $\mathcal{A}$ then knowing D_i does not make $f[i]$ any more or less attractive than $x_{f[i]}$, no matter what happens for all the other outcomes of the randomization device. If $\mathcal{D}$ is isolated from $\mathcal{A}$ then we have $f = (f[1], \ldots, f[n]) \sim (x_{f[1]}, \ldots, f[n]) \sim (x_{f[1]}, x_{f[2]}, \ldots, f[n]) \sim \cdots \sim (x_{f[1]}, \ldots x_{f[n]})$ (**Fact 1**). This observation is used in the proofs of Lemmas 1 and 2:

Lemma 1 *Assume that $\succsim$ is monotonic in the sense that $f \succsim f'$ holds when $f(\omega) \succsim f'(\omega)$ holds for all ω. Then $\succsim$ is transparent if and only if $\mathcal{D}$ is isolated from $\mathcal{A}$.*

Proof The statement can be formalized as $(I) \Leftrightarrow (II)$ with

$$(I): \quad (f^* \succsim f \ \forall f \in S) \Leftrightarrow (f^*[i] \succsim f[i] \ \forall \ f[i] \in S_i, i) \quad \forall (S, f^* \in S)$$
$$(II): \quad (x_{D_i} f \sim f) \Leftrightarrow (x \sim f[i]) \quad \forall (x, i, f).$$

To see $(I) \Rightarrow (II)$, fix a triple (x, i^*, f^*) and assume that (II) does not hold. To simplify notation let $g : = f^*[i^*]$ and $D_{i^*} : = D$.

First assume that $f^\circ : = x_D f^* \sim f^*$ and $x \not\sim g$ hold. Define S through $S_{i^*} = \{x, g\}$ and $S_i = \{f^*[i]\}$ for all $i \neq i^*$. So $f^\circ \sim f^* \succsim f$ holds for all $f \in S$, however $x \not\sim g$ implies that one of the two acts $f^\circ[i^*] = x$ and $f^*[i^*] = g$ must be strictly preferred. So (I) is violated. Next assume that $f^\circ : = x_{g_D} f^* \not\sim f^*$ holds. Define S through $S_{i^*} = \{x_g, g\}$ and $S_i = \{f^*[i]\}$ for all $i \neq i^*$. Observe that $f^\circ[i] \sim f^*[i] \succsim f[i]$ holds for all $f[i] \in S_i$ and all

i (including i^*), however $f^\circ \not\sim f^*$ implies that one of these two acts must be preferred to the other. So - once again - (I) is violated.

To see $(II) \Rightarrow (I)$ fix a tuple $(S, f^* \in S)$ and assume that (II) holds.

First assume that $f^* \succsim f$ holds for all $f \in S$ while $f^*[i^*] \prec g$ holds for some i^* and $g \in S_{i^*}$. Let $D_{i^*} :\, = D$ and define x such that $f^* \sim x_D f^*$ holds. By (II) x is uniquely defined through $x \sim f^*[i^*]$. This, $f^*[i^*] \prec g$, and monotonicity imply $x_D f^* \prec x_{g_D} f^*$. Applying (II) once again we obtain $x_{g_D} f^* \sim g_D f^*$ and therefore $g_D f^* \succ f^*$, which stands in contradiction with $f^* \succsim f$ for all $f \in S$. We can conclude that $(f^* \succsim f \; \forall f \in S)$ implies $(f^*[i] \succsim f[i] \; \forall \; f[i] \in S_i, i)$ if (II) holds. Now assume that $f^*[i] \succsim f[i]$ holds for all $f[i] \in S_i$ and all i. So we have that $x_{f^*[i]} \succsim x_{f[i]}$ holds for all $f[i] \in S_i$ and all i. Monotonicity and Fact 1 (which follows from (II)) imply that $f^* \sim (x_{f^*[1]}, \ldots, x_{f^*[n]}) \succsim (x_{f[1]}, \ldots, x_{f[n]}) \sim f$ holds for all $f \in S$. In sum we obtain that $(f^*[i] \succsim f[i] \; \forall \; f[i] \in S_i, i)$ implies $(f^* \succsim f \; \forall f \in S)$ when (II) holds. □

Lemma 2 *Assume that $\succsim$ has a MMEU representation $U(f) = \min_{\pi \in C} \sum_\Omega f(\omega)\pi(\omega)$ and let $C^* :\, = co\big(\{\pi \mid \pi_{\mathcal{D}} \in C_{\mathcal{D}} \text{ and } \pi_{\mathcal{A}}(\dot{|}D_i) \in C_{\mathcal{A}} \text{ for all } i = 1, \ldots, n\}\big)$. Then $\mathcal{D}$ is isolated from $\mathcal{A}$ if and only if $C^* = C$ as well as $\pi(D_i) > 0$ for all i and $\pi \in C$.*

Proof Let $C = C^*$ as well as $\pi(D_i) > 0$ for all i and $\pi \in C$. Fix any act f and i^*, define $D_{i^*} :\, = D$ and $f[i^*] :\, = g$, so $U(f) = U(g_D f)$ can be calculated as

$$\min_{\pi \in C} \Big(\pi(D) \sum_{A \in \mathcal{A}} g(A)\pi(A \mid D) + \sum_{i \neq i^*} \pi(D_i) \sum_{A \in \mathcal{A}} f[i](A)\pi(A \mid D_i)\Big) =$$

$$\pi^*_{\mathcal{D}}(D) \min_{\pi_{\mathcal{A}} \in C_{\mathcal{A}}} \sum_{A \in \mathcal{A}} g(A)\pi_{\mathcal{A}}(A) + \sum_{i \neq i^*} \pi^*_{\mathcal{D}}(D_i) \min_{\pi^i_{\mathcal{A}} \in C_{\mathcal{A}}} \sum_{A \in \mathcal{A}} f[i](A)\pi^i_{\mathcal{A}}(A) =$$

$$\pi^*_{\mathcal{D}}(D) x_g + \sum_{i \neq i^*} \pi^*_{\mathcal{D}}(D_i) \min_{\pi^i_{\mathcal{A}} \in C_{\mathcal{A}}} \sum_{A \in \mathcal{A}} f[i](A)\pi^i_{\mathcal{A}}(A) = U\big(x_{g_D} f\big).$$

Letting $\pi^*_{\mathcal{D}} \in C_{\mathcal{D}}$ be the marginal on $\mathcal{D}$ at which the sum is minimized, the first first equality follows from $C = C^*$. The second follows from the definition of the certainty equivalent of g. The third is once again implied by the definition of $C = C^*$. Since $\pi^*_{\mathcal{D}}(D)$ is positive $U(f) = U\big(x_D f\big)$ holds only if $x = x_g$. So $\mathcal{D}$ is isolated from $\mathcal{A}$.

To see the necessity of the conditions for isolation, observe that isolation is violated if $\pi^*(D) = 0$ held for some $D \in \mathcal{D}$ and $\pi^* \in C$. In that case we have $U(1_D 0) = 0 = U(0) =$

$U(0_D0)$ and $1 \not\sim 0$, a violation of isolation. [6]

Next suppose that $C^* \neq C$. First suppose there exists a $\pi^* \in C^* \setminus C$. Since C and C^* are both convex the separating hyperplane theorem implies the existence of an act f^* such that $\min_{\pi\in C^*} \sum_\Omega f^*(\omega)\pi(\omega) < \min_{\pi\in C} \sum_\Omega f^*(\omega)\pi(\omega)$. The arguments in the sufficiency part of the proof imply that $\min_{\pi\in C^*} \sum_\Omega f^*(\omega)\pi(\omega) = \min_{\pi_\mathcal{D}\in C_\mathcal{D}} \sum_{i=1}^n \pi_\mathcal{D}(D_i)x_{f^*[i]} = U\big((x_{f^*[1]},\dots,x_{f^*[n]})\big)$ and we obtain $(x_{f^*[1]}, x_{f^*[2]},\dots,x_{f^*[n]}) \prec f^*$. But since $\mathcal{D}$ is isolated from $\mathcal{A}$, Fact 1 implies $(x_{f^*[1]}, x_{f^*[2]},\dots,x_{f^*[n]}) \sim f^*$, a contradiction. The case that there exists a $\pi^* \in C \setminus C^*$ is covered by the same arguments mutatis mutandis. □

Lemma 3 *Assume that $\succsim$ has a SAA representation $V(f) = \int_{\Delta\Omega} \phi\big(\sum_\Omega f(\omega)\pi(\omega)\big)d\mu(\pi)$ and that the preference $\succsim^\mathcal{A}$ over $\mathcal{A}$-acts is strictly ambiguity-averse. Then $\mathcal{D}$ is isolated of $\mathcal{A}$ if and only if $\pi(D)\pi(\overline{D}) = 0$, $\mu(\{\pi \mid \pi(D) = 1\}) > 0$ and $\mu_{\Sigma_\mathcal{A}} = \mu_{\Sigma_\mathcal{A}}(\cdot \mid \pi(D) = 1)$ hold for all $D \in \mathcal{D}$ and all $\pi \in supp(\mu)$.*

Proof First assume that $\pi(D)\pi(\overline{D}) = 0$, $\mu(\{\pi \mid \pi(D) = 1\}) > 0$ and $\mu_{\Sigma_\mathcal{A}} = \mu_{\Sigma_\mathcal{A}}(\cdot \mid \pi(D) = 1)$ hold for all $D \in \mathcal{D}$ and all $\pi \in supp(\mu)$. Fix any act f and i^*, define $D_{i^*} :\ = D$ and $f[i^*]:\ = g$. Since $\pi(D)\pi(\overline{D}) = 0$ holds for all $\pi \in supp(\mu)$ we can represent $V(f) = V(g_Df)$ as $\int_{\Delta\Omega,\pi(D)=1} \phi\Big(\sum_\Omega f(\omega)\pi(\omega)\Big)d\mu(\pi) + \int_{\Delta\Omega,\pi(D)=0} \phi\Big(\sum_\Omega f(\omega)\pi(\omega)\Big)d\mu(\pi)$. Rewrite the first term of the sum as follows:

$$\begin{aligned}\int_{\Delta\Omega,\pi(D)=1} \phi\Big(\sum_\Omega f(\omega)\pi(\omega)\Big)d\mu(\pi) &= \int_{\Delta\Omega,\pi(D)=1} \phi\Big(\sum_{A\in\mathcal{A}} g(A)\pi(A)\Big)d\mu(\pi) = \\ \mu(\{\pi \mid \pi(D) = 1\}) &\int_{\Delta\Omega} \phi\Big(\sum_{A\in\mathcal{A}} g(A)\pi(A)\Big)d\mu(\pi \mid \pi(D) = 1) = \\ \mu(\{\pi \mid \pi(D) = 1\}) &\int_{\Delta\Omega} \phi\Big(\sum_{A\in\mathcal{A}} g(A)\pi(A)\Big)d\mu(\pi) = \\ \mu(\{\pi \mid \pi(D) = 1\})\phi\big(x_g\big) &= \int_{\Delta\Omega,\pi(D)=1} \phi\big(x_g\big)d\mu(\pi).\end{aligned}$$

The first and second equality follow from the restriction to probability measures π with $\pi(D) = 1$ and the definition of the conditional probability $\mu(\cdot \mid \pi(D) = 1)$. The third equality holds since the marginal $\mu_{\Sigma_\mathcal{A}}$ is equal to the conditional marginal $\mu_{\Sigma_\mathcal{A}}(\cdot \mid \pi(D) = 1)$. The fourth equality uses the definition of the certainty equivalent of g. In sum we

[6] This argument also applies to the ambiguity neutral case: $\succsim$ is not transparent if the agent has an expected utility representation with a prior π such that $\pi(D) = 0$ holds for some $D \in \mathcal{D}$.

have $V\big(f\big) = V\big(g_D f\big) = V\big(x_{g_D} f\big)$. Since $\mu(\{\pi \mid \pi(D) = 1\})$ is positive the equality $V\big(f\big) = V\big(x_D f\big)$ holds if and only if $x = x_g$. So $\mathcal{D}$ is isolated from $\mathcal{A}$. To see the necessity of the conditions for isolation, suppose first of all that $\mu(\{\pi \mid \pi(D) = 1\}) = 0$ held for some D. Then we have $V(1_D 0) = 0 = V(0) = V(0_D 0)$ even though $1 \not\sim 0$. So isolation is violated.

Next suppose that $\pi^*(D)\pi^*(\overline{D}) \neq 0$ held for some $\pi^* \in supp(\mu)$ and $D \in \mathcal{D}$. Since $\succsim^{\mathcal{A}}$ is strictly ambiguity-averse there must exist an $\mathcal{A}$-event A such that $\{\pi(A) \mid \pi \in supp(\mu)\}$ is neither a singleton set, nor a subset of $\{0,1\}$, moreover ϕ must be strictly concave. Normalize ϕ such that $\phi(0) = 0$ and ϕ is strictly concave in some neighborhood around 0. Implicitly define a continuous function $y : \mathbb{R}_0^+ \to \mathbb{R}_0^-$ through $V\big(x_A y(x)\big) = 0$ and let $b[x] :=x_A y(x)$, so $y(0) = 0$.

Since $\mathcal{D}$ is isolated from $\mathcal{A}$, $b[x]_D 0 \sim 0 \sim 0_D b[x]$ must hold for all $x \geq 0$:

$$0 = V\big(b[x]_D 0\big) + V\big(0_D b[x]\big) =$$
$$\int_{\Delta\Omega} \phi\big(x\pi(A\cap D) + y(x)\pi(\overline{A}\cap D)\big)d\mu(\pi) + \int_{\Delta\Omega} \phi\big(x\pi(A\cap \overline{D}) + y(x)\pi(\overline{A}\cap \overline{D})\big)d\mu(\pi) =$$
$$\int_{\Delta\Omega} \phi\big(x\pi(A\cap D) + y(x)\pi(\overline{A}\cap D)\big) + \phi\big(x\pi(A\cap \overline{D}) + y(x)\pi(\overline{A}\cap \overline{D})\big)d\mu(\pi) \geq$$
$$\int_{\Delta\Omega} \phi\big(x\pi(A\cap D) + y(x)\pi(\overline{A}\cap D) + x\pi(A\cap \overline{D}) + y(x)\pi(\overline{A}\cap \overline{D})\big)d\mu(\pi) =$$
$$\int_{\Delta\Omega} \phi\big(x\pi(A) + y(x)\pi(\overline{A})\big)d\mu(\pi) = V(b[x]) = 0$$

The concavity of ϕ implies the weak inequality.

The strict concavity of ϕ at 0 together with the assumption that $\phi(0) = 0$ implies that $\phi(\alpha) + \phi(\beta) > \phi(\alpha + \beta)$ holds, if and only if $\alpha \neq 0 \neq \beta$. So either $x\pi(A\cap D) + y(x)\pi(\overline{A}\cap D) = 0$ or $x\pi(A\cap \overline{D}) + y(x)\pi(\overline{A}\cap \overline{D}) = 0$ or both must must hold for any $\pi \in supp(\mu)$, in particular for π^*. Since y is a continuously decreasing function with $y(0) = 0$, $y(x) = -\rho x$ must hold for ρ either $\frac{\pi^*(A\cap D)}{\pi^*(\overline{A}\cap D)}$ or $\frac{\pi^*(A\cap \overline{D})}{\pi^*(\overline{A}\cap \overline{D})}$. For any arbitrary x we then obtain

$$0 = V(b[\frac{x}{2}]) = \int_{\Delta\Omega} \phi\big((\frac{1}{2}0 + \frac{1}{2}x)(\pi(A) - \rho\pi(\overline{A}))\big)d\mu(\pi) >$$
$$\frac{1}{2}\int_{\Delta\Omega} \phi(0)d\mu(\pi) + \frac{1}{2}\int_{\Delta\Omega} \phi(x(\pi(A) - \rho\pi(\overline{A})))d\mu(\pi) = 0 + \frac{1}{2}V(b[x]) = 0.$$

The inequality follows since ϕ is strictly concave around 0, and since $\pi(A) - \rho\pi(\overline{A}) \neq 0$ must hold for a set of π that has positive measure according to μ (given that $\{\pi(A) \mid \pi \in supp(\mu)\}$ is not a singleton). The conclusion follows from $\phi(0) = 0$ and the definition of $b[x]$. In sum $\pi(D)\pi(\overline{D}) = 0$ must hold for all $\pi \in supp(\mu)$ and all $D \in \mathcal{D}$.

If $\mu_{\Sigma_{\mathcal{A}}} \neq \mu_{\Sigma_{\mathcal{A}}}(\cdot \mid \pi(D) = 1)$ holds for some $D \in \mathcal{D}$ then there exists an $\mathcal{A}$-act f with

$$\int_{\Delta\Omega} \phi\Big(\sum_{A\in\mathcal{A}} f(A) d\pi(\omega)\Big) d\mu(\pi) \neq \int_{\Delta\Omega} \phi\Big(\sum_{A\in\mathcal{A}} f(A) d\pi(\omega)\Big) d\mu(\pi \mid \pi(D) = 1).$$

For f we have the contradiction $V\big(f_D x_f\big) = \int_{\Delta\Omega} \phi\Big(\sum_{\mathcal{A}} (f_D x_f)(A)\pi(A)\Big) d\mu(\pi) =$

$$\mu(\{\pi \mid \pi(D) = 1\}) \int_{\Delta\Omega} \phi\Big(\sum_{A\in\mathcal{A}} f(A) d\pi(\omega)\Big) d\mu(\pi \mid \pi(D) = 1) + \mu(\{\pi \mid \pi(D) = 0\})\phi(x_f) \neq$$
$$\mu(\{\pi \mid \pi(D) = 1\})\phi(x_f) + \mu(\{\pi \mid \pi(D) = 0\})\phi(x_f) = V(f).$$

The inequality is implied by $\mu(\{\pi \mid \pi(D) = 1\}) > 0$. □

Lemma 4 *Assume that $\succsim$ either has a MMEU or an SAA representation, that $\mathcal{D}$ is isolated from $\mathcal{A}$ and that $\succsim^{\mathcal{A}}$ is strictly ambiguity-averse. Then $\mathcal{D}$ and $\mathcal{A}$ cannot be independent.*

Proof Since $\succsim^{\mathcal{A}}$ is strictly ambiguity-averse there exists an ambiguous $\mathcal{A}$-event A. Fix a $D \in \mathcal{D}$ and define a value x and a bet b such that $b = 1_D x \sim 0$. I show that $b_A 0 \not\sim 0$ holds for either representation contradicting the independence of $\mathcal{D}$ and $\mathcal{A}$ which requires $b_A 0 \sim 0$.

If the preference has the MMEU representation $U(f) = \min_{\pi\in C} \sum_\Omega f(\omega)\pi(\omega)$, then $x = -\min_{\pi\in C} \Big(\pi(D)/\pi(\overline{D})\Big) := \pi^*(D)/\pi^*(\overline{D})$ holds; Lemma 2 implies $\pi^*(D) \neq 0 \neq \pi^*(\overline{D})$ and $x \neq 0$. We can calculate $U\big(b_A 0\big)$ as

$$\begin{aligned}
\min_{\pi\in C} \Big(\pi(A\cap D) + x\pi(A\cap \overline{D}) + 0\pi(\overline{A})\Big) &= \min_{\pi^1\in C_{\mathcal{D}}, \pi^2, \pi^3 \in C_{\mathcal{A}}} \Big(\pi^2(A)\pi^1(D) + x\pi^3(A)\pi^1(\overline{D})\Big) = \\
\min_{\pi^2\in C_{\mathcal{A}}} \pi^2(A)\pi^*(D) + x \max_{\pi^3\in C_{\mathcal{A}}} \pi^3(A)\pi^*(\overline{D}) &= \pi^*(D)(\min_{\pi^2\in C_{\mathcal{A}}} \pi^2(A) - \max_{\pi^3\in C_{\mathcal{A}}} \pi^3(A)) < 0
\end{aligned}$$

The first equality follows from Lemma 2 which shows that C must be defined as $co(\{\pi \mid \pi_{\mathcal{D}} \in C_{\mathcal{D}}$ and $\pi_{\mathcal{A}}(\cdot \mid D_i) \in C_{\mathcal{A}}$ for all $i = 1, \ldots, n\})$ for $\mathcal{D}$ to be isolated from $\mathcal{A}$. The second equality recognizes the fact that a difference is minimized through minimising the minuend and maximizing the subtrahend. The third uses the definition of x. The inequality holds since A is ambiguous, meaning that $\{\pi(A) \mid \pi \in C\}$ is not a singleton.

Now assume that $\succsim$ has the SAA representation $V(f) = \int_{\Delta\Omega} \phi(\sum_\Omega f(\omega)\pi(\omega)) d\mu(\pi)$ with $\phi(0) = 0$ and ϕ strictly concave in some open interval around 0. Since $b \sim 0$, we have $V(b) = \lambda\phi(1) + (1-\lambda)\phi(x) = 0$ for $\lambda\colon = \mu(\{\pi \mid \pi(D) = 1\})$.

$$V(b_A 0) = \int_{\Delta\Omega} \phi\Big((\pi(A\cap D) + x\pi(A\cap \overline{D}) + 0\pi(\overline{A}\cap D) + 0\pi(\overline{A}\cap \overline{D})))\pi(\omega)\Big)d\mu(\pi) =$$

$$\lambda \int_{\Delta\Omega} \phi\Big(\pi(A) + 0\pi(\overline{A})\Big)d\mu(\pi) + (1-\lambda)\int_{\Delta\Omega} \phi\Big(x\pi(A) + 0\pi(\overline{A}))\Big)d\mu(\pi) >$$

$$\lambda \int_{\Delta\Omega} \Big(\phi(1)\pi(A) + \phi(0)\pi(\overline{A}))\Big)d\mu(\pi) + (1-\lambda)\int_{\Delta\Omega} \Big(\phi(x)\pi(A) + \phi(0)\pi(\overline{A})\Big)d\mu(\pi) =$$

$$\int_{\Delta\Omega} \Big((\lambda\phi(1) + (1-\lambda)\phi(x))\pi(A)\Big)d\mu(\pi) = \int_{\Delta\Omega} \Big(0\pi(A)\Big)d\mu(\pi) = 0.$$

The first equality follows from the definition of $b_A 0$. The second equality is implied by Lemma 3 which shows that $\pi(D)\pi(\overline{D}) = 0$ and $\mu_{\Sigma_{\mathcal{A}}} = \mu_{\Sigma_{\mathcal{A}}}(\cdot \mid \pi(D) = 1)$ must hold for all $D \in \mathcal{D}$ and all $\pi \in supp(\mu)$ for $\mathcal{D}$ to be isolated from $\mathcal{A}$ when $\succsim^{\mathcal{A}}$ is strictly ambiguity-averse. The inequality follows from the assumption that ϕ is strictly concave around 0 and $\mu(\{\pi \mid 0 < \pi(A) < 1\}) > 0$ as implied by A being ambiguous. The next equality follows from $\phi(0) = 0$, finally $V(b) = \lambda\phi(1) + (1-\lambda)\phi(x) = 0$ yields the conclusion. □

To prove Theorem 1 all preceding Lemmas need to be combined: Assume that $\succsim$ has a MMEU or a SAA representation. Given that $\succsim$ is monotonic, Lemma 1 applies; $\succsim$ is transparent if and only if $\mathcal{D}$ is isolated from $\mathcal{A}$. Lemma 4 shows that $\mathcal{D}$ and $\mathcal{A}$ cannot be independent if $\mathcal{D}$ is isolated from $\mathcal{A}$ and if $\succsim^{\mathcal{A}}$ is strictly ambiguity-averse.

References

[1] Ahn, D., S. Choi, D. Gale, and S. Kariv: "Estimating Ambiguity Aversion in a Portfolio Choice Experiment", *Quantitative Economics*, 5, (2014), pp. 195223.

[2] Azrieli, Y., C. Chambers and P. Healy: "Incentives in Experiments: a Theoretical Analysis", *mimeo* (2014) Ohio State University.

[3] Bade, S. "Independent Randomization Devices and the Elicitation of Ambiguity Averse Preferences" Max Planck Insitut Bonn, working paper, (2011).

[4] Baillon, A., Y. Halevy, and C. Li. "Experimental elicitation of ambiguity attitude using the random incentive system." *mimeo* (2014) Erasmus University Rotterdam.

[5] Camerer, C. and M. Weber: ”Recent Developments in Modeling Preferences: Uncertainty and Ambiguity”, *Journal of Risk and Uncertainty*, 5, (1992), pp. 325-370.

[6] Cerreia-Vioglio S., F. Maccheroni, M. Marinacci, and L. Montrucchio “Uncertainty Averse Preferences” *Journal of Economic Theory*, 146, (2011), pp. 1275-1330.

[7] Epstein, L. and M. Schneider “Recursive Multiple-Priors”, *Journal of Economic Theory*, 113, (2003), pp. 1-31.

[8] Gilboa, I. and D., Schmeidler, “Maxmin Expected Utility with Non-Unique Prior”, *Journal of Mathematical Economics*, 18, (1989), pp. 141-153.

[9] Grether, D. and C. Plott, “Economic Theory of Choice and the Preference Reversal Phenomenon,” *American Economic Review*, 69, (1979), pp. 623-638.

[10] Halevy, Y. “Ellsberg Revisited: An Experimental Study”, *Econometrica*, 75, (2007), pp. 503-536.

[11] Holt, C., “Preference Reversals and the Independence Axiom” *The American Economic Review*, 76, (1986), pp. 508-515.

[12] Johnson, C., A. Baillon, H. Bleichrodtb, Z. Lic, D. van Dolder, P. Wakker, “Prince: An Improved Method for Measuring Incentivized Preferences” *mimeo* (2014) University of Arizona.

[13] Karni, E. and Z. Safra, “ ’Preference Reversal’ and the Observability of Preferences by Experimental Methods” *Econometrica*, 55, (1987), pp. 675-685.

[14] Klibanoff, P., “Stochastically Independent Randomization and Uncertainty Aversion”, *Economic Theory*, 18, (2001), pp. 605-620.

[15] Kuzmics, C., “A rational ambiguity averse person will never display her ambiguity aversion”, *mimeo* Univeristy of Bielefed (2013).

[16] Klibanoff, P., M. Marinacci, and S. Mukerji, “A Smooth Model of Decision Making under Ambiguity”, *Econometrica*, 73, (2005), pp. 1849-1892.

[17] Lichtenstein S. and P. Slovic, “Reversals of Preferences Between Bids and Choices in Gambling Decisions,” *Journal of Experimental Psychology*, 89, (1971), pp. 46-55.

[18] Oechssler, J. and A. Roomets: "Unintended Hedging in Ambiguity Experiments", *Economics Letters* 122 (2014), pp. 243-246.

[19] Saito, K. "Preferences for Flexibility and Randomization under Uncertainty" *American Economic Review* forthcoming.

[20] Schmeidler, D. "Subjective Probability and Expected Utility without Additivity", *Econometrica*, 57, (1989), pp. 571-587.

[21] Stahl, D. "Heterogeneity of Ambiguity Preferences", *Review of Economics and Statistics*, (2013), forthcoming.

[22] Starmer, C. and R. Sugden "Does the Random-Lottery Incentive System Elicit True Preferences? An Experimental Investigation," *The American Economic Review*, 81, (1991), pp. 971-978