Essays on Decision-Making Under Risk

by

David Freeman

B.A., The University of Guelph, 2006
M.A., The University of British Columbia, 2007

A THESIS SUBMITTED IN PARTIAL FULFILLMENT OF
THE REQUIREMENTS FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

in

The Faculty of Graduate and Postdoctoral Studies

(Economics)

THE UNIVERSITY OF BRITISH COLUMBIA

(Vancouver)

October 2013

Abstract

Chapter 2 of this thesis studies the testable content of models of expectations-based reference-dependence. The main results of this chapter characterize a model based on Kőszegi and Rabin's (2006) preferred personal equilibrium. This model is shown to be behaviourally equivalent to a version of the shortlisting model of Manzini and Mariotti (2007) in environments without risk. Environments with risk motivate novel axioms that are conceptually consistent with expectations-based reference-dependence. These axioms are shown to behaviourally characterize a restricted version of the preferred personal equilibrium model of decision-making. Additional results characterize the choice behaviour of Kőszegi and Rabin's (2006, 2007) personal equilibrium and choice-acclimating personal equilibrium concepts.

In the presence of background risk and under the Reduction of Compound Lotteries axiom, non-expected utility preferences cannot capture descriptively reasonable risk aversion over small stakes without producing implausible risk aversion over large stakes (Safra and Segal 2008). Motivated by experimental evidence, Chapter 3 assumes that compound lotteries are evaluated recursively. The main results of this chapter show that non-expected utility theories generate 'as-if' narrow bracketing over small-stakes gambles despite defining utility over final wealth, and can be consistent with empirically reasonable risk aversion over both small and large stakes.

Chapter 4 uses list elicitation with varying probabilities to experimentally study choices between lotteries for a population of online workers. We document that list elicitation significantly diminishes risk aversion compared to binary choice elicitation. We show that this observation is consistent with a decision maker who has non-expected utility preferences, but when list elicitation is employed, reduces the compound lottery induced by her choices and the external randomization device used to determine payment. As a result, list elicitation distorts the inferences that can be drawn about non-expected utility preferences.

Preface

All chapters of this thesis are original and unpublished work. I am the sole author of Chapters 2 and 3 of this thesis.

Chapter 4 is coauthored with Yoram Halevy and Terri Kneeland, and is a part of a larger collaborative project. The experiment was jointly designed by all three authors, and the interpretation of our results and editing the paper were also collaborative efforts. The online interface was designed and implemented by Terri Kneeland. The experiments were run by myself and Terri Kneeland. I was responsible for data analysis and the details of the theoretical explanation of our results, and was primarily responsible for writing the paper. Grant applications, funding, and ethics approval were the sole responsibility of Yoram Halevy. Ethics approval under the project title "Mechanical Turk - Decision Theory 1" approved by the Behavioural Research Ethics Board of the University of British Columbia (certificate number H11-01719).

Table of Contents

Appendices

List of Tables

Acknowledgements

Numerous faculty members at UBC have been extremely helpful at different times during my years as both an MA and PhD student. I'm extremely grateful to my professors and classmates for their support on my journey through graduate school. I would like to especially thank my committee members Mike Peters and Li Hao for their suggestions and feedback.

I would like to thank my advisor, Yoram Halevy. Yoram dedicated a substantial amount of time to helping me develop as an economist, and struck a fine balance between providing encouragement and pushing me to raise the bar for myself.

My mother and father, Cathy and Michael Freeman, have always encouraged and supported my education from day one all the way through to my PhD. I'm incredibly lucky to have had such support.

To Luba: we made it. I am thankful for your love and support, for your having faith in me and in us.

I graciously acknowledge financial support from SSHRC in the form of a Canada Graduate Scholarship.

Chapter 1

Introduction

Many important economic decisions involve elements of risk. Examples include insurance, portfolio, and occupational choices. Thought experiments and laboratory experiments in economics and psychology have identified robust evidence of behaviour that is not consistent with economists' workhorse model - expected utility. Three particularly robust findings are that people make decisions that: (i) depend on reference points (Kahneman and Tversky, 1979), (ii) violate the Independence Axiom (Allais, 1953), and (iii) exhibit risk aversion over small stakes (Rabin, 2000). Yet the challenge of incorporating this evidence into economic models has raised issues that remain unsolved.

Reference-dependence

Experiments have identified that choice behaviour can depend on how alternatives can compare to a reference point both in risky and riskless choice tasks (Kahneman and Tversky, 1979; Kahneman, Knetsch, and Thaler, 1990; Tversky and Kahneman, 1991). Kahneman and Tversky's (1979) prospect theory, its extension to riskless environments (Tversky and Kahneman, 1991), and subsequent work (Masatlioglu and Ok, 2005; Kőszegi and Rabin, 2006; Sagi, 2006) show that reference-dependence can be incorporated in economic models of decision-making. These models drop the standard assumption that a decision-maker has a single complete and transitive preference relation and instead allow the decision-maker's preferences to depend on her reference point. However, reference points are not a directly observable economic variable. In lab experiments, the researcher is able to control the framing of decisions in a way that would induce subjects' reference points. However, researchers cannot do this in natural economic environments. This poses a challenge to the testing and use of these models in real economic environments, or in lab settings where the reference point is not explicitly primed.

The traditional assumption in applying reference-dependent models is that a decision-maker's status quo determines her reference point, as in prospect-theoretic models that were designed to explain behaviour in lab experiments. But Kahneman and Tversky (1979, p. 286) acknowledge the limitations of this approach in real economic settings:

> "So far in this paper, gains and losses were defined by the amounts of money that are obtained or paid when a prospect is played, and the reference point was taken to be the status quo, or one's current assets. Although this is probably true for most choice problems, there are situations in which gains and losses are coded relative to an expectation or aspiration level that differs from the status quo. For example, an unexpected tax withdrawal from a monthly pay check is experienced as a loss, not as a reduced gain."

Recently, Kőszegi and Rabin (2006) proposed a model of expectations-based reference-dependent decision-making, consistent with Kahneman and Tversky's suggestion that expectations may determine reference points in environments where a decision-maker expects her status-quo to change. However, expectations are not observed in standard economic data, and so the critique remains: models of reference-dependence in general, and expectations-based reference-dependence in particular, allow the decision-maker's utility function to depend on unobserved reference points. Do these models have testable implications, or can they explain almost anything?

The Independence Axiom

The normatively-appealing Independence Axiom is the substantive axiom that gives complete, transitive, and continuous preferences an expected utility representation. But a set of thought experiments by Allais (1953) suggests that many people would fail to satisfy the Independence Axiom. The central feature of Allais' thought experiments is a distinction between choice between a pair of risks, and choice when a certain reward is available.

The decision theory literature in economics responded to the "Allais paradox" by relaxing the Independence Axiom to accommodate the choice patterns suggested by Allais. Such non-expected utility models include rank-dependent utility (Quiggin, 1982; Yaari, 1987), weighted utility (Chew, 1983), and disappointment aversion (Gul, 1991).[1]

In models that do not satisfy the Independence Axiom, there are multiple ways to model how a decision-maker treats multiple sources of risk. When applying non-expected utility models, it is natural to view multiple risks as forming a multi-stage compound lottery in the sense of Segal (1990). Suppose we take a non-expected utility preference over single-stage lotteries as a starting point. Segal suggests two ways to extend single-stage lottery preferences to the domain of compound lotteries, the Reduction of Compound Lotteries and Compound Independence axioms. Under the Reduction of Compound Lotteries axiom, a decision-maker behaves as-if she reduces a compound lottery to its probabilistically-equivalent single-stage lottery, and then applies her lottery preferences to evaluate the resulting single-stage lottery. Under the Compound Independence Axiom, a decision-maker behaves as-if she evaluates a compound lottery recursively, that is, she first applies her lottery preferences to evaluate the certainty equivalent of each second-stage lottery, and then applies her lottery preferences again to the resulting single-stage lottery.

There is evidence for and against both axioms. The Reduction of Compound Lotteries axiom is used in Karni and Safra's (1987) explanation of preference reversals in the Becker, DeGroot, and Marschak (1964) mechanism for eliciting certainty equivalents, and a failure of Compound Independence is central to their argument. Segal (1990) shows that the Compound Independence is consistent with Kahneman and Tversky's isolation effect. Halevy (2007) finds additional evidence against the Reduction of Compound Lotteries axiom. Segal (1990) shows that preferences on single-stage lotteries satisfy both assumptions simultaneously if and only if they are expected utility. This discussion suggests that in some environments, the Compound Independence axiom is the more descriptively appropriate assumption, while in other

[1] I use the term "non-expected utility" to refer to these axiomatic models, as distinct from psychologically-derived models like Kahneman and Tversky's (1979) prospect theory which make the more radical departure of defining utility over gains and losses, rather than over final outcomes.

environments the Reduction of Compound Lotteries axiom may be more appropriate.

Small-stakes risk aversion and the Rabin critique

Thought experiments (Samuelson, 1963; Rabin, 2000), lab experiments, and insurance choices suggests that people demonstrate quantitatively significant risk aversion over stakes of tens and hundreds of dollars. For example, most people would turn down a gamble with a 50% chance of winning $11, and a 50% chance of losing $10. This choice pattern seems reasonable, and on its own does not contradict any axioms of expected utility.

However, suppose that a decision-maker would make the same choice even if she were much richer or much poorer. If she would turn down this gamble at any possible wealth level, Rabin shows that if she has expected utility preferences she would also turn down *any* gamble with a 50% chance of losing $100, even if it involved a 50% chance of winning an infinite amount of money. This level of risk aversion is descriptively implausible.

In expected utility, risk aversion is captured in a concave utility-for-wealth function. Rabin suggests that either reference-dependent models or non-expected utility models could address his critique of expected utility. In reference-dependent models that incorporate loss aversion and narrow bracketing, the decision-maker weighs losses more than equal-sized gains in evaluating a gamble, and ignores all other risks in her economic environment. Loss aversion can generate substantial small-stakes risk aversion over gambles that involve both gain and loss outcomes; probability weighting (as in prospect theory) can generate small-stakes risk aversion over gambles that involve only gains. Popular non-expected utility models, including rank-dependent utility and disappointment aversion, can also display small-stakes risk aversion without being susceptible to Rabin's argument, since in these models risk aversion over small-stakes comes primarily from non-linear probability weighting.

Because most people face substantial pre-existing risks, and the decision to take or turn down an offered small-stakes gamble involves multiple sources of risk. As previously discussed, reference points are not directly observable, and reference-dependent models require assumptions about choice bracketing. Non-expected utility models avoid the former issue, and allow for two disciplined approaches to the latter issue. Multiple sources of risk are naturally modelled as a compound lottery, and either the Reduction of Compound Lotteries axiom or the Compound Independence axiom can tractably extend the domain of non-expected utility models to compound lotteries. As reviewed previously, it is not a-priori obvious which assumption is the more natural descriptive assumption in real economic decisions.

Safra and Segal (2008) show that Rabin's argument extends to smooth versions of non-expected utility; but most commonly applied non-expected utility models, including rank-dependent utility and disappointment aversion, do not satisfy their smoothness assumption. However, Safra and Segal offer an additional result that shows that that in the presence of background risk, a major class of non-expected utility preferences are subject to a Rabin critique. While their result applies to rank-dependent utility and disappointment aversion, their argument implicitly relies on the assumption that a decision-maker satisfies the Reduction of Compound Lotteries axiom when evaluating the compound lottery formed by a gamble and her pre-existing risks.

Thesis outline

Chapter 2 of this thesis provides behavioural foundations for models of expectations-based reference-dependence (à la Kőszegi and Rabin (2006)). When restricted to axioms and model restrictions apply equally in environments with and without risk, each commonly-used model of preferred personal equilibrium decision-making of Kőszegi and Rabin (2006) cannot be distinguished from a versions of the shortlisting model of Manzini and Mariotti (2007). The analysis shows that environments with risk provide the natural testing ground for models of expectations-based reference-dependence. The chapter offers three new axioms in environments with risk that are conceptually motivated by expectations-based reference-dependence. Theorem 2.1 uses these axioms to provide a tight characterization of the preferred personal equilibrium model of expectations-based reference-dependence with functional form restrictions that are natural to environments with risk. Related results characterize alternative models of expectations-based reference-dependence in environments with and without risk. An additional contribution of the analysis is a choice-based definition of expectations-dependence based on a type of violation of the Independence Axiom.

Chapter 3 of this thesis revisits the problem of risk-taking in the presence of background risk under non-expected utility preferences over single-stage lotteries. The analysis assumes that a decision-maker recursively evaluates the compound lottery formed by a gamble offered on top of pre-existing risks. The results of this chapter show that this assumption allows non-expected utility models like rank-dependent utility and disappointment aversion to simultaneously produce descriptively reasonable levels of risk aversion over small and large stakes for a decision-maker who faces background risk. Indeed, a decision-maker with recursive non-expected utility preferences over compound lotteries behaves *as-if* she brackets narrowly over small-stakes, a result that is made precise in Theorems 3.1 and 3.2. These positive results stand in stark contrast to Safra and Segal's (2008) negative results assuming reduction of compound lotteries.

List elicitation has become a popular experimental design for eliciting precise information about risk preferences. In list elicitation, a subject makes a list of binary choices between lotteries, and one of those binary choices is randomly selected and the choice played for real to determine the subject's payment from the list. This experimental design makes payment the result of a compound lottery. If subjects evaluate this compound lottery recursively, then they answer each question in the list the same way they would answer if they made a single binary choice that determined their payment for sure. However, if subjects have non-expected utility preferences and do not evaluate compound lotteries recursively, perhaps because they reduce compound lotteries, then the design of list elicitation will lead subjects to make different binary choices than if they faced a single choice. Chapter 4 visits this problem of experimental design for eliciting risk preferences. The main result of this chapter is that list elicitation sigificantly affects subjects' choices. This result suggests that it is difficult to draw unambiguous inferences about risk preferences from data elicited using list elicitation.

Chapter 2

Revealed Preference Foundations of Expectations-Based Reference-Dependence

2.1 Introduction

Seminal work by Kahneman and Tversky introduced psychologically and experimentally motivated models of *reference-dependence* to economics. A limitation preventing the adoption of reference-dependent models is that reference points are not a directly observable economic variable. Kahneman and Tversky (1979) acknowledge that while it may be natural to assume that a decision-maker's status quo determines her reference point in their experiments, it is not appropriate in many interesting economic environments. The lack of a generally applicable model of reference point formation in economic environments has hindered applications of reference-dependence to economic settings.

Kőszegi and Rabin (2006) propose a model in which a decision-maker's recently-held expectations determine her reference point. Their solution concept for endogenously determined reference points has made their model convenient in numerous economic applications, including risk-taking and insurance decisions, consumption planning and informational preferences, firm pricing, short-run labour supply, labour market search, contracting under both moral hazard and adverse selection, and domestic violence.[2] In many of these applications, observed behaviour that appears im-

[2]Kőszegi and Rabin (2007); Sydnor (2010); Kőszegi and Rabin (2009); Heidhues and Kőszegi (2008, Forthcoming); Karle and Peitz (2012); Crawford and Meng (2011); Abeler, Falk, Götte, and Huffman (2011); Pope and Schweitzer (2011); Eliaz and Spiegler (2013); Herweg, Muller, and Weinschenk (2010); Carbajal and Ely (2012); Card and Dahl (2011).

possible to explain using standard models naturally fits the intuition of expectations-based reference-dependence.

Little is known about the testable implications of expectations-based reference-dependence in more general settings in spite of the large number of applications. It has been suggested that models of expectations-based reference-dependence may have no meaningful revealed preference implications, and that their success comes from adding in an unobservable variable, the reference point, used at the modeller's discretion (Gul and Pesendorfer, 2008). The results here confront this claim: models of expectations-based reference-dependence do have economically meaningful and testable implications for standard economic data. The revealed preference axioms of this paper completely summarize the implications of a widely-applied version of the model.

The main contribution of this paper is to provide a set of revealed preference axioms that constitute necessary and sufficient conditions for a model of expectations-based reference-dependence. Commonly-used cases of Kőszegi and Rabin's model are special cases of the model studied here. The revealed preference axioms clarify how the model can be tested against both the standard rational model and against alternative behavioural theories.

As in existing models of reference-dependence, behaviour is consistent with maximizing preferences conditional on the decision-maker's reference point. The main challenge of the analysis is that expectations are not observed in standard economic data. Under expectations-based reference-dependence, the interaction between optimality given a reference point and the determination of the reference point as rational expectations can generate behaviour that appears unusual since expectations are not observed. Axioms justified by the logic of expectation-dependent decisions are shown to summarize the testable content of this unusual behaviour.

2.1.1 Background: expectations-based reference-dependence

The logic of reference-dependence suggests that rather than using a single utility function, a reference-dependent decision-maker has a set of reference-dependent utility functions. The utility function $v(\cdot|r)$ defines the decision-maker's utility function given reference lottery r. When the reference lottery r is observable, as in the case where a decision-maker's status quo is her referent, standard techniques can be applied to study $v(\cdot|r)$. But when the reference lottery is determined endogenously and is unobserved, as in the case where the reference lottery is determined by the decision-maker's recent expectations, an additional modelling assumption is needed. To that end, Kőszegi and Rabin (2006) introduce two solution concepts - personal equilibrium and preferred personal equilibrium - that capture the endogenous determination of the reference lottery for models with expectations as the reference lottery.

In an environment in which a decision-maker faces a fully-anticipated choice set D, rational expectations require that the decision-maker's reference lottery corresponds with her actual choice from D. In such an environment, the set of *personal equilibria* of D provides a natural set of predictions of a decision-maker's choice from a set D:

$$PE_v(D) = \{p \in D : \ v(p|p) \geq v(q|p) \ \forall q \in D\} \tag{2.1}$$

The personal equilibrium concept has the following interpretation. When choosing from choice set D, a decision-maker uses her reference-dependent preferences $v(\cdot|r)$ given her reference lottery (r) and chooses $\arg\max_{p \in D} v(p|r)$. When forming expectations, the decision-maker recognizes that her expected choice p will determine the reference lottery that applies when she chooses from D. Thus, she would only expect a $p \in D$ if it would be chosen by the reference-dependent utility function $v(\cdot|p)$, that is, if $p \in \arg\max_{q \in D} v(q|p)$. The set of personal equilibria of D in (2.1) is the set of all such p.

There may be a multiplicity of personal equilibria for a given choice set. Indeed, if reference-dependence tends to bias a decision-maker towards her reference lottery, multiplicity is natural. At the time of forming her expectations, a decision-maker evaluates the lottery p according to $v(p|p)$, which reflects that she will evaluate outcomes of lottery as gains and losses relative to outcomes of p itself. The *preferred personal equilibrium* concept is a natural refinement of the set of personal equilibria based on a decision-maker picking her best personal equilibrium expectation according to $v(p|p)$:

$$PPE_v(D) = \arg\max_{p \in PE(D)} v(p|p) \tag{2.2}$$

Kőszegi and Rabin (2006) adopt a particular functional form for v. They assume that given probabilistic expectations summarized by the lottery r, a decision-maker ranks a lottery p according to:

$$v^{KR}(p|r) = \sum_k \sum_i p_i m^k(x_i^k) + \sum_k \sum_i \sum_j p_i r_j \mu\left(m^k(x_i^k) - m^k(x_j^k)\right) \tag{2.3}$$

In (2.3), m^k is a consumption utility function in "hedonic dimension" k; different hedonic dimensions are akin to different goods in a consumption bundle, but specified based on "psychological principles". The function μ is a gain-loss utility function which captures reference-dependent outcome evaluations.

The Kőszegi-Rabin model with the preferred personal equilibrium concept has been particularly amenable to applications, since the model's predictions are pinned down by (2.3) and (2.2). However, little is known about how the Kőszegi-Rabin model behaves except in very specific applications.

This paper focuses on expectations-based reference-dependent preferences with the preferred personal equilibrium concept as in (2.2). Theorem 2.1 provides a complete revealed preference characterization of the choice correspondence c that equals the set of all preferred personal equilibria of a choice set, $c(D) = PPE_v(D)$. The model of decision-making equivalent to the axioms does not restrict v to the form in (2.3) but does require that v be jointly continuous in its arguments, $v(\cdot|r)$ satisfy expected utility, and v satisfy a property related to disliking mixtures of lotteries.

The tight characterization of the PPE model of expectations-based reference-dependence in Theorem 2.1 may come as a surprise relative to previous work (e.g. Gul and Pesendorfer 2008; Kőszegi 2010).[3] The analysis here also provides additional surprising connections. First, the PPE representation is related to the shortlisting representation of Manzini and Mariotti (2007), a connection clarified in Proposition 2.2.

[3]Gul and Pesendorfer (2008) show that with the personal equilibrium concept and without using any lottery structure, the reference-dependent preferences of Kőszegi and Rabin (2006) have no testable implications beyond an equivalence with a choice correspondence generated by a binary relation. Kőszegi (2010) initially proposed the personal equilibrium concept studied here but provides only a limited set of testable implications, and suggested that a complete revealed preference may not be possible: "I do not offer a revealed-preference foundation for the enriched preferences—it is not clear to what extent the decisionmaker's utility function can be extracted from her behavior."

Second, there is a tight connection between expectations-based reference-dependence and failures of the Mixture Independence Axiom; violations of Independence of Irrelevant Alternatives (IIA) are sufficient but not necessary for expectations-dependent behaviour in the model (Proposition 2.3).

2.1.2 Outline

Section 2.2 provides two examples that motivate expectations-based reference-dependence, and a result that illustrates the limits to the model's testable implications in environments without risk. Section 2.3 provides axioms and a representation theorem for PPE decision-making, and suggests a way of defining expectations-dependence in terms of observable behaviour. Section 2.4 explores special cases of the model, including Kőszegi-Rabin and a new axiomatic model of expectations-based reference lottery bias. Section 2.5 shows how the analysis can be adapted to study PE decision-making and also to decision-making under Kőszegi and Rabin's (2007) choice-acclimating personal equilibrium (CPE).

2.2 Two examples and a motivating result

2.2.1 Formal setup

Let Δ denote the set of all *lotteries* with support on a given finite set X, with typical elements $p, q, r \in \Delta$. Let $\mathcal{D}$ denote the set of all finite subsets of Δ, a typical $D \in \mathcal{D}$ is called a *choice set*. The starting point for analysis is a *choice correspondence*, $c : \mathcal{D} \to \mathcal{D}$, which is taken as the set of elements we might observe a decision-maker choose from a set D. Assume $\emptyset \neq c(D) \subseteq D$, that is, a decision-maker always chooses something from her choice set.

Define the mixture operation $(1-\lambda)D + \lambda D' := \{(1-\lambda)p + \lambda q : p \in D, q \in D'\}$.

2.2.2 Mugs, pens, and expectations-based reference-dependence

The classic experimental motivation for loss-aversion in riskless choice comes from the *endowment effect*. An example of an endowment effect comes from the experimental finding that randomly-selected subjects given a mug have a median willingness-to-accept for a mug that is double the median willingness-to-pay of subjects who were not given a mug (Kahneman, Knetsch, and Thaler, 1990). This classic experiment provides no separation between status-quo-based and expectations-based theories of reference-dependence since subjects given a mug could expect to be able to keep it at the end of the experiment.

To separate expectations-based theories of reference-dependence from status-quo based theories, Ericson and Fuster (2011) design an experiment in which all subjects are endowed with a mug, and subjects are told that there is a fixed probability (either 10% or 90%) they will receive their choice between a retaining the mug or instead obtaining a pen, and with the remaining probability they will retain the mug; the conditional choice must be made before uncertainty is resolved.[4] Subjects in a treatment with a 10% chance of receiving their choice must expect to receive a mug with at least a 90% chance, and consistent with expectations-based reference-dependence, 77% of these subjects' conditionally choose the mug. In contrast, only 43% of subjects conditionally choose the mug in the treatment in which subjects received their chosen item with a 90% chance.

The Mixture Independence axiom below adapts of von-Neuman and Morgenstern's axiom to a choice correspondence.

Mixture Independence. $(1-\alpha)c(D) + \alpha c(D') = c((1-\alpha)D + \alpha D')\ \forall \alpha \in (0,1)$

The median choice pattern in Ericson and Fuster's experiment has $\{\langle \text{mug, } 1\rangle\} = c(.9\{\langle \text{mug, } 1\rangle\} + .1\{\langle \text{mug, } 1\rangle, \langle \text{pen, } 1\rangle\})$ but $\{\langle \text{mug, } .1; \text{ pen, } .9\rangle\} = c(.1\{\langle \text{mug, } 1\rangle\} +$

[4] This paper interprets the subjects' choice as being between two lotteries, each of which involves the prize of the mug with a fixed probability (10% or 90%) and the prize chosen by the subject with the remaining probability. An alternative interpretation of the experimental setup is that subjects face a lottery over choice sets, one of which is a singleton, and must choose from the non-singleton choice set before the lottery is resolved. For a result on the formal relationship between these choice spaces, see Ortoleva (2013).

Table 2.1: Example of reference-dependent preferences

	$v(p\|\cdot)$	$v(q\|\cdot)$	$v(r\|\cdot)$
$v(\cdot\|p)$	1000	900	1050
$v(\cdot\|q)$	-1350	0	-75
$v(\cdot\|r)$	-1575	-450	-262.50

$.9\{\langle \text{mug}, 1\rangle, \langle \text{pen}, 1\rangle\})$. This choice pattern suggests an intuitive and empirically supported violation of Mixture Independence that is consistent with expectation-bias.

2.2.3 IIA violations under Kőszegi-Rabin under PPE

Consider a decision-maker with a Kőszegi-Rabin v as in (2.3), with linear utility and linear loss aversion:[5][6]

$$m(x) = x, \quad \mu(x) = \begin{cases} x & \text{if } x \geq 0 \\ 3x & \text{if } x < 0 \end{cases}$$

When faced with a set of lotteries, suppose that our decision-maker chooses his preferred personal equilibrium lottery as in (2.2).

Consider the three lotteries $p = \langle \$1000, 1\rangle$, $q = \langle \$0, .5; \$2900, .5\rangle$, and $r = \langle \$0, .5; \$2000, .25; \$4100, .25\rangle$. As broken down in Table 2.1, the decision-maker's choice correspondence, c, is given by $\{p\} = c(\{p, q\})$, $\{q\} = c(\{q, r\})$, $\{r\} = c(\{p, r\})$, and $\{q\} = c(\{p, q, r\})$.

Choice from binary sets reveals an intransitive *cycle*. Because of this, there is no possible choice from $\{p, q, r\}$ is consistent with preference-maximization! Consider the Independence of Irrelevant Alternatives (IIA) axiom below, which Arrow (1959) shows is equivalent to maximization of a complete and transitive preference relation.

IIA. $D' \subset D$ and $c(D) \cap D' \neq \emptyset \implies c(D') = c(D) \cap D'$.

[5] I would like to specially thank Matthew Rabin for suggesting this example.

[6] Linear loss aversion is used in most applications of Kőszegi-Rabin, and the chosen parameterization is broadly within the range implied by experimental studies.

In the Kőszegi-Rabin PPE example, adding the lottery r to the set $\{p, q\}$ generates a violation of IIA, since r is not chosen yet affects choice from the larger set. Given fixed expectations r, our decision-maker's behaviour would be consistent with the standard model: she would maximize $v(\cdot|r)$. The decision-maker exhibits novel behaviour because her expectations, and hence preferences, are determined endogenously in a choice set. However, rational expectations combined with preferred personal equilibrium put quite a bit of structure on the decision-maker's novel behaviour. The axiomatic analysis that follows will clarify the nature of such structure.

2.2.4 The testable implications of Kőszegi-Rabin under PE: a negative result

The preceding example demonstrates that the Kőszegi-Rabin model with PPE generates choice behaviour that cannot be rationalized by a complete and transitive preference relation. Gul and Pesendorfer (2008) suggest that compared to the standard rational model, this may be the *only* revealed preference implication of the Kőszegi-Rabin model when paired with the personal equilibrium solution criteria in (2.1). Gul and Pesendorfer take as a starting point a finite set X of riskless elements, a reference-dependent utility $v : X \times X \to \Re$, and offer the following result:

Proposition 2.1. (Gul and Pesendorfer 2008). *The following are equivalent: (i) c is induced by a complete binary relation, (ii) there is a v such that $c(D) = PE_v(D)$ for any choice set D, (iii) there is a v that satisfies (2.3) such that $c(D) = PE_v(D)$ for any choice set D.*

Proof. (partial sketch)

If $c(D) = \{x \in D : xRy \, \forall y \in D\}$ then define v by: $v(x|x) \geq v(y|x)$ if xRy, and $v(y|x) > v(x|x)$ otherwise. Then, $\{xRy \; \forall y \in D\} \iff \{v(x|x) \geq v(y|x) \; \forall y \in D\}$. By reversing the process, we could construct R from v. Thus (i) holds if and only if (ii) holds.

Gul and Pesendorfer cite Kőszegi and Rabin's (2006) argument that the set of hedonic dimensions in a given problem should be specified based on "psychological principles". Since X has no assumed structure, Gul and Pesendorfer infer hedonic dimensions from c and the structure imposed by (2.3). Their construction shows any v has a representation in terms of the functional form in (2.3).

□

The analysis that follows uses two assumptions that allow for a rich set of testable implications of expectations-based reference-dependence. First, c is defined on a subsets of *lotteries* over a finite set. The structure of lotteries in choice sets places additional observable restrictions on expectations in a choice set and additional information on behaviour relative to expectations. New axioms make particular use of this lottery structure to trace the observable implications of expectations-based reference-dependence.

Second, the main analysis looks for the revealed preference implications of *preferred* personal equilibrium. The sharper predictions of preferred personal equilibrium lead to different testable implications of the PPE based model expectations-based reference-dependence in the absence of risk.

This choice space does not allow the analysis to say anything insightful about the set of hedonic dimensions of the problem. In light of Gul and Pesendorfer's result, the representation here does not seek any particular structure on the v that represents reference-dependent preferences. The analysis considers the particular structure imposed by the functional form (2.3) as a secondary issue for future work.

2.3 Revealed preference analysis of PPE

2.3.1 Technical prelude

Define distance on lotteries using the Euclidean distance metric, $d^E(p,q) := \sqrt{\sum_i (p_i - q_i)^2}$, and the distance between choice sets using the Hausdorff metric,
$d^H(D,D') := \max\left(\max_{p\in D}\left[\min_{q\in D'} d^E(p,q)\right], \max_{q\in D'}\left[\min_{p\in D} d^E(p,q)\right]\right)$.

It will be useful to offer a few definitions in advance of the analysis. For any set T with typical element t, let $\{t^\epsilon\}$ denote a *convergent net* indexed by a set $(0,\bar{\epsilon}]$ and with limit point t; t^ϵ will be used to denote the ϵ term in the net.[7] Define

[7] A *net* in a set T is a function $t : S \to T$ for some directed set S (Aliprantis and Border, 1999).

$c^U(D)$ as the upper hemicontinuous extension of c; that is, $c^U(D) := \{p \in D : \exists\{p^\epsilon, D^\epsilon\}$ such that $p^\epsilon \in c(D^\epsilon)$, $p^\epsilon \to p$, $D^\epsilon \to D\}$. For $p \in \Delta$ and $\epsilon > 0$, let $N_p^\epsilon := \{p^\epsilon \in \Delta : \ d^E(p, p^\epsilon) < \epsilon\}$ denote a ϵ-neighbourhood of p. For any binary relation R, let cl$R$ denote its closure, defined by: $p(\mathrm{cl}R)q$ if $\exists\{p^\epsilon\} \to p$, $\{q^\epsilon\} \to q$ such that for each $\epsilon > 0$, $p^\epsilon R q^\epsilon$. For any finite set D and binary relation R, define $m(D,\, R) := \{p \in D : \nexists q \in D$ such that qRp but not $pRq\}$ as the set of undominated elements in D according to binary relation R.

2.3.2 Revealed preference analysis without risk

Ignoring restrictions specific to risks, the classic IIA axiom provides the point of departure from standard models. The two axioms below allow for failures of IIA that can arise from the endogenous determination of expectations and preferences in each choice set. For this section, restrict attention to axioms and restrictions on the representation in (2.2) that do not make use of the particular economic structure of lotteries, except for the continuity of Δ.

The following Expansion axiom is due to Sen (1971).

Expansion. $p \in c(D) \cap c(D') \implies p \in c(D \cup D')$

Expansion says that if a lottery p is chosen in both D and D' then it is chosen in $D \cup D'$. This seems weak as both a normative and a descriptive property, and is an implication of variations on the Weak Axiom of Revealed Preference (see Sen (1971)). Expansion rules out the attraction and compromise effects, in which an agent chooses p over both q and r in pairwise choices, but chooses q from $\{p,\, q,\, r\}$.[8] In the attraction effect, r is similar to, but dominated by q and attracts the decision-maker to p in $\{p,\, q,\, r\}$; in the compromise effect, q is a compromise between more extreme options p and r in the choice set $\{p,\, q,\, r\}$.

The Weak RARP (RARP for Richter's (1966) Axiom of Revealed Preference[9]) is in the spirit of the classic axioms of revealed preference (like WARP, SARP, and

[8]See Simonson (1989) for evidence on attraction and compromise effects. Ok, Ortoleva, and Riella (2012) provide a model of the attraction effect that captures this phenomenon.

[9]Richter refers to his axiom as "Congruence". I use RARP to emphasize the close connection with WARP, SARP, GARP, etc. For more on the connection between these axioms, see Sen (1971).

GARP) albeit with an embedded continuity requirement. In particular, the axiom weakens (a suitably continuous version of) RARP.

Define $p\tilde{\tilde{R}}q$ if $p \in c(D)$ and $q \in c^U(\bar{D})$ for some $D, \bar{D}$ with $\{p,q\} \subseteq D \subseteq \bar{D}$. The relation $\tilde{\tilde{R}}$ is defined whenever sometimes p is chosen when q is available, and sometimes q is choosable (in the sense that $q \in c^U(\bar{D})$) when p is available. The statement $p\tilde{\tilde{R}}q$ holds when p is weakly chosen over q in a smaller set, but q is weakly choosable over p in a set that is larger in the sense of set inclusion. Define $p\tilde{\tilde{W}}q$ if there exist $p^0 = p,\ p^1, ..., p_{n-1}, p_n = q$ such that $(p^{i-1}, p^i) \in \mathrm{cl}\tilde{\tilde{R}}$ for $i = 1, ..., n$. That is, $\tilde{\tilde{W}}$ is the continuous and transitive extension of $\tilde{\tilde{R}}$.

Weak RARP. $p \in c(D),\ q \in c^U(\bar{D}),\ q \in D \subseteq \bar{D},$ and $q\tilde{\tilde{W}}p \implies q \in c(D)$

The crucial implication of Weak RARP is captured by its main economic implication, *Weak WARP*: if $p = c(\{p,q\})$ and $p \in c(D)$ then $q \notin c(D')$ whenever $p \in D' \subseteq D$.[10] Manzini and Mariotti (2007) offer an interpretation in terms of constraining *reasons*: an agent might choose p over q in a smaller set, like $\{p,q\}$, yet might have a constraining reason against choosing p in a larger set D. However, if we observe p chosen from a large set D, then any D' that is a subset of D contains no constraining reason against choosing p. Thus, her choice in D' should be minimally consistent with her choice in $\{p,q\}$ and she should not choose q.

Weak RARP strengthens the logic of Weak WARP in two ways. Weak WARP allows only WARP violations consistent with the existence of constraining reasons, and takes choices from smaller sets - which can fewer constraining reasons - as the determinant of choice in the absence of constraining reasons. The main way Weak RARP strengthens Weak WARP is by imposing that choice among unconstrained options is determined by a transitive procedure.[11]

Weak RARP as stated also strengthens a transitive version of Weak WARP by imposing continuity in two ways. Taking the topological closure of $\tilde{\tilde{R}}$ and then taking the transitive closure imposes that choice among unconstrained options is determined

[10] The following proof that Weak RARP implies Weak WARP may help clarify the connection. Suppose $p \in c(D)$, $p \in D' \subset D$, and $q \in c(D')$. Then $q\tilde{\tilde{W}}p$, and so if $p \in c(\{p,q\})$, Weak RARP implies that $q \in c(\{p,q\})$ as well. Thus Weak RARP implies that if $p = c(\{p,q\})$ and $p \in c(D)$ hold, $q \in c(D')$ could not hold.

[11] In this regard, Weak RARP is closely related to the "No Binary Cycle Chains" axiom of Cherepanov, Feddersen, and Sandroni (Forthcoming).

by a rationale that is both transitive and continuous. This imposes a restriction that is economically natural relative to the topological structure of lotteries. The second continuity aspect of Weak RARP is that if $p \in c^U(D)$, p is seen as chooseable from D. That is, if it is revealed that there is no reason to reject p^ϵ from D^ϵ when p^ϵ and D^ϵ are 'arbitrarily close' to p and D respectively, then Weak RARP assumes that there is no reason revaled to reject p from D (even if p is not chosen at D). These two strengthenings in Weak RARP are natural given the topological structure of the space of lotteries (and many other choice spaces).

Formally, say that a PPE representation in (2.2) is *continuous* if v is jointly continuous. Proposition 2.2 (i) $\iff$ (ii), clarifies the link between the Expansion and Weak RARP axioms on one hand, and the PPE decision-making on the other hand.

Manzini and Mariotti (2007) characterize a *shortlisting representation*, $c(D) = m(m(D,P_1),P_2)$ for two asymmetric binary relations P_1, P_2, in terms of two axioms, Expansion and Weak WARP.[12] If P_2 is transitive and both P_1 and P_2 are continuous, say that P_1, P_2 is a *continuous and transitive shortlisting representation.*[13] Proposition 2.2 (ii) $\iff$ (iii), provides a link between a version of the shortlisting model of Manzini and Mariotti and the PPE representation in (2.2).

Proposition 2.2. *(i)-(iii) are equivalent: (i) c satisfies Expansion and Weak RARP, (ii) c has a continuous PPE representation, (iii) c has a continuous and transitive shortlisting representation.*

Proof. (ii) $\iff$ (iii)

Consider the following mapping between a continuous PPE representation v and a continuous and transitive shortlisting representation:

$v(q|p) > v(p|p) \iff qP_1p$

$v(p|p) > v(q|q) \iff pP_2q$

For v and $P_1,\ P_2$ that satisfy this mapping, $m(D,\ P_1) = PE_v(D)$, and $m(m(D,\ P_1),\ P_2) = PPE_v(D)$.

[12] Manzini and Mariotti (2007) and follow-up papers assume that c is a single-valued choice function, which simplifies their analysis.

[13] This terminology is different from Au and Kawai (2011) and Horan (2012) who discuss shortlisting representations in which both P_1 and P_2 are transitive.

It remains to verify that joint continuity in v is equivalent to continuity of P_1 and P_2 - the full argument is in the appendix. □

The v in a PPE representation characterized by Proposition 2.2 is highly non-unique: *any* $\hat{v}$ that satisfies $\hat{v}(q|p) > \hat{v}(p|p) \iff v(q|p) > v(p|p)$ and has $\hat{v}(p|p) = u(p)$ for *some* u that represents P_2 in the shortlisting representation also represents the same c. Put another way, v includes information about how a decision-maker would choose between any two lotteries p and q given any reference lottery r. However, if the decision-maker's rational expectations determine her reference lottery, as in a PPE representation, choices give us no direct information about a decision-maker would choose between p and q given any reference lottery $r \notin \{p, q\}$.

2.3.3 Revealed preference analysis with risk

The result in Proposition 2.2 did not consider the possibility of adopting stronger axioms or restrictions on v that are suitable when working with choice among lotteries but may not be economically sensible in other domains. But the evidence supporting expectations-based reference-dependence in Ericson and Fuster (2011) suggests that environments with risk provide a natural environment for studying expectations-based reference-dependence. This section explores the possibility of a stronger characterization in environments with risk.

Environments with risk enable a partial separation between expectations and choice. Suppose we view the mixture $(1-\alpha)q + \alpha D$ as arising from a lottery over choice sets that gives the singleton choice set $\{q\}$ with probability $1-\alpha$ and gives choice set D with probability α. Under this interpretation, fraction $1-\alpha$ of expectations are fixed at expecting q and we also observe the decision-maker's conditional choice from D. The three axioms below make use of variations on this interpretation.

The Induced Reference Lottery Bias Axiom uses this partial separation between expectations and choice. The axiom requires that if p is chosen in a choice set D, then p would also be conditionally chosen from D when some of the expectations are fixed at p, as in any mixture of the form $(1-\alpha)p + \alpha D$. This is a natural axiom to adopt under expectations-based reference-dependence: fixing expectations at p at least partially fixes the reference-lottery weakly towards p; if the decision-maker is biased towards her reference-lottery, this should bias her towards choosing p.

Induced Reference Lottery Bias. $p \in c(D)$ implies $p \in c((1-\alpha)p + \alpha D)$ $\forall \alpha \in (0,1)$.

Notice that Induced Reference Lottery Bias allows for the violation of Mixture Independence observed by Ericson and Fuster (2011), but rules out a violation in the opposite direction.

IIA Independence weakens the Mixture Independence Axiom to a variation that only implies a restriction on behaviour in the presence of IIA violations, with an embedded continuity requirement.

IIA Independence. If $p \in c(D)$ and $\exists \alpha \in (0,1]$ such that $p \notin c(D \cup ((1-\alpha)p + \alpha q)) \ni r$ and $p\tilde{\tilde{W}}r$, then $\exists \epsilon > 0$ such that $\forall \alpha' \in (0,1]$, $\forall \hat{p} \in N_p^\epsilon$, $\forall \hat{q} \in N_q^\epsilon$, and $\forall D' \ni (1-\alpha')\hat{p} + \alpha'\hat{q}$, $\hat{p} \notin c(D')$.

The spirit of Weak RARP is the requirement that in the absence of constraining reasons, c is consistent with maximizing $\tilde{\tilde{W}}$, derived from choice from smaller choice sets. The choice pattern $p \in c(D)$, $p \notin c(D \cup q) \ni r$, and $p\tilde{\tilde{W}}r$ then reveals that q blocks p.[14] This revealed blocking behaviour only appears when the model violates IIA. The IIA Independence axiom requires that in this case, any mixture between q and p also prevents p from being chosen from any choice set. The logic of expectations-dependence then requires that the agent would not choose p when it involves a conditional choice of p over q.

Remark 2.1. A simple test of IIA Independence that could detect behaviour inconsistent with expectations-dependence would be to find p, q, α, D with $p \in c(D)$, $\{p,q\} \cap c(D \cup q) = \emptyset$ but $p \in c(D \cup ((1-\alpha)p + \alpha q))$. Table 2.2 shows two possible choice correspondences that describe a decision-maker who finds candy too tempting to turn down for an apple whenever she had been expecting to eat but who can avoid temptation by planning in advance to abstain from snacking. Choice correspondence c captures a decision-maker who can exert limited self-control against the expectations-induced temptation to go for candy, and is inconsistent with the IIA Independence axiom. Choice correspondence $\hat{c}$ cannot exert this limited self-control, and is consistent with the axiom.

[14] In the appendix, it is shown that this choice pattern is ruled out by Weak RARP and Expansion.

Table 2.2: Testing IIA Independence

D	$c(D)$	$\hat{c}(D)$
$\{\langle\text{apple, } 1\rangle, \langle\text{don't eat, } 1\rangle\}$	$\{\langle\text{apple, } 1\rangle\}$	$\{\langle\text{apple, } 1\rangle\}$
$\{\langle\text{candy, } 1\rangle, \langle\text{apple, } 1\rangle, \langle\text{don't eat, } 1\rangle\}$	$\{\langle\text{don't eat, } 1\rangle\}$	$\{\langle\text{don't eat, } 1\rangle\}$
$\{\langle\text{apple, .9; candy, .1}\rangle, \langle\text{apple, } 1\rangle\}$	$\{\langle\text{apple, } 1\rangle\}$	$\{\langle\text{apple, .9; candy, .1}\rangle\}$

The continuity requirement embedded in IIA Independence slightly strengthens restriction on c when adding q to the choice set prevents p from being conditionally chosen. The IIA Independence axiom requires that in this case, lotteries close to p prevent lotteries close to q from being conditionally chosen as well.

Say that q is a *weak conditional choice over r given p, $q\bar{R}_p r$,* if there exists a net $\{p^\epsilon, q^\epsilon, r^\epsilon\} \to p, q, r$ such that $(1-\epsilon)p^\epsilon + \epsilon q^\epsilon \in c((1-\epsilon)p^\epsilon + \epsilon\{q^\epsilon, r^\epsilon\})$ for each ϵ. A conditional choice involves a choice between q and r for when expectations are close to p.

Transitive Limit. $q\bar{R}_p r$ and $r\bar{R}_p s \implies q\bar{R}_p s$.

If IIA violations are only driven by the behavioural influence of expectations and their endogenous determination, then the agent's behaviour should be consistent with the standard model when her expectations are fixed. The Transitive Limit axiom says that conditional choice behaviour should look like the standard model when expectations are almost fixed, although the axiom only imposes this restriction on weak conditional choices.

Remark 2.2. As with continuity axioms, the Transitive Limit axiom is not *exactly* testable. However, the axiom is *approximately* testable. The choice sets in Table 2.3 provide an approximate test of Transitive Limit; $\hat{c}$ is consistent with what we would expect if the choice correspondence satisfies Transitive Limit. However, the choice pattern displayed by c is approximately inconsistent with Transitive Limit, and suggests that c would violate this axiom.

Formally, say that a PPE representation is an *EU-PPE representation* if $v(\cdot|p)$ takes an expected utility form for any $p \in \Delta$. Say that v *dislikes mixtures* if $v(p|p) \geq v(q|p)$ and $v(q|q) \leq \max[v(p|p), v(p|q)]$ imply that $\forall\alpha \in (0,1)$, $v((1-\alpha)p + \alpha q|(1-\alpha)p + \alpha q) \leq \max[v(p|p), v(p|(1-\alpha)p + \alpha q)]$.

Table 2.3: Testing Transitive Limit

$D = .9\{\langle\text{mug, }1\rangle\} + .1$	$c(D)$	$\hat{c}(D)$
$\{\langle\text{pen, }1\rangle, \langle\text{mug, }1\rangle\}$	$\{\langle\text{mug, }1\rangle\}$	$\{\langle\text{mug, }1\rangle\}$
$\{\langle\text{candy, }1\rangle, \langle\text{mug, }1\rangle\}$	$.9\{\langle\text{mug, }1\rangle\} + .1\{\langle\text{candy, }1\rangle\}$	$\{\langle\text{mug, }1\rangle\}$
$\{\langle\text{candy, }1\rangle, \langle\text{pen, }1\rangle\}$	$.9\{\langle\text{mug, }1\rangle\} + .1\{\langle\text{pen, }1\rangle\}$	$.9\{\langle\text{mug, }1\rangle\} + .1\{\langle\text{pen, }1\rangle\}$

Theorem 2.1. *c satisfies Weak RARP, Expansion, IIA Independence, Induced Reference Lottery Bias, and Transitive Limit if and only if it has a continuous EU-PPE representation in which v dislikes mixtures.*

The full proof is in the appendix, and is discussed in the next subsection.

Corollary 2.1. *Given a continuous EU-PPE representation v for c, any other continuous EU-PPE representation $\hat{v}$ for c satisfies $\hat{v}(q|p) \geq \hat{v}(r|p) \iff v(q|p) \geq v(r|p)$ and $\hat{v}(p|p) \geq \hat{v}(q|q)$ whenever $p\tilde{\tilde{W}}q$.*

Corollary 2.1 clarifies that a continuous EU-PPE is unique in the sense that any v, $\hat{v}$ that represent the same c must represent the same reference-dependent preferences.[15] This definition of uniqueness captures that the underlying reference-dependent preferences are uniquely identified, but says nothing about the cardinal properties of reference-dependent utility functions. In an EBRD, v plays roles in both determining the set of of personal equilibria, and selecting from personal equilibria. The second part of Corollary 2.1 clarifies that this second role places a restriction that any v representing c must represent the same ranking of personal equilibria, at least when that ranking is revealed from choices.

Remark 2.3. In the representation in Theorem 2.1, any p chosen in D is (i) an element of D, and (ii) is in $\arg\max_{q\in D} v(\cdot|p)$. A more general model might allow a decision-maker to randomize among elements of her choice set. An alternative representation might have the decision-maker's reference lottery involve a randomization among elements in D, or perhaps only elements in $c(D)$. However, Theorem 2.1 proves that if c satisfies the five axioms it has a representation in which it is *as-if* the decision-maker never views herself as randomizing among elements of D.

[15] A stronger uniqueness result is possible, since (i) each $v(\cdot|p)$ satisfies expected utility and thus has an affinely unique representation, (ii) joint continuity of v in the representation restricts the allowable class of transformations of v.

2.3.4 Sketch of proof and an intermediate result

The first part of the proof takes $\bar{R}_p$ and characterizes a v such that $v(\cdot|p)$ represents $\bar{R}_p$. By Transitive Limit and because $\bar{R}_p$ is continuous by construction, such a $v(\cdot|p)$ exists. A sequence of lemmas show that the definition of $\bar{R}_p$ and Transitive Limit axiom imply the existence of a jointly continuous v such that $v(\cdot|p)$ represents $\bar{R}_p$ and satisfies expected utility.

Crucial to proof is providing a link between behaviour captured by v and behaviour in arbitrary choice sets. Consider an alternative axiom, Limit Consistency, which was not assumed in Theorem 2.1 but which would have been a reasonable axiom to adopt. First, define R_p as the asymmetric part of $\bar{R}_p$.

Limit Consistency. qR_pp implies $p \notin c(D)$ whenever $q \in D$.

The statement qR_pp says that q is always conditionally chosen over p when expectations are almost fixed at p. Limit Consistency requires that a decision-maker who always conditionally chooses q over p when her expectations are almost fixed at p would also never choose p when q is available. This is consistent with the logic of expectations-dependence. If instead qR_pp but p were chosen over q in some set D, then the decision-maker would choose p over q when her expectations are p even though she always conditionally chooses q over p when her expectations are almost fixed at p; such behaviour would be inconsistent with expectations-dependence and is ruled out.

The lemma below establishes that the axioms in Theorem 2.1 imply Limit Consistency.

Lemma. *Expansion, Weak RARP, and Induced Reference Lottery Bias imply Limit Consistency.*

The sufficiency part of the proof of Theorem 2.1 proceeds by using Expansion, Weak RARP, Limit Consistency, and v constructed from $\bar{R}_p$ to show that $c(D) = PPE_v(D)$. This gives the following intemediate result, a characterization of an EU-PPE representation in terms of Weak RARP, Expansion, IIA Independence, Limit Consistency, and Transitive Limit.

Table 2.4: Two choice correspondences

	c	$\hat{c}$
$.9\{\langle \text{pen, } 1\rangle\} + .1\{\langle \text{pen, } 1\rangle\ ,\ \langle \text{mug, } 1\rangle\}$	$\{\langle \text{pen, } 1\rangle\}$	$\{\langle \text{pen, } 1\rangle\}$
$.9\{\langle \text{mug, } 1\rangle\} + .1\{\langle \text{pen, } 1\rangle\ ,\ \langle \text{mug, } 1\rangle\}$	$\{\langle \text{mug, } 1\rangle\}$	$\{\langle \text{mug}, .9;\ \text{pen}, .1\rangle\}$

Theorem 2.2. *c satisfies Weak RARP, Expansion, IIA Independence, Limit Consistency, and Transitive Limit if and only if it has a continuous EU-PPE representation.*

Notice than in any EU-PPE representation, expected utility of $v(\cdot|p)$ and joint continuity of v will imply that $v(q|p) > v(r|p) \implies qR_pr$. With this observation in hand, the necessity of Limit Consistency follows obviously from the representation. The remainder of the proof of the above Theorem follows from the proof of Theorem 2.1.

2.3.5 A definition of expectations-dependence and its implications

Say that c exhibits *expectations-dependence* at D, α, p, q, r for $\alpha \in (0,1)$ and $p, q, r \in \Delta$ if $(1-\alpha)p + \alpha r \in c((1-\alpha)p + \alpha D)$ but $(1-\alpha)q + \alpha r \notin c((1-\alpha)q + \alpha D)$. Interpret $(1-\alpha)p + \alpha r \in c((1-\alpha)p + \alpha D)$ as involving a *conditional choice* of r from D, conditional on fraction $1-\alpha$ of expectations being fixed by p. Say that *c exhibits strict expectations-dependence* at D, α, p, q, r for $D \in \mathcal{D}$, $\alpha \in (0,1)$, and $p, q, r \in \Delta$ if there is a $\bar{\epsilon} > 0$ such that for all r^ϵ, D^ϵ pairs such that $r^\epsilon \in D^\epsilon$ and $\max\left[d^E(r^\epsilon, r), d^H(D^\epsilon, D)\right] < \epsilon$, $(1-\alpha)p + \alpha r^\epsilon \in c((1-\alpha)p + \alpha D^\epsilon)$ for all $\epsilon < \bar{\epsilon}$ but $(1-\alpha)q + \alpha r^\epsilon \notin c((1-\alpha)q + \alpha D^\epsilon)$ for all $\epsilon < \bar{\epsilon}$. This behavioural definition of expectations-dependence provides a tool for identifying and eliciting expectations-dependence, as illustrated by the example below.

Example (mugs and pens). Fix $\alpha = .1$, let $p = \langle \text{pen, } 1\rangle$; $q = \langle \text{mug, } 1\rangle$, $r = p$, and $D = \{p,\ q\}$.

Table 2.4 shows the values that two choice correspondences, c and $\hat{c}$, take on the menus $(1-\alpha)p + \alpha D = \{\langle \text{mug, } 1\rangle\ ,\ \langle \text{mug, } .9\,;\text{pen, } .1\rangle\}$ and $(1-\alpha)q + \alpha D =$

$\{\langle \text{mug}, .1; \text{ pen}, .9\rangle, \langle \text{pen}, 1\rangle\}$. Of these two choice correspondences, c exhibits expectations-dependence given D, α, p, q, r, while $\hat{c}$ does not.

$\diamond$

The definition of exhibiting expectations-dependence bears striking relation to the Mixture Independence axiom. Indeed, expectations-dependence as defined is a type of violation of Mixture Independence. Proposition 2.3 below clarifies the link between a exhibiting expectations-dependence, properties of a continuous EU-PPE representation, and violations of the IIA axiom.

Proposition 2.3. *c with a continuous EU-PPE representation strictly exhibits expectations-dependence if and only if $v(\cdot|p)$ is not ordinally equivalent to $v(\cdot|q)$ for some $p, q \in \Delta$. In addition, c with a continuous EU-PPE representation that violates IIA exhibits strict expectations-dependence.*

The first part of Proposition 2.3 highlights how expectations-dependence in c is captured in a PPE representation. There is a tight tie between expectations-dependence and failures of Mixture Independence in a PPE representation, and the second part of Proposition 2.3 shows that a failure of IIA implies, but is not necessary for, expectations-dependence.

The mugs and pens example shows how one might study expectations-dependence based on the definition. Ericson and Fuster's (2011) data violate Mixture Independence in a way consistent with expectations-based reference-dependence, and Proposition 2.3 shows that any PPE representation representing their median subject's behaviour must exhibit expectations-dependence.

2.3.6 Limited cycle property of a PPE representation

The characterization in Theorem 2.1 is tight. However, it is possible that some structure already imposed on the problem implies additional structure on v. Proposition 2.4 shows that this is indeed the case.

Say that a PPE representation satisfies the *limited cycle inequalities* if for any $p^0, p^1, ..., p^n \in \Delta$, $v(p^i|p^{i-1}) > v(p^{i-1}|p^{i-1})$ for $i = 1, ..., n$, then $v(p^n|p^n) \geq v(p^0|p^n)$.

Proposition 2.4. *Any PPE representation satisfies the limited cycle inequalities. Moreover, if v is jointly continuous, satisfies the limited-cycle inequalities, dislikes mixtures, and $v(\cdot|p)$ is EU for each $p \in \Delta$, then v defines an EU-PPE representation by (2.2).*

Proof. Take any $p^0, p^1, ..., p^n \in \Delta$, with $v(p^i|p^{i-1}) > v(p^{i-1}|p^{i-1})$. The i^{th} term in this sequence implies by the representation that $p^{i-1} \notin c(\{p^0, ..., p^n\})$; since $c(\{p^0, ..., p^n\}) \neq \emptyset$ by assumption it follows that $p^n = c(\{p^0, ..., p^n\})$. This implies, by the representation, that $v(p^n|p^n) \geq v(p^i|p^n)$ for all $i = 0, 1, ..., n-1$, which implies the desired result.

Conversely, for any v that satisfies the three given restrictions, the limited cycle inequalities imply that $PE(D)$ is non-empty for any $D \in \mathcal{D}$. Thus by Theorem 2.2, v defines a EU-PPE representation.

□

Munro and Sugden (2003) mention the limited cycle inequalities (their Axiom C7), and defend the limited cycle inequalities based on a money-pump argument. In contrast, the limited cycle inequalities emerge here as a consequence of the assumption that $c(D)$ is always non-empty combined with the reference-dependent preference representation. If one considers a class of choice problems in which the agent always makes a choice, the limited cycle inequalities are a basic consequence of this and the agent's endogenous determination of her reference lottery, regardless of the normative interpretation of the inequalities.

2.4 Special cases of PPE representations

2.4.1 Kőszegi-Rabin reference-dependent preferences

It may not be apparent at first glance whether Kőszegi-Rabin preferences in (2.3) satisfy the limited-cycle inequalities that a PPE representation must satisfy to generate a non-empty choice correspondence. Kőszegi and Rabin (2006) cite a result due to Kőszegi (2010, Theorem 1) that a personal equilibrium exists whenever D is convex, or equivalently, an agent is free to randomize among elements of any non-convex choice set. It is unclear whether or when this restriction is necessary to guarantee the existence of a non-empty choice correspondence.

Kőszegi and Rabin suggest restrictions on (2.3). In particular, applications of Kőszegi-Rabin have typically assumed *linear loss aversion*, which holds when there are η and λ such that:

$$\mu(x) = \begin{cases} \eta x & \text{if } x \geq 0 \\ \eta \lambda x & \text{if } x < 0 \end{cases} \tag{2.4}$$

where $\lambda > 1$ captures loss aversion and $\eta \geq 0$ determines the relative weight on gain/loss utility. Proposition 2.5 shows that under linear loss aversion, Kőszegi-Rabin preferences with the PPE solution concept are a special case of the more general continuous EU-PPE representation.

Proposition 2.5. *Kőszegi-Rabin preferences that satisfy linear loss aversion satisfy the limited cycle inequalities and dislike mixtures.*

Proposition 2.5 is alternative result to Kőszegi and Rabin's (2006) Proposition 1.3, and to my knowledge provides the first general proof that a personal equilibrium that does not involve randomization always exists in finite sets for this subclass of Kőszegi-Rabin preferences.

While commonly used versions of Kőszegi-Rabin preferences can provide the v in a PPE representation, there are (pathological?) cases of Kőszegi-Rabin preferences that cannot.

Proposition 2.6. *Not all Kőszegi-Rabin preferences consistent with (2.3) satisfy the limited cycle inequalities.*

2.4.2 Reference lottery bias and dynamically consistent non-expected utility

Expectations-based reference-dependence is the central motivation to considering the PPE representation. Now equipped with some understanding of the revealed preference implications of a PPE representation, we might take the preference relations $\succeq_L$

and $\{\succeq_p\}_{p\in\Delta}$ as primitives, where $\succeq_p$ is the preference relation corresponding to $v(\cdot|p)$, and $p \succeq_L q$ corresponds to the ranking $v(p|p) \geq v(q|q)$. With these primitives, we can study axioms that capture reference lottery bias. This is similar to the standard exercise in the axiomatic literature on reference-dependent behaviour (e.g. Tversky and Kahneman (1991; 1992); Masatlioglu and Ok (2005; 2012); Sagi (2006)). In that vein, consider the *Reference Lottery Bias* axiom below, which is closely related to the "Weak Axiom of Status Quo Bias" in Masatlioglu and Ok (2012).

Reference Lottery Bias. $p \succeq_L q \implies p \succeq_p q$

I offer three interpretations of Reference Lottery Bias. The first interprets $\succeq_L$ as representing the preferences that take into account that expecting to choose and then choosing lottery p leads to p being evaluated against itself as the reference lottery. Under this interpretation, if an agent would want to choose p over q, knowing that this choice would also determine the reference-lottery against which they would evaluate outcomes, then the agent would also choose p over q when p is the reference lottery. The second interpretation (along the lines of Masatlioglu and Ok (2012)) is that $\succeq_L$ captures reference-independent preferences; in this second interpretation, if p is preferred to q in a reference-independent comparison, then when p is the reference lottery, p is also preferred to q. According to either interpretation, Reference Lottery Bias imposes that $\succeq_p$ biases an agent towards p relative to $\succeq_L$. This seems like a natural generalization of the endowment effect for expectations-based reference-dependence.

A third interpretation emphasizes $\succeq_L$ as the ranking of lotteries induced by the agent's ex-ante ranking of choice sets when restricted to singleton choice sets. Under this interpretation, an agent who wants to choose a lottery from a choice set according to her ex-ante ranking would also want to choose it from that choice set if she then expected that lottery, and it subsequently acted as her reference point.

What implications does the Reference Lottery Bias axiom have? Kőszegi-Rabin preferences do not satisfy Reference Lottery Bias; recall the example in Section 2.2.2 in which $v(p|p) > v(r|r)$ but $v(r|p) > v(p|p)$. This suggests a conflict between the psychology of reference-dependent loss aversion captured by the Kőszegi-Rabin model and the notion of Reference Lottery Bias defined in the axiom. No experimental evidence to my knowledge sheds light on this matter.

Proposition 2.7. *A PPE representation satisfies Reference Lottery Bias if and only if* $c(D) = m(D, \succeq_L)$.

Proposition 2.7 implies (recalling Proposition 2.3) that under Reference Lottery Bias, reference-dependent behaviour in a PPE representation is tightly connected to non-expected utility behaviour in $\succeq_L$.

The non-expected utility literature has provided numerous models of decision-making under risk based on complete and transitive preferences that, motivated by the Allais paradox, satisfy a relaxed version of the Mixture Independence axiom (e.g. Quiggin (1982); Chew (1983); Dekel (1986); Gul (1991)). The model of expectations-based reference-dependence based on the Reference Lottery Bias axiom is based on a dynamically consistent implementation of non-expected utility preferences (as in Machina (1989)). I offer two examples of PPE representations that satisfy Reference-Lottery Bias and capture expectations-based reference-dependence.

Example (Disappointment Aversion). Suppose $\succeq_L$ satisfies Gul's (1991) disappointment aversion; that is (letting $u(x)$ denote $u(\langle x,\, 1\rangle)$), $u(p) = \frac{1}{1+\beta}\sum_i p_i\,(u(x_i) + \beta \min[u(x_i),\, u(p)])$ represents $\succeq_L$ for some $\beta \geq 0$. Then Reference Lottery Bias implies:

$$v^{DA}(p|r) = \frac{1}{1+\beta}\sum_i p_i\,(u(x_i) + \beta \min[u(x_i),\, u(r)]) \tag{2.5}$$

In cases of lotteries over multidimensional choice objects, it is not hard to see how to extend (2.5) via additive separability across dimensions. The resulting functional form captures loss aversion relative to past expectations (as in Kőszegi-Rabin) but does not generate IIA violations.

$\diamond$

Example (Mixture Symmetry). Suppose $\succeq_L$ satisfies Chew, Epstein, and Segal's (1991) mixture symmetric utility; that is, there is a symmetric function ϕ such that $u(p) = \sum_i\sum_j \phi(x_i,\, x_j)$ represents $\succeq_L$. Then Reference Lottery Bias implies:

$$v^{MS}(p|r) = \sum_i\sum_j p_i r_j \phi(x_i,\, x_j) \tag{2.6}$$

While the functional form for v^{MS} in 2.6 does capture the Kőszegi-Rabin functional form in (2.3), but the ϕ function corresponding to v^{KR} is generally not symmetric.

◇

2.5 Alternative models of expectations-based reference-dependence: analysis of PE and CPE representations

2.5.1 Characterization of PE

In addition to the PPE representation in (2.2) which is used in most applications of expectations-based reference-dependence, Kőszegi and Rabin (2006) also discuss the PE as a solution concept as in (2.1). The analysis below shows that the PE representation can be axiomatized similar to the PPE representation, by replacing Weak RARP with Sen's α, changing the continuity assumptions, and modifying IIA Independence.

Sen's α. $p \in D' \subset D$ and $p \in c(D)$ implies $p \in c(D')$

Sen's α requires that if an item p is choosable in a larger set D, then it is also deemed choosable in any subset D' of D where p is available. Sen's α is strictly weaker than IIA.[16]

The Upper Hemicontinuity axiom is the continuity property satisfied by continuous versions of the standard model, in which choice is determined by a continuous binary relation.

UHC. $c(D) = c^U(D)$

Proposition 2.8 (i) $\iff$ (ii) provides an axiomatic characterizating of PE decision-making that does not make use of the structure of environments with risk; (ii) $\iff$ (iii) is a continuous version of Gul and Pesendorfer's (2008) result (Proposition 2.1 in this paper).[17]

[16] Sen's α and Sen's β are jointly equivalent to IIA; see Sen (1971) and Arrow (1959).

[17] The result (i) $\iff$ (iii) is a continuous version of Theorem 9 in Sen (1971).

Proposition 2.8. *(i)-(iii) are equivalent: (i) c satisfies Expansion, Sen's a, and UHC, (ii) c has a continuous PE representation, (iii) c is induced by a continuous binary relation.*

IIA Independence 2 modifies the antecedent in the IIA Independence axiom to PE. Under PE, a lottery q is revealed to block p if there is a D such that $p \in c(D)$ but $p \notin c(D \cup q)$. IIA Independence 2 has a different antecedent from IIA Independence that reflects the differences in how constraining lottery pairs are revealed in the two models. IIA Independence 2 also embeds a continuity requirement.

IIA Independence 2. If $p \in c(D)$ and $\exists \alpha \in (0,1]$ such that $p \notin c(D \cup (1-\alpha)p + \alpha q))$, then $\exists \epsilon > 0$ such that $\forall \alpha' \in (0,1]$, $\forall \hat{p} \in N_p^\epsilon$, $\forall \hat{q} \in N_q^\epsilon$, and $\forall D' \ni (1-\alpha')\hat{p} + \alpha' \hat{q}$, $\hat{p} \notin c(D')$.

Theorem 2.3 provides a characterization of a continuous EU-PE representation.

Theorem 2.3. *c satisfies Expansion, Sen's a, UHC, IIA Independence 2, and Transitive Limit if and only if c has a continuous EU-PE representation. These axioms jointly imply that Induced Reference Lottery Bias holds as well.*

2.5.2 Characterization of CPE

Kőszegi and Rabin (2007) also introduce the choice-acclimating personal equilibrium (CPE) concept:

$$CPE_v(D) = \arg\max_{p \in D} v(p|p) \tag{2.7}$$

While most applications of expectations-based reference-dependence use the PPE solution concept, many use CPE. Theorem 2.4 clarifies the revealed preference foundations of CPE decision-making.

Theorem 2.4. *(i)-(iii) are equivalent. (i) c satisfies IIA and UHC, (ii) c has a continuous EU-CPE representation in which v is continuous, (iii) there is a complete, transitive, and continuous binary relation $\succeq$ such that $c(D) = m(D, \succeq)\ \forall D$.*

Theorem 2.4 appears to be a negative result - it suggests that expectations-based reference-dependence combined with CPE has no testable implications beyond the standard model of preference maximization! However, CPE decision-making can fail the Mixture Independence Axiom in ways that are consistent with expectations-based reference-dependent behaviour. This raises the question of what restrictions the Induced Reference Lottery Bias impose on the representation. Say that a binary relation $\succeq$ is *quasiconvex* if $p \succeq q \implies p \succeq (1-\alpha)p + \alpha q \, \forall \alpha \in (0,1)$.

Proposition 2.9. *Suppose* $\exists \succeq, v$ *such that* $c(D) = m(D, \succeq) = CPE_v(D)$*. (i)-(iii) are equivalent: (i)* c *satisfies Induced Reference Lottery Bias, (ii)* $\succeq$ *is quasiconvex, (iii)* $v(p|p) \geq v(q|q) \implies v(p|p) \geq v((1-\alpha)p + \alpha q | (1-\alpha)p + \alpha q) \, \forall \alpha \in (0,1)$*.*

Remark 2.4. Proposition 2.7 and Theorem 2.4 establish that if c has a PPE representation that satisfies the Reference Lottery Bias axiom, then $PPE_v(D) = CPE_v(D)$.

Example (Kőszegi-Rabin and Mixture Symmetry). Under CPE concept, the requirement that ϕ in 2.6 be symmetric is without loss of generality. Thus the Kőszegi-Rabin functional form in 2.3 corresponds to a special case of the mixture symmetric utility functional form in 2.6 under CPE.

$\diamond$

Chapter 3

Calibration without Reduction for Non-Expected Utility

Recent calibration critiques of Rabin (2000) and Safra and Segal (2008) show that whenever expected utility (EU) and non-expected utility (non-EU) define utility over final wealth states, they cannot simultaneously exhibit nonnegligible risk aversion over small stakes and can only exhibit moderate risk aversion over large stakes. Introspection and empirical evidence suggest that even if the stakes are small, most people would rather not take a small risk with a positive expected value if it could involve a loss of money. Yet most people still take substantial risks over large stakes, for instance, by investing in stocks. As a result, these calibration critiques have been widely understood as suggesting the demise of descriptive theories that define utility over final wealth except as a normative benchmark—further suggesting that descriptive models must define utility over gains and losses. But defining utility over final wealth gives non-EU theories a tractability and modelling discipline that more psychologically based theories such as prospect theory lack.

This paper shows that non-expected utility can generate both nonnegligible small-stakes risk aversion as well as the moderate large-stakes risk aversion. The crucial assumption made here is that a decision maker (DM) who faces preexisting risks treats a gamble that is offered as the first stage of a two-stage *compound lottery*, which is then not treated as equivalent to the one-stage lottery that gives the same probability distribution over final wealth but is evaluated recursively (Segal, 1990). This contrasts sharply with existing proposals for solving the Rabin critique, which have all relied on abandoning consequentialism—the assumption that utility is defined over final wealth levels—assuming instead that utility is also evaluated over gains and losses or lab income.

The intuition for the rank-dependent utility (RDU) (Quiggin, 1982) case of the main results of the paper is as follows: without background risk, RDU can produce descriptively reasonable risk aversion at a range of stakes through probability weighting even if utility is defined over wealth levels. Now suppose DM faces background risk $\tilde{w}$, and the utility-for-wealth function u is linear. Under recursive RDU, a compound lottery is evaluated by a folding back procedure, and DM evaluates the offered gamble $(-L, .5; +G, .5)$ by folding back the compound lottery $[\tilde{w} - L, .5; \tilde{w} + G, .5]$ to $[c(\tilde{w} - L), .5; c(\tilde{w} + G), .5]$ where c is DM's certainty-equivalent function. When u is linear, $c(\tilde{w} + x) = c(\tilde{w}) + x$; this compound lottery is evaluated as $[c(\tilde{w}) - L, .5; c(\tilde{w}) + G, .5]$, and probability weighting produces small-stakes risk aversion over the offered gamble the way it would without background risk. With a linear u, DM turns down $(-L, .5; +G, .5)$ if and only if DM turns down $(-tL, .5; +tG, .5)$ for all $t > 0$; therefore, small-stakes risk aversion due to probability weighting is compatible with reasonable large stakes risk aversion.

Relation to previous literature Since EU maximizers are approximately risk neutral over small stakes, they would only be willing to pay a trivial amount to avoid small risks. Popular alternatives to EU that define utility over final wealth levels are either (i) 'smooth' as in (Machina, 1982), locally risk neutral, and subject to the same criticism as EU (Safra and Segal, 2008), or (ii) obtain nonnegligible risk aversion over small stakes because they weigh probabilities nonlinearly (as suggested by the Allais paradox) and are hence immune from Rabin's critique.

However, most people face substantial lifetime wealth risk (e.g., employment-income risk and ownership of risky assets). The combination of a gamble offered in the lab in the presence of preexisting wealth risk is naturally viewed as a two-stage compound lottery in which the offered gamble resolves first. When DM reduces compound lotteries to single-stage lotteries by multiplying out probabilities, any small-stakes gamble offered adds only minimally to lifetime wealth risk; therefore, probability weighting is mostly determined by preexisting lifetime wealth risk and does not produce substantial risk aversion over offered small-stakes gambles (Safra and Segal, 2008; Barberis, Huang, and Thaler, 2006). This argument relies crucially on the assumption that DM satisfies reduction of compound lotteries—an assumption that is not consistent with experimental evidence. Instead, this paper assumes recursive preferences over compound lotteries.

Recursive non-EU (RNEU) preferences over compound lotteries are used in this paper as a descriptive model of decision making, following Segal (1990). The theoretical distinction between compound versus single-stage lotteries was first suggested by Samuelson (1952). RNEU preferences have been applied by Segal (1987b) to explain ambiguity aversion (see also Dillenberger and Segal (2012)) and by Dillenberger (2010) to explain preferences for one-shot resolution of uncertainty. Dillenberger also remarks that an RNEU DM behaves as if they bracket narrowly; section 1.4 of this paper makes a precise connection between RNEU and narrow bracketing.

The theoretical tradition of RNEU preferences following Segal (1990) is related to but distinct from the use of recursive utility due to preferences over the timing of resolution of uncertainty (Kreps and Porteus, 1978; Epstein and Zin, 1989). When recursive preferences are used only because of preferences over the resolution of uncertainty then DM applies reduction of compound lotteries to an offered delayed risk combined with income risk that resolves at the same time, and will not demonstrate small-stakes risk aversion over such gambles (Barberis, Huang, and Thaler, 2006).

Existing theoretical approaches that avoid a calibration critique (Kahneman and Tversky (1979); Cox and Sadiraj (2006); Barberis, Huang, and Thaler (2006)) directly incorporate narrow bracketing by assuming that the value function is defined over the outcomes of a gamble as well as (possibly) over final wealth states. RNEU is formally very different from these nonconsequentialist models in that in RNEU the utility function is only defined over final wealth states and not directly over the outcomes of a gamble. While RNEU does not assume narrow bracketing and is fully consequentialist, Theorems 3.1 and 3.2 precisely establish a connection between RNEU and models of narrow bracketing: an RNEU DM behaves as-if she engages in a form of narrow bracketing, at least over small-stakes or if her lottery preferences satisfy constant absolute risk aversion.

The RDU special case of the RNEU preferences studied in this paper are still subject to calibration arguments by Neilson (2001) and by Safra and Segal (2008, Theorem 1)[18]; however, the assumed risk-averse choice patterns behind these critiques

[18]See also a related critique of RDU by Sadiraj (2012)

have limited experimental evidence and lack field evidence. For a literature review of calibration critiques, see Section 8.6 of Wakker (2010).

Experimental evidence on compound lotteries Halevy (2007) finds that 80 percent of subjects violate reduction of compound lotteries, while 59 percent of subjects' choices are best explained by RNEU. Previous experimental work also found substantial violations of reduction of compound lotteries that suggest the use of RNEU preferences; for example, Carlin (1992); Camerer and Ho (1994). Recursive preferences over compound lotteries are also consistent with experimental findings that randomly picking one of the subject's many decisions to determine payment is an incentive compatible mechanism for eliciting preferences (Cubitt and Sugden (1998)).

3.1 Theory: RNEU risk preferences with background wealth risk

3.1.1 Non-expected utility over lotteries

Notation Let $W = \Re_+$ denote the set of feasible final wealth levels, and consider a preference over one-stage lotteries, $V : \Delta(W) \to \Re$ with utility-for-wealth function $u : W \to \Re$ and associated with certainty equivalent function c. A *one-stage lottery* over W can be written as $q = [w_1, q_1; ...; w_m, q_m] \in \Delta(W)$ whenever q has finite support, where q_i denotes the probability of receiving prize w_i. Assume V is increasing in the sense of first-order stochastic dominance. Adopt the convention that $w_1 \leq ... \leq w_m$. Say that V is *risk averse* if it is averse to mean-preserving spreads.

Popular models The two most commonly used non-EU theories are RDU (Quiggin (1982), Yaari (1987), Segal (1990)), and disappointment aversion (DA) (Gul, 1991). Table 3.1 reviews these and EU; it notes conditions under which RDU and DA demonstrate the Allais paradox and small-stakes risk aversion not present under EU. DA preferences are a special case of the larger class of betweenness-satisfying preferences (Dekel (1986), Chew (1989)).

Table 3.1: Non-expected utility theories

Theory	$V([w_1, q_1; ...; w_m, q_m])$	Allais?	Small-stakes risk averse?
EU	$\sum_{i=1}^{m} q_i u(w_i)$	No	Not if u' exists
RDU	$\sum_{i=1}^{m} [g(\sum_{j=1}^{i} q_j) - g(\sum_{j=1}^{i-1} q_j)] u(w_i)$	g concave	g concave
DA	$\sum_{i=1}^{m} \frac{1+\beta \mathbb{I}_{u(w_i)<V(q)}}{1+\beta \sum_{i=1}^{n} q_j \mathbb{I}_{u(w_j)<V(q)}} q_i u(w_i)$	$\beta > 0$	$\beta > 0$

Sources: Gul (1991); Segal and Spivak (1990); Segal (1987a)

3.1.2 Recursive non-expected utility

RNEU extends non-EU preferences over single-stage lotteries to the domain of compound lotteries.

Define a *compound lottery* as a finite lottery over lotteries over final wealth levels; a compound lottery can be written as $Q = [q^1, p_1; ...; q^n, p_n]$ where $q^i \in \Delta(W)$ and p_i is the probability of receiving lottery q^i; let $\Delta(\Delta(W))$ denote the set of compound lotteries. The utility function U is defined over compound lotteries over final wealth levels. Without loss of generality, adopt the convention that for a compound lottery Q as above, $V([q^1, 1]) \leq ... \leq V([q^n, 1])$.[19]

A RNEU maximizer evaluates a compound lottery Q via a simple two step procedure:

1. Compute the certainty equivalent of each lottery q^i that is a possible prize of Q:
$$c(q^i) = u^{-1} \circ V(q^i)$$

2. Recursively compute the value of the compound lottery as the non-expected

[19] Readers familiar with Segal (1990) will note that I assume time neutrality here. This is not essential for the conclusions.

utility of the one-step lottery $[c(q^1), p_1; ...; c(q^n), p_n]$:

$$U(Q) = c([c(q^1), p_1; ...; c(q^n), p_n]) \tag{3.1}$$

An alternative to the recursivity assumption is that a DM immediately reduces the compound lottery to a single-stage lottery, which is then evaluated according to V. Such a reduction of compound lotteries assumption is *not* consistent with evidence that subjects fail to reduce compound lotteries to one-stage lotteries presented earlier. If a non-EU DM reduces compound lotteries, compound lottery Q is evaluated as equivalent to the one-stage lottery $Q^R = [w_1, \sum_{i=1}^{n} p_i q_1^i; ...; w_K, \sum_{i=1}^{n} p_i q_K^i]$. For purposes of comparison, a *non-expected utility with reduction* DM evaluates a compound lottery p by:

$$U^{ROCL}(Q) = c(Q^R) \tag{3.2}$$

3.1.3 Wealth risk as a compound lottery

A one-time choice is never the only thing going on in a DM's life. Empirical work shows that people face substantial risks in their lives (Guiso, Jappelli, and Pistaferri, 2002). If we want to retain the modeling discipline of defining utility over final wealth levels, then we have to make a choice about how to model a DM's attitude to risk from a one-time gamble and from everything else in life. The combination of a one-time gamble offered (like those offered in lab experiments) and background wealth risk constitutes a compound lottery composed of two distinct and independent sources of risk in which the one-time gamble resolves first, and the rest of life's uncertainties resolve in due course.

Consider a DM who faces background wealth risk described by the random variable $\tilde{w} = [w_1, q_1; ...; w_m, q_m]$, which is not the subject of choice, and who is offered the gamble over prizes $\hat{p} = (y_1, p_1; ...; y_n, p_n)$ where $y_i \in Y \subset \Re$ is a monetary prize added to or taken away from the DM's final wealth after lottery $\hat{p}$ resolves. Let $\hat{p} \oplus \tilde{w}$ denote the compound lottery formed by simple gamble over prizes $\hat{p}$, which resolves first, and independent background risk $\tilde{w}$, which resolves second. The compound lottery $\hat{p} \oplus \tilde{w}$

is given by:

$$\hat{p} \oplus \tilde{w} = [\tilde{w} + y_1, p_1; ...; \tilde{w} + y_n, p_n] \tag{3.3}$$

where $\tilde{w} + y_i = [w_1 + y_i, q_1; ...; w_m + y_i, q_m]$ denotes the lottery over final wealth states that the DM faces if prize y_i is won in the gamble $\hat{p}$. The compound lottery $\hat{p} \oplus \tilde{w}$ is well defined whenever $w + y_i \in W$ for each w in the support of $\tilde{w}$ and each y_i in the support of $\hat{p}$.[20]

Say that a DM with utility function U defined on $\Delta(\Delta(W))$ treats an offered gamble $\hat{p}$ in the presence of background risk $\tilde{w}$ as a compound lottery in which $\hat{p}$ resolves first if for any offered gambles $\hat{p} \in \Delta(Y)$, DM evaluates the utility of $\hat{p}$ according to $U(\hat{p} \oplus \tilde{w})$.

3.1.4 Nonreduction and narrow bracketing

Segal (1990) first suggested replacing the reduction of compound lotteries axiom with recursivity as a consequentialist alternative to prospect theory that captures Kahneman and Tversky's (1979) "isolation effect" (a particular example of a failure of the reduction of compound lotteries axiom). Rabin (2000) noted that prospect theory is immune to his calibration critique. Existing theoretical approaches that avoid a calibration critique (Kahneman and Tversky (1979); Cox and Sadiraj (2006); Barberis, Huang, and Thaler (2006)) directly incorporate narrow bracketing by assuming that the value function is defined over the outcomes of a gamble as well as (possibly) over final wealth states. RNEU is formally very different from these nonconsequentialist models in that the value function is only defined over final wealth states and not directly over the outcomes of a gamble.

While RNEU does not assume that a gamble is framed narrowly, Theorem 3.1 demonstrates that when c satisfies constant absolute risk aversion, an RNEU DM behaves as if she brackets narrowly: that is, her choices among offered gambles are independent of the background risk she faces. Formally, c satisfies *constant absolute*

[20] An alternative but less intuitive assumption is that DM distinguishes between one-stage and compound risks but treats the gamble $\hat{p}$ as resolving in the second stage. The main results of the paper are not sensitive to this assumption.

risk aversion if for any $\tilde{w} \in \Delta(W)$ and $y \in \Re$ such that $\tilde{w} + y \in \Delta(W)$, $c(\tilde{w} + y) = c(\tilde{w}) + y$.

Theorem 3.1. *Suppose a recursive non-EU decision maker treats an offered gamble in the presence of background risk as a compound lottery as in (3.3), and has lottery preferences that satisfy constant absolute risk aversion. Then c has a unique extension to $\Delta(\Re)$, $\hat{c}$, such that, $U(\hat{p} \oplus \tilde{w}) = \hat{c}(\hat{p}) + c(\tilde{w})$ represents preferences whenenver $\hat{p} \oplus \tilde{w} \in \Delta(\Delta(W))$.*

Proof. Extend c to the set of the set $\Delta(\Re)$ of lotteries on $\Delta(\Re)$ with support bounded from below by the identification that for any $q \in \Delta(\Re)$ define $\underline{w} = -\inf\{\text{support} q\}$, and extend c to $\Delta(\Re)$ by $\hat{c}(q) = c(q + \underline{w}) - \underline{w}$. Since c satisfies constant absolute risk aversion, this extension is unique. Under RNEU,

$U(\hat{p} \oplus \tilde{w}) = c([c(\tilde{w} + y_1), p_1; ...; c(\tilde{w} + y_n), p_n])$

Applying constant absolute risk aversion at the second stage, and then at the first stage making use of $\hat{c}$,

$= c([c(\tilde{w}) + y_1, p_1; ...; c(\tilde{w}) + y_n, p_n])$

$= \hat{c}(\hat{p}) + c(\tilde{w})$ □

3.1.5 Non-reduction and small-stakes risk aversion

Theorem 3.1 suggests that demonstrating small-stakes risk aversion under RNEU does not lead to calibration implications for risk aversion over larger stakes gambles.

Say that c satisfies *constant relative risk aversion* if for any $\tilde{w} \in \Delta(W)$ and $t \in \Re_+$, $c([tw_1, q_1; ...; tw_m, q_m]) = tc(\tilde{w})$. Preferences that satisfy both constant absolute risk aversion and constant relative risk aversion have been termed *constant risk averse* by Safra and Segal (1998). Special cases of constant risk averse preferences include linear u cases of RDU and DA. Corollary 3.1 shows that for the class of non-EU preferences satisfying constant risk aversion, fairly tight calibration implications can be drawn from turning a gamble $\hat{p}$, but these implications seem reasonable.

Corollary 3.1. *Suppose c satisfies constant risk aversion. For $\hat{p}^t = (ty_1, p_1; ...; ty_n, p_n)$, $U(\hat{p}^t \oplus \tilde{w}) = t\hat{c}(\hat{p}) + c(\tilde{w})$ represents preferences whenenver $\hat{p}^t \oplus \tilde{w} \in \Delta(\Delta(W))$, where $\hat{c}$ is the extension of c from Theorem 3.1.*

An implication of the corollary is that whenever we know that lottery preferences satisfy constant risk aversion, then DM turns down $\hat{p}$ at $\tilde{w}$ at which $U(\hat{p} \oplus \tilde{w})$ is well defined if and only if she turns down $\hat{p}^t$ for all $\tilde{w}$ and for all $t > 0$.

Typical applications of non-EU preferences do not assume constant risk aversion, but rather allow diminishing marginal utility of wealth. However, one might expect that over small stakes, non-expected utility preferences behave like constant risk averse preferences since u is almost linear locally. I define a notion of dual differentiability, as a way of approximating preferences by a derivative taken with respect to a dual mixture of lotteries in the sense of Yaari (1987), but that is only defined around mixtures in which one of the lotteries is degenerate. In the Appendix (Proposition B.1), I show that RDU, semi-weighted utility preferences including DA (Chew, 1989), and Fréchet differentiable preferences are all weakly dually differentiable in the sense described below. Theorem 3.2 shows that for weakly dually differentiable preferences, behavior over small-stakes gambles given a fixed distribution of background wealth risk is well approximated by a constant risk averse certainty equivalent function.

Formally, say that c is *dually differentiable at wealth* w if for each $\hat{p} \in \Delta(Y)$ and $w \in W$ if there is a linear-in-money $\hat{c}(\cdot)$ such that $c(\hat{p}^t + w) = w + t\hat{c}(\hat{p}) + o(t)$. Say that c is *weakly dually differentiable given wealth risk* $\tilde{w}$ if there are $\bar{c}_{y^+}(\tilde{w}), \bar{c}_{y^-}(\tilde{w})$ such that $c(\tilde{w} + ty) = c(\tilde{w}) + ty\bar{c}_{y^{\text{sign}(y)}}(\tilde{w}) + o(t)$. Additionally, say that a certainty equivalent function c is *first-order risk averse* at wealth w if for any $\hat{p}^t$ with an expected value of zero, $\frac{dc(w+\hat{p}^t)}{dt}|_{t=0^+} < 0$ (Segal and Spivak, 1990).

Theorem 3.2. *If c is weakly dually differentiable at $\tilde{w}$ with $\bar{c}_{y^+}(\tilde{w}) \leq \bar{c}_{y^-}(\tilde{w})$, and is also dually differentiable at $w = c(\tilde{w})$, then for any $a \in \Re$, $U(\hat{p}^t \oplus \tilde{w}) = t\hat{c}(\hat{p} + a) - ta + c(\tilde{w}) + o(t)$. Moreover, if c is first-order risk averse then so is $\hat{c}$.*

3.1.6 The Rabin critique?

An immediate corollary of Theorem 3.1 in the spirit of Rabin's Theorem is that if c is constant absolute risk averse, $U(\hat{p} \oplus \tilde{w})$, DM turns down gamble $\hat{p}$ when her background wealth distribution is $\tilde{w}$ if and only if she turns down $\hat{p}$ for any $\tilde{w}$. This gives no sense of DM's attitude towards larger stakes gambles, leading to the following remark:

Remark 3.1. Suppose a recursive non-EU decision maker treats an offered gamble in the presence of background risk as a compound lottery as in (3.3), and always turns down an actuarially favorable gamble $\hat{p}$ for any distribution of background risk $\tilde{w}$. Knowing only that V is globally risk averse, the strongest conclusion that can be drawn is that DM will always turn down any mean-preserving spread of $\hat{p}$.

Remark 3.1 shows that weak assumptions lead to weak conclusions, but what if we made stronger assumptions about c? Theorem 3.2 and Corollary 3.1 give strong intuition for what attitudes towards a small-stakes gamble $\hat{p}$ will imply about attitudes towards gambles at larger stakes for preferences for classes of preferences that include a constant risk averse subclass. Theorem 3.3 offers an counterpart to Rabin's (2000) calibration theorem, assuming instead that DM has risk averse RDU/DA RNEU preferences and faces background risk.

Theorem 3.3. *Suppose a recursive non-EU decision maker treats an offered gamble in the presence of background risk as a compound lottery as in (3.3), and always turns down a gamble $\hat{p} = (y_1, p_1; ...; y_n, p_n)$ for any distribution of background risk $\tilde{w}$. Knowing only that c is risk averse and is RDU, then the strongest restriction on large-stakes gambles that can be drawn without further assumptions is that DM will turn down any gamble $\hat{p}^t = (ty_1, p_1; ...; ty_n, p_n)$ for all $t > 1$ and for all $\tilde{w}$. The same result applies if "RDU" is replaced with "DA".*

Proof. DM turns down $\hat{p}^t$ if $U(\hat{p}^t \oplus \tilde{w}) - c(\tilde{w}) < 0$. Turning down $\hat{p}$ implies $U(\hat{p} \oplus \tilde{w}) - c(\tilde{w}) < 0$. If u is linear, then under DA and RDU this implies that DM turns down $\hat{p}^t$ for all t but cannot rule out accepting more favorable gambles. If u is concave, then it can be shown that under risk aversion and either DA or RDU that $(u \circ U)(\hat{p}^t \oplus \tilde{w})$ is concave in t; therefore, DM will still turn down $\hat{p}^t$ for any $t > 1$. However, EU remains a special case of DA and RDU, and if DM had EU preferences, $\hat{p}^t$ would be accepted for a sufficiently small t whenever $0 \notin \text{support}(\tilde{w})$.

Under DA and RDU, if V is globally risk averse u must be weakly concave, so the strongest calibration result possible comes from the case where u is linear for $t > 1$. That is, DM will turn down $\hat{p}^t$ for all $t > 1$. □

Since recursive DA is a special case of the more general class of betweenness preferences (Dekel, 1986; Chew, 1989), if recursive DA is immune to calibration critiques, then the class of recursive betweenness preferences is immune to Rabin-style calibration critiques, as are more general classes of preferences. Thus, recursive versions of a wide range of non-EU theories are not susceptible to calibration critiques, and are potentially suitable for modeling risk preferences over both small and large stakes.

3.2 Calibration

What constitutes descriptively reasonable risk aversion is a quantitative question. This section calibrates a version of recursive RDU and shows that this calibration can produce descriptively reasonable risk aversion, while EU and RDU with reduction cannot.

Convenient functional forms for g and u should have as few parameters as possible to calibrate and should be easily comparable to commonly used models. I adopt the standard power utility-for-wealth function:

$u(w) = \frac{w^{1-\gamma}}{1-\gamma}$

and the probability weighting function:

$g(p) = p^{\nu}$

axiomatized by Grant and Kajii (1998) and used in Safra and Segal (2008) since it is only one parameter richer than EU, is consistent with small-stakes risk aversion and Allais-type choices when $0 < \nu < 1$, and captures expected utility as a special case when $\nu = 1$.

Chetty (2006) points out that the curvature of the utility-for-wealth function also governs how an individual makes trade-offs between labour and leisure. I use $\gamma = .71$, suggested by Chetty based on previous studies of labor supply responses to wage changes.[21] I then calibrate ν to match modal choices in Holt and Laury's (2002) experimental data on small-stakes risk aversion to the extent possible. While EU cannot avoid mispredicting the modal choice in their data when most subjects demonstrate risk aversion, if $\nu \in [.5, .64]$, the calibrated recursive RDU model fits the data reasonably but not perfectly.

[21] While Chetty assumes expected utility in his calculations, the approach he takes fully carries through to RDU in the case where utility is separable in consumption and leisure; I use Chetty's estimates from this case.

Table 3.2: Calibration results - small and large stakes risk aversion

Loss	$\nu = .5$	$\nu = .64$	$\nu = .5$, reduction	$\nu = 1$ (EU)
10	24.14	17.91	10.10	10.00
100	241.60	179.21	103.36	100.03
200	483.59	358.62	209.61	200.13
500	1211.84	898.12	538.77	500.83
1000	2433	1801	1112	1003
2000	4904	3623	2328	2013
5000	12555	9219	6403	5084
10000	26133	18992	14364	10343
25000	73931	52052	46738	27270
50000	185392	123239	137678	60105

Gain required for DM to take (-Loss, .5; Gain, .5)

While the risk in $\tilde{w}$ only has a second-order effect on decisions among offered gambles in recursive RDU, the risk in $\tilde{w}$ reduces risk aversion RDU with reduction. To allow for comparison, take $\tilde{w} = U[\$100000, \$500000]$ to capture background wealth risk.[22]

Table 3.2 summarizes how different calibrated models discussed above would predict that a DM would make choices in $(-L, .5;\ G, .5)$ gambles. In each row of the table, L is fixed at the level in the left-hand column, while the entry in the table lists the G at which a DM would be indifferent to either taking or turning down the listed gamble.

Table 3.2 (Columns 1 and 2) indicates that for $\nu = .5$, $.64$ recursive RDU can produce descriptively reasonable risk aversion over both small and large stakes. RDU with reduction produces barely any risk aversion over small stakes (Column 3). Even for stakes into the thousands of dollars, EU induces preferences over gambles that are extremely close to expected value maximization (Column 4). Even with a higher value for γ, EU would induce preferences over gambles that are extremely close to expected value maximization over stakes of hundreds of dollars. These quantitative

[22] I derive quantitative results using a discrete approximation of the uniform distribution. I assume that lifetime wealth has an expected present value of $300,000, since this figure is emphasized in Rabin (2000), but the quantitative results are not particularly sensitive to this assumption.

results are not sensitive to the choice of a distribution for background wealth risk.

3.3 Conclusion

This paper has shown that RNEU can produce non-negligible small-stakes risk aversion without implying ridiculous large-stakes risk aversion, and generate a form of 'as-if' narrow bracketing over small-stakes gambles. A calibration exercise demonstrated that recursive RDU can be calibrated to provide descriptively reasonable levels of risk aversion in the small and in the large. The non-expected utility theories studied in this paper are attractive. Non-expected utility theories, rank-dependent utility and disappointment aversion in particular, have clear axiomatic foundations, have been well studied, and have proven tractable in applications. The RNEU approach to applying non-expected utility preserves this tractability for a decision-maker who faces multiple risks. Furthermore, the two departures this model does make from expected utility theory are each well supported by experimental evidence on the Allais paradox and nonreduction of compound lotteries.

Chapter 4

List Elicitation of Risk Preferences

An ideal experimental study of individual decision-making would elicit precise data while providing subjects with monetary incentives that allow for unambiguous inferences about individual preferences. List elicitation with the random incentive scheme (RIS) is a standard technique to achieve this goal and has become the workhorse method for experimental economists. In list elicitation, subjects respond to a series of binary choices lined up as a list. List elicitation is typically combined with the RIS, in which one of these choices is randomly selected to determine the subject's payment. The unambiguous interpretation of choice behavior in terms of individual preferences relies on the assumption that the list elicitation combined with the RIS does not impact subjects' choices.

A common hypothesis is that subjects behave as-if they make each choice in *isolation* from all other choices and independent of the details of the incentive scheme they face. Experimental data obtained from list elicitation with the RIS is frequently interpreted under the assumption that the isolation hypothesis holds. But when is the isolation hypothesis appropriate?

The theoretical literature provides conflicting guidance as to whether we should expect list elicitation combined with the RIS to impact subjects' choices. Formally, isolation holds if and only if the subjects' preferences over compound lotteries induced by her choices and the external randomization device satisfy Segal's (1990) compound independence axiom.[23] In particular, isolation does not hold if subjects have non-expected utility preferences and obey the reduction of compound lotteries axiom (Karni and Safra, 1987). An earlier literature showed that the Becker, DeGroot, and Marschak (1964) (BDM) elicitation mechanism can lead to preference reversals Grether and Plott (1979) in the direction predicted by Karni and Safra (1987). This

[23] Cox and Epstein (1990) elaborate on this point.

Table 4.1: Binary choice versus list elicitation

	Binary choice	List	p-value
$(\$y, 1)$	23%	50%	$<.001$
$(\$y, p)$	30%	45%	.02

Percent choosing the riskier option.
p-values are for a Fisher's exact test of association.

suggests that the possibility of a failure of isolation due to a failure of the independence axiom is more than a theoretical curiousity.

The experimental literature provides only partial guidance on whether the theoretical possibility suggested by Karni and Safra is empirically relevant in the context of list elicitation. Experimentalists often invoke the evidence for isolation from binary choice experiments (Starmer and Sugden, 1991; Cubitt and Sugden, 1998) as support for isolation in the context of list elicitation. However, no previous study examined whether this interpretation is appropriate. If list presentation invokes reduction and subjects have non-expected utility preferences, then we should not expect subjects to isolate each choice in a list of binary choices.

Whether isolation holds under list elicitation is ultimately an empirical question. We answer this question by using a between-subjects design to compare behaviour under list elicitation to behaviour in binary choice tasks. In one group of treatments, subjects respond to a list (or two) of binary choices. In a second group of treatments, subjects make a single (or two) binary choice(s); the binary choice questions correspond exactly to the questions faced in line 11 of the lists.

Table 4.1 demonstrates our main finding. Since the binary choice tasks correspond to unique lines in each list, we can compare the frequency of risky choices (with the higher payment of $\$x$) under the alternative treatments. We find that the list increases the probability of choosing the risky lottery from 23% to 50% when the safer alternative is certain, and from 30% to 45% when the safer alternative is risky. This suggests that subjects do not isolate their choices under list elicitation.

An additional contribution of this paper is to evaluate the behaviour of a relatively new subject pool, that of online workers on Amazon's Mechanical Turk, whose behavior has barely been studied using real monetary incentives in standard decision-making tasks often studied by experimental economists. Our paper provides a new interface for managing subjects that allows us to use Mechanical Turk as a virtual recruitment tool for running fully incentivized individual decision-making experiments. We also replicate our findings using a standard subject pool of university students.

4.1 Experiment

Subjects were recruited from Amazon's Mechanical Turk online labour market. In Appendix A (to be written), we outline our procedures and relevant aspects of Mechanical Turk in detail. To recruit subjects, we released an ad for a task ("HIT" for Human Intelligence Task) that could be seen by online workers on Amazon's Mechanical Turk site ("turkers"). Any turker with a US based account and who had at least 95% of their past HITs approved was eligible to view a brief description of our study that mentioned the fixed payment of one dollar (which worked like the show-up fee in lab experiments) and the possibility of a bonus (which corresponded to the incentive scheme). From the description, interested turkers could accept a HIT and proceed to our website with a unique identifying code from the HIT. Each HIT would disappear once accepted by a turker.

On our website, turkers input their Mechanical Turk ID as well as the HIT ID, then completed a standard consent form, followed by a description of how the experiment would work and how payment would be determined, which was accompanied by a multiple choice comprehension quiz which subjects had to answer correctly before proceeding.

The main experiment consisted of sixteen different treatments, based on varying the payment mechanism and the order of the two questions (Q1 and Q2). Table 4.2 shows the questions and subquestions when using list elicitation, and Table 4.3 shows the sixteen different treatments. In treatments B1 and B2, subjects only answered one question (a binary choice) and were paid based on that question while in B12 and B21 subjects answered two binary choice questions, one of which was randomly selected to determine payment; the questions corresponded to line 11 of the lists in Table 4.2. In treatments beginning with L, each question consisted of a list of subquestions, each subquestion involved a binary choice, where the probability of winning the prize in the right hand side option decreased each line as subjects proceeded down the list. In treatments L1 and L2, subjects answered only one list question while in treatments LO12, LO21, LA12, and LA21, subjects answered two list questions. In treatments LO12 and LO21, the instructions made it clear that one of the two list questions would be randomly selected to determine payment. In treatments LA12 and LA21, the instructions made it clear that both of the two list questions would determine payment. Subjects were informed that whenever a list was used to determine payment (the sole list in L1 and L2, both lists in LA12 and LA21, and a randomly chosen list in LO12 and LO21), one line from the list would be randomly selected to be played out to determine the subject's bonus payment. Treatments LO21 and LA21 reversed the order of the list questions of treatments LO12 and LA12 respectively. In the L treatments, we allowed subjects to switch from Option A to Option B at any number of points on the list, but used a javascript pop-up to warn subjects who switched from Option B to Option A and then from Option A to Option B. The S (separate screens) treatments mirror the L treatments, except that before completing each list subjects responded to a sequence of (non-incentivized) binary choices that appeared on separate screens. In the S treatments, the binary choices tasks are chosen so as to spiral towards finding the switching point for a monotone subject. Subjects then responded to an incentivized list that was already filled in using their responses to the binary choice tasks but was otherwise identical to that in the corresponding L treatment (crucially, subjects were free to change their answers in the list).

Subjects completed the HIT by submitting a completion code generated by our website to the Mechanical Turk interface. A random number generator was used to determine the outcomes of all risk automatically, and subjects were informed of how much of a bonus would be paid after completing the study. Subject payments were credited to subjects' Mechanical Turk accounts within 30 minutes of completing the HIT. Our \$1 base payment is somewhat high compared to other experiments using Mechanical Turk given that our experiment should take at most 15 minutes (e.g. Horton, Rand, and Zeckhauser (2011)). Bonus payments of \$3 or \$4 provided relatively high stakes for this subject pool.

Table 4.2: Questions

	Q1		Q2	
Line	Option A	Option B	Option A	Option B
1	$(3, 1)$	$(4, 1)$	$(3, .5)$	$(4, .50)$
2	$(3, 1)$	$(4, .98)$	$(3, .5)$	$(4, .49)$
3	$(3, 1)$	$(4, .96)$	$(3, .5)$	$(4, .48)$
4	$(3, 1)$	$(4, .94)$	$(3, .5)$	$(4, .47)$
5	$(3, 1)$	$(4, .92)$	$(3, .5)$	$(4, .46)$
6	$(3, 1)$	$(4, .90)$	$(3, .5)$	$(4, .45)$
7	$(3, 1)$	$(4, 88)$	$(3, .5)$	$(4, .44)$
8	$(3, 1)$	$(4, .86)$	$(3, .5)$	$(4, .43)$
9	$(3, 1)$	$(4, .84)$	$(3, .5)$	$(4, .42)$
10	$(3, 1)$	$(4, .82)$	$(3, .5)$	$(4, .41)$
11	$(3, 1)$	$(4, .80)$	$(3, .5)$	$(4, .40)$
12	$(3, 1)$	$(4, .78)$	$(3, .5)$	$(4, .39)$
13	$(3, 1)$	$(4, .76)$	$(3, .5)$	$(4, .38)$
14	$(3, 1)$	$(4, .74)$	$(3, .5)$	$(4, .37)$
15	$(3, 1)$	$(4, .72)$	$(3, .5)$	$(4, .36)$
16	$(3, 1)$	$(4, .70)$	$(3, .5)$	$(4, .35)$
17	$(3, 1)$	$(4, .68)$	$(3, .5)$	$(4, .34)$
18	$(3, 1)$	$(4, .66)$	$(3, .5)$	$(4, .33)$
19	$(3, 1)$	$(4, .64)$	$(3, .5)$	$(4, .32)$
20	$(3, 1)$	$(4, .62)$	$(3, .5)$	$(4, .31)$
21	$(3, 1)$	$(4, .60)$	$(3, .5)$	$(4, .30)$
22	$(3, 1)$	$(4, .58)$	$(3, .5)$	$(4, .29)$
23	$(3, 1)$	$(4, .56)$	$(3, .5)$	$(4, .28)$
24	$(3, 1)$	$(4, .54)$	$(3, .5)$	$(4, .27)$
25	$(3, 1)$	$(4, .52)$	$(3, .5)$	$(4, .26)$
26	$(3, 1)$	$(4, .50)$	$(3, .5)$	$(4, .25)$

Table 4.3: Treatments

Order	Binary choice	One list	Pay one list	Pay both lists
Q1 only	B1	L1, S1		
Q2 only	B2	L2, S2		
Q1 then Q2	B12		LO12, SO12	LA12, SA12
Q2 then Q1	B21		LO21, SO21	LA21, SA21

S, L, and B respectively denote separate screen, standard list, and binary choice treatments

Table 4.4: Subjects by treatment

	B1	B2	B12	B21	L1	L2	LO12	LO21
n	39	41	20	22	47	45	36	35
n monotone					43	41	27	29
n regular					32	27	18	21

	LA12	LA21	S1	S2	SO12	SO21	SA12	SA21
n	36	33	48	49	37	32	26	25
n monotone	31	26	46	48	36	28	25	25
n regular	18	17	34	47	27	24	17	20

Table 4.5: Answers to line 11

	Binary choice		List elicitation	
	One choice	Two choices	One list	Two lists
Q1	23%	24%	47%	51%
Q2	27%	33%	46%	44%
n	39/41	42	66/74	162
Treatments	B1,B2	B12, B21	L1,L2,S1,S2	All O and A

Fraction choosing the riskier option

4.2 Results

All analysis of results focuses solely on monotone subjects: those who exhibit single-switching in each list and who do not choose a dominated option in the first line of a list. Except where specified otherwise, our results focus on *regular* subjects: subjects who never exhibit extreme risk seeking by sticking with B throughout the list, nor exhibit extreme risk aversion by switching immediately to A on the second line.

4.2.1 List elicitation versus binary choice

Table 4.5 shows the distribution of answers in binary choice for line 11 of the list, grouped by the incentives provided.

There are no significant differences between asking one question or asking two questions shown on separate screens when binary choice is used ($p = .92$ for Q1 and $p = .53$ for Q2), a finding consistent with the literature supporting the incentive compatibility of the RIS.

Table 4.6: Fanning in vs out

	SCR	EU	RCR
L,S	34%	20%	46%
B	19%	71%	10%

Fisher's test: $p < .001$

The most obvious difference in Table 4.5 is that in Q1 23% of subjects choose the risky option in binary choice, but 50% choose the risky option in list elicitation, a significant difference ($p < .001$, exact test). In Q2, 30% of subjects choose the risky option in binary choice, but 45% choose the risky option in list elicitation, a significant difference ($p = .02$). Comparable results would hold up if we included all monotone subjects or only focused on subsets of list treatments.

4.2.2 The independence axiom

Under both list elicitation and binary choice, responses are close to expected utility, though exhibiting a slight reverse common ratio effect when using list elicitation and a slight common ratio effect with binary choice elicitation. In binary choice, the violation of expected utility is not significant ($p = .80$ for an exact test for B1 vs B2, .47 for an exact aggregate test for B12 and B21). However, since these two questions only look in a very particular region of the Marschak-Machina triangle, we do not view this as providing strong evidence in favour of the independence axiom.

Pooling all the list treatments, the median choice pattern is $(\$4, .82) \succ (\$3, 1) \succ (\$4, .80)$ and $(\$4, .41) \succ (\$3, .5) \succ (\$4, .40)$, consistent with EU. A rank-sum test for equality of distributions of Q1 and Q2 has $p = .06$, suggesting an aggregate level deviation from EU that is borderline statistically significant (and is in the opposite direction of the standard common-ratio effect). A within-subject test suggests that violations of the independence axiom are significant ($p = .01$, signed-rank test).

Aggregate analysis of behaviour masks substantial heterogeneity of individual decisions that list elicitation picks up: under list elicitation, we detect violations of the independence axiom for 79% of subjects split between standard common ratio and reverse common ratio violations with the latter type of violation being slightly more frequent (Table 4.6). Binary choice data only detects violations of the independence

Table 4.7: Tests for payment mechanism and order effects

	Q1	Q2
Order effect (12=21)	.97	.32
Payment mechanism effect (O=A)	.06	.58
Separate screens effect (L=S)	.47	.38

p-values reported for a rank-sum test of equality of distribution

axiom (AB and BA choice patterns) for 29% of subjects, and using only data from line 11 would detect a similar fraction of EU violations for the L and S treatments.

4.2.3 Treatment effects

By using list elicitation, the possibility of within-list contamination is equally present in all L and S treatments. Our treatments allow us to test whether any cross-list contamination occurs. The 2x2x2 design embedded in the treatments ($\{$L,S$\}\times\{$O,A$\}\times\{$12,21$\}$) allows us to separately test for the presence of separate screen effects, payment mechanism effects, and order effects in each question.

Applying non-parametric analysis, we do find a mildly significant payment mechanism effect in Q1 (Table 4.7), however this effect would not be significant if we corrected for multiple hypothesis tests using a Bonferroni correction.[24]

So far, our results have focused on regular subjects. However, the most striking treatment effect is that the proportion of regular subjects is much higher in the S treatments than in the L treatments (78% vs. 57%, $p < .001$, see Table 4.4). As we might expect, there are relatively more (79% vs. 62%, $p = .01$) regular subjects in the treatments in which subjects faced only one list (as opposed to two). Neither order nor payment mechanism significantly affect the proportion of regular subjects ($p = .31, .52$ respectively, exact tests). We suspected that a combination of isolation in binary choice, a lack of influence of hypothetical versus real incentives, and a status-quo bias when the actual list was displayed filled in, might bridge the difference between standard list elicitation and binary choice. Table 4.7 shows that the incentivized choice data do not tend to support this view.

[24] The payment mechanism effect is driven entirely by the O12 treatments. With enough specification searching this may appear significant in some tests, but would not be after correcting for multiple hypothesis tests.

Table 4.8: mTurk vs. student subjects

	Binary choice		List elicitation	
	One choice	Students	mTurk	Students
Q1	23%	33%	50%	52%
n	81	27	228	21
Treatments	All B		All L,S,O,A	

Fraction choosing the riskier option

4.2.4 Comparison to a student subject pool

A limited pair of treatments confirms that student subjects exhibit behaviour that is similar to turkers. Students were recruited through the UBC Economics Lab subject pool using ORSEE (Greiner, 2004), and offered the opportunity to do an experiment online for payment by Interac money transfer. Since student subjects tend to be paid much higher amounts per hour than turkers typically earn, we changed the higher payoffs to \$13, \$10, and \$0 from \$4, \$3, and \$0.[25] Table 4.8 shows the results of the two student treatments, one corresponds to Q1 under binary choice (B1), and one corresponds to Q1 under list elicitation (L1).

The student data demonstrates exactly the same pattern as the data from turkers - students are more likely to choose the safer lottery in a binary choice task than when the choice is embedded in a list. Students' responses to Q1 in binary choice were not significantly different from turkers' answers to Q1 under binary choice ($p = .32$, exact test), nor were students' responses to Q1 in the list significantly different from turkers' responses to Q1 under list elicitation ($p = .82$, exact test).

4.3 Theory: binary choice versus list elicitation

How is it that subjects' preferences demonstrate more risk aversion when we ask subjects one question at a time? How is it that subject preferences appear close to consistent with EU when we use list elicitation, in spite of its precise data being ideally-suited to test the Independence Axiom?

[25] All major Canadian banks have a \$10 minimum transfer, there was no show-up fee, which was made clear to subjects in advance.

A subject who chooses option B for the last time at line i of the list version of Q1 receives the two-stage compound lottery:

$$\left[(\$4,\ 1)\,,\ \frac{1}{26};\ ...;\ (\$4,\ 1.02 - .02i)\,,\ \frac{1}{26};\ (\$3,\ 1)\,,\ \frac{26-i}{26}\right] \tag{4.1}$$

One possible explanation for our results is that there is some cross-question contamination within the elicitation list, which biased the answers of non-EU subjects under list elicitation. A subject satisfies Segal's (1990) Compound Independence axiom and evaluates the compound lottery in (4.1) by folding back, if and only if the binary choice at each 'branch' of the compound lottery is unaffected by the other branches (see Segal (1988); Starmer and Sugden (1991)). Karni and Safra (1987) show that if instead subjects reduces the compound lottery formed by her choices in the list and has non-EU preferences, then list elicitation will distort her choices. Segal (1988) shows that if subjects takes an alternative view of the compound lottery formed by her choices in the list, then list elicitation will distort her choices even if she satisfies Compound Independence. We sketch these two approaches below in the context of our experiment.

4.3.1 Karni and Safra (1987): Reduction of Compound Lotteries

If the subject satisfies the Reduction of Compound Lotteries axiom, then she evaluates (4.1) as equivalent to the single-stage lottery:

$$\left[\$4,\ \frac{1.01i - .01i^2}{26};\ \$3,\ \frac{26-i}{26}\right] \tag{4.2}$$

While a binary choice between (\$3, 1) and (\$4, .8) involves a choice between a certain and a risky option, the choice of a switching line in a list involves choice among multiple risky alternatives and a dominated certain one (for $i = 0$) when the subject views the list according to (4.2). The logic of the certainty effect suggests that a subject may rank (\$3, 1) $\succ$ (\$4, .8) and choose accordingly in a binary choice task, yet choose (\$4, .8) over (\$3, 1) on line 11 of a list since this conditional choice involves a choice between a riskier and a less risky (but not certain) alternatives. This exact logic could be extended, but does not directly apply to the reversal observed in Q2, although the reversal in Q2 is empirically weaker.

Some non-EU theories predict, or can accommodate this type of behaviour. For example, the Negative Certainty Independence (NCI) axiom of Cerreia-Vioglio, Dillenberger, and Ortoleva (2013) allows exactly the type of reversal observed in Q1, while disallowing the opposite pattern of behaviour. NCI does not make a direct prediction that relates to Q2. Some (but not all) functional forms for rank-dependent utility can also capture the observed reversals. For example, the power weighting function:

$f(p) = p^{\beta}$, $\beta > 1$

and the neo-additive weighting function (Chateauneuf, Eichberger, and Grant, 2007; Webb and Zank, 2011):

$$f(p) = \begin{cases} 1 & \text{if } p = 1 \\ ap + b & \text{if } p \in (0,1) \, f(p) = ap + b \\ 0 & \text{if } p = 0 \end{cases}$$

are both able to accommodate observed reversal behaviour for reasonable parameter values.[26] In the Appendix, we sketch how these weighting functions can be used to rationalize the data.

4.3.2 Segal (1988): non-standard view of the compound lottery formed by the list

Segal (1988) assumes that a subject satisfies Compound Independence, but views a BDM certainty equivalent elicitation scheme as the two-stage lottery with two prizes after the first stage (i) lottery at all "lines" where it is chosen, (ii) the uniform lottery over all $ amounts for "lines" where the dollar amount is chosen.

The version of list elicitation used in this paper corresponds to BDM with varying probabilities rather than the traditional BDM. To adapt Segal's (1988) approach to this setting, suppose a subject views the list elicitation with the RIS for Q1 according

[26] Some other weighting functions are less successful at accommodating the observed behaviour. For example, the Prelec (1998) weighting function does not predict the reversal observed in Q1 for standard parameter values.

to the three-stage compound lottery:

$$\left(\left[(\$4,\ 1)\,,\ \frac{1}{i};\ ...;\ (\$4,\ 1.02-.02i)\,,\ \frac{1}{i}\right],\ \frac{i}{26};\ [\$3,\ 1]\,,\ \frac{26-i}{26}\right) \tag{4.3}$$

A subject who evaluates her choices in the list by applying the Compound Independence axiom to (4.3) will tend to violate isolation unless her preferences are EU. As in the Karni and Safra (1987) explanation for the preference reversal, some (but not all) versions of rank-dependent utility, including some parameter values for the power weighting function, are able to account for the preference reversal observed in our data. The Appendix shows how RDU with the power weighting function can be used to rationalize our main results.

4.4 Discussion

4.4.1 Related literature on elicitation mechanisms and incentive-compatibility

Incentivized list elicitation was pioneered by Becker, DeGroot, and Marschak (1964), who introduced it as a way of eliciting certainty equivalents. Unfortunately, use of BDM induced preference reversals in which subjects' reported certainty equivalents of two lotteries were inconsistent with their choices in a binary choice task (Grether and Plott, 1979). Holt (1986), Karni and Safra (1987), and Segal (1988) showed that unless subjects evaluate their choices by applying the compound independence axiom to evaluate the compound lottery generated by the external randomizing device and their choices, BDM will not elicit subjects' true certainty equivalents if subjects violate the independence axiom. Safra, Segal, and Spivak (1990) show that assumptions that generate the common-ratio effect will tend to generate the observed preference reversals.

Starmer and Sugden (1991) and Cubitt and Sugden (1998) are widely interpreted as providing evidence in favour of the use of the RIS in experiments. They have some subjects answer multiple questions, with one randomly selected for payment, and compare their responses to those of subjects who answer only one incentivized question. They find only weak evidence for any "contamination effect" of RIS. In recent work, Cox, Sadiraj, and Schmidt (2011) and Harrison and Swarthout (2012) use comparable designs but consider a larger number of ways of paying subjects who answer multiple questions. None of these four papers, however, uses list elicitation. Harrison and Swarthout (2012) suggest that asking 30 questions and paying one does lead to different structural utility function estimates as compared to asking and paying one question.

While recent variants on list elicitation have been seen as avoiding some of the early problems with BDM, we show that the analysis of Karni and Safra (1987) can be extended to another variant on list elicitation, with an approach that would also apply to other variations on list elicitation. Recent experiments using list elicitation widely cite Starmer and Sugden (1991) and Cubitt and Sugden (1998) in support of combining list elicitation with the RIS; our results suggests that this inference is not warranted. List elicitation experiments typically have 10 to 100 questions per list on a single screen or sheet of paper at a time, and then subjects respond to multiple lists; this contrasts with experiments by Starmer and Sugden (1991) and Cubitt and Sugden (1998) in which each subject completes at most a small number of binary choice tasks, which are not given together as a list.

4.4.2 Discussion of our results

Variations on list elicitation have been extremely popular in recent years (e.g. Holt and Laury (2002), Andersen, Harrison, Lau, and Rutström (2008), Bruhin, Fehr-Duda, and Epper (2010), see also Andersen, Harrison, Lau, and Rutström (2006)). Probability list elicitation was first used (without incentives) by Davidson, Suppes, and Siegel (1957), revisited by McCord and De Neufville (1986), and was revisited (with the RIS) by Andreoni and Sprenger (2013) and **?**.

List elicitation supposedly provides an efficient method for collecting precise data on individual preferences. But its attractiveness relies on the assumption that the isolation hypothesis holds. Without the assumption that isolation holds, one must take into account the structure of the list and the RIS in order to interpret subjects' choices in terms of their preferences, which complicates the inferences that can be drawn from choice data collected using list elicitation.

Our results are clear evidence against the isolation hypothesis. Our results are in the direction consistent with non-expected utility theory combined with assumptions about how subjects view and evaluate the their choices and the external randomizing device as a compound lottery. This theoretical approach is attractive, since it already provides a unified understanding of preference reversals in BDM and of the Allais paradox. However, our experimental design cannot rule out other possible explanations for our findings.

We have suggested two different ways in which non-expected utility preferences would lead to the biases we observe in our experiment. Each of these explanations has different implications for how we might recover preferences from list elicitation. Given that there are other possible explanations for our findings, we do not see any way to draw unambiguous inferences about preferences from list elicitation. Future work might distinguish between different possible explanations, as Keller, Segal, and Wang (1993) do for explanations of preference reversals in BDM.

The experimental economics literature has developed many other ways of eliciting information about preferences. Alternatives to list elicitation include binary choice (Hey and Orme, 1994) and convex budget sets (Choi, Fisman, Gale, and Kariv, 2007). A typical implementation of each of these designs still uses RIS to select one of multiple questions to determine payment. However, in these designs choice tasks are always displayed on separate screens, with no RIS applying within a screen, which might induce subjects to isolate their choices.

4.4.3 Experiments on Mechanical Turk

As a large online labour market, Mechanical Turk provides a convenient way to recruit and pay subjects over the internet. Mechanical Turk allows researchers to economize on costs and experiment on a different population from undergraduates. Mechanical Turk has been advocated as a platform for recruiting subjects by psychologists studying judgement and decision-making (Mason and Suri (2011), Paolacci, Chandler, and Ipeirotis (2010), Buhrmester, Kwang, and Gosling (2011)), political scientists (Berinsky, Huber, Lenz, et al., 2012), and economists (Horton, Rand, and Zeckhauser, 2011). A potential downside of running on experiments on Mechanical Turk is that subjects complete the experiment from their home computer, and not in a controlled lab environment, making it difficult to know for sure who the subjects really are and how much attention they're paying to the tasks. Paolacci, Chandler, and Ipeirotis (2010) find that the population of US-based turkers who participate in experiments is heterogeneous and is more representative of the US population than typical undergraduate samples, and that turkers pay as much attention to experimental tasks as undergraduates in a lab. Paolacci, Chandler, and Ipeirotis (2010) and Horton, Rand, and Zeckhauser (2011) show that some standard experimental results in the judgement and decision-making literature can be qualitatively and quantitatively replicated using turkers. Our paper also replicates our result with a set of student subjects.

Bibliography

Abeler, J., A. Falk, L. Götte, and D. Huffman (2011): "Reference points and effort provision," *American Economic Review*, 101(2), 470–492.

Aliprantis, C. D., and K. C. Border (1999): *Infinite dimensional analysis: a hitchhiker's guide*. Springer.

Allais, M. (1953): "Le comportement de l'homme rationnel devant le risque: Critique des postulats et axiomes de l'école Américaine," *Econometrica*, 21(4), 503–546.

Andersen, S., G. Harrison, M. Lau, and E. Rutström (2006): "Elicitation using multiple price list formats," *Experimental Economics*, 9(4), 383–405.

——— (2008): "Eliciting risk and time preferences," *Econometrica*, 76(3), 583–618.

Andreoni, J., and C. Sprenger (2013): "Uncertainty Equivalents: Linear Tests of the Independence Axiom," Working Paper, University of California, San Diego.

Arrow, K. (1959): "Rational choice functions and orderings," *Economica*, 26(102), 121–127.

Au, P., and K. Kawai (2011): "Sequentially rationalizable choice with transitive rationales," *Games and Economic Behavior*, 73(2), 608–614.

Barberis, N., M. Huang, and R. Thaler (2006): "Individual preferences, monetary gambles, and stock market participation: A case for narrow framing," *American Economic Review*, 96(4), 1069–1090.

Becker, G., M. DeGroot, and J. Marschak (1964): "Measuring utility by a single-response sequential method," *Behavioral Science*, 9(3), 226–232.

Berinsky, A., G. Huber, G. Lenz, et al. (2012): "Evaluating Online Labor Markets for Experimental Research: Amazon.com's Mechanical Turk," *Political Analysis*, 20(3), 351–368.

Bruhin, A., H. Fehr-Duda, and T. Epper (2010): "Risk and rationality: Uncovering heterogeneity in probability distortion," *Econometrica*, 78(4), 1375–1412.

Buhrmester, M., T. Kwang, and S. Gosling (2011): "Amazon's Mechanical Turk: A New Source of Inexpensive, Yet High-Quality, Data?," *Perspectives on Psychological Science*, 6(1), 3–5.

Camerer, C. F., and T. H. Ho (1994): "Violations of the betweenness axiom and nonlinearity in probability," *Journal of Risk and Uncertainty*, 8(2), 167–196.

Carbajal, J., and J. Ely (2012): "Optimal Contracts for Loss Averse Consumers," Working Paper, University of Queensland.

Card, D., and G. Dahl (2011): "Family Violence and Football: The Effect of Unexpected Emotional Cues on Violent Behavior," *Quarterly Journal of Economics*, 126(1), 103–143.

Carlin, P. S. (1992): "Violations of the reduction and independence axioms in Allais-type and common-ratio effect experiments," *Journal of Economic Behavior & Organization*, 19(2), 213–235.

Cerreia-Vioglio, S., D. Dillenberger, and P. Ortoleva (2013): "Cautious Expected Utility and the Certainty Effect," Working Paper, University of Pennsylvania.

Chateauneuf, A., J. Eichberger, and S. Grant (2007): "Choice under uncertainty with the best and worst in mind: Neo-additive capacities," *Journal of Economic Theory*, 137(1), 538–567.

Cherepanov, V., T. Feddersen, and A. Sandroni (Forthcoming): "Rationalization," *Theoretical Economics*.

Chetty, R. (2006): "A new method of estimating risk aversion," *American Economic Review*, 96(5), 1821–1834.

Chew, Soo Hong, E. K., and Z. Safra (1987): "Risk aversion in the theory of expected utility with rank dependent probabilities," *Journal of Economic Theory*, 42(2), 370–381.

Chew, S. (1983): "A generalization of the quasilinear mean with applications to the measurement of income inequality and decision theory resolving the Allais paradox," *Econometrica*, pp. 1065–1092.

Chew, S. H. (1989): "Axiomatic utility theories with the betweenness property," *Annals of Operations Research*, 19(1), 273–298.

Chew, S. H., L. G. Epstein, and U. Segal (1991): "Mixture symmetry and quadratic utility," *Econometrica*, 59(1), 139–163.

Choi, S., R. Fisman, D. Gale, and S. Kariv (2007): "Consistency and heterogeneity of individual behavior under uncertainty," *American Economic Review*, 97(5), 1921–1938.

Cox, J., V. Sadiraj, and U. Schmidt (2011): "Paradoxes and mechanisms for choice under risk," Working Paper, Georgia State University.

Cox, J. C., and V. Sadiraj (2006): "Small-and large-stakes risk aversion: Implications of concavity calibration for decision theory," *Games and Economic Behavior*, 56(1), 45–60.

Crawford, V., and J. Meng (2011): "New York City Cab Drivers' Labor Supply Revisited: Reference-Dependent Preferences with Rational Expectations Targets for Hours and Income," *American Economic Review*, 101(5), 1912–1932.

Cubitt, Robin P., C. S., and R. Sugden (1998): "On the validity of the random lottery incentive system," *Experimental Economics*, 1(2), 115–131.

Davidson, D., P. Suppes, and S. Siegel (1957): *Decision making: an experimental approach*. Stanford University Press.

DEKEL, E. (1986): "An axiomatic characterization of preferences under uncertainty: weakening the independence axiom," *Journal of Economic Theory*, 40(2), 304–318.

DILLENBERGER, D. (2010): "Preferences for one-shot resolution of uncertainty," *Econometrica*, 78(6), 1973–2004.

DILLENBERGER, D., AND U. SEGAL (2012): "Recursive Ambiguity and Machina's Puzzles," Working Paper, University of Pennsylvania.

ELIAZ, K., AND R. SPIEGLER (2013): "Reference Dependence and Labor-Market Fluctuations," Working Paper, University College of London.

EPSTEIN, L., AND S. ZIN (1989): "Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework," *Econometrica*, pp. 937–969.

ERICSON, K., AND A. FUSTER (2011): "Expectations as endowments: Evidence on reference-dependent preferences from exchange and valuation experiments," *Quarterly Journal of Economics*, 126(4), 1879–1907.

FISHBURN, P. C. (1970): *Utility theory for decision making*. Wiley.

GRANT, S., AND A. KAJII (1998): "AUSI expected utility: an anticipated utility theory of relative disappointment aversion," *Journal of Economic Behavior & Organization*, 37(3), 277–290.

GREINER, B. (2004): "The Online Recruitment System ORSEE 2.0 - A Guide for the Organization of Experiments in Economics," Working Paper, University of Cologne.

GRETHER, D., AND C. PLOTT (1979): "Economic theory of choice and the preference reversal phenomenon," *American Economic Review*, 69(4), 623–638.

GUISO, L., T. JAPPELLI, AND L. PISTAFERRI (2002): "An empirical analysis of earnings and employment risk," *Journal of Business and Economic Statistics*, 20(2), 241–253.

GUL, F. (1991): "A theory of disappointment aversion," *Econometrica*, 59(3), 667–686.

GUL, F., AND W. PESENDORFER (2008): "The case for mindless economics," in *The Foundations of Positive and Normative Economics: A Handbook*, ed. by A. Caplin, and A. Schotter. Oxford University Press New York.

HALEVY, Y. (2007): "Ellsberg revisited: An experimental study," *Econometrica*, 75(2), 503–536.

HARRISON, G., AND J. SWARTHOUT (2012): "The independence axiom and the bipolar behaviorist," Working Paper, Georgia State University.

HEIDHUES, P., AND B. KŐSZEGI (2008): "Competition and price variation when consumers are loss averse," *American Economic Review*, pp. 1245–1268.

——— (Forthcoming): "Regular prices and sales," *Theoretical Economics*.

HERWEG, F., D. MULLER, AND P. WEINSCHENK (2010): "Binary payment schemes: Moral hazard and loss aversion," *American Economic Review*, 100(5), 2451–2477.

HEY, J., AND C. ORME (1994): "Investigating generalizations of expected utility theory using experimental data," *Econometrica*, pp. 1291–1326.

HOLT, C. (1986): "Preference reversals and the independence axiom," *American Economic Review*, 76(3), 508–515.

HOLT, C. A., AND S. K. LAURY (2002): "Risk aversion and incentive effects," *American Economic Review*, 92(5), 1644–1655.

HORAN, S. (2012): "A Simple Model of Two-Stage Maximization," Working Paper, Université du Québec à Montréal.

HORTON, J., D. RAND, AND R. ZECKHAUSER (2011): "The online laboratory: conducting experiments in a real labor market," *Experimental Economics*, 14(3), 399–425.

KAHNEMAN, D., J. KNETSCH, AND R. THALER (1990): "Experimental tests of the endowment effect and the Coase theorem," *Journal of Political Economy*, pp. 1325–1348.

Kahneman, D., and A. Tversky (1979): "Prospect theory: an analysis of decision under risk," *Econometrica*, 47(2), 263–291.

Karle, H., and M. Peitz (2012): "Pricing and Information Disclosure in Markets with Loss-Averse Consumers," Working Paper, University of Mannheim.

Karni, E., and Z. Safra (1987): ""Preference reversal" and the observability of preferences by experimental methods," *Econometrica*, 55(3), 675–685.

Keller, L., U. Segal, and T. Wang (1993): "The Becker-DeGroot-Marschak mechanism and generalized utility theories: Theoretical predictions and empirical observations," *Theory and Decision*, 34(2), 83–97.

Kőszegi, B. (2010): "Utility from anticipation and personal equilibrium," *Economic Theory*, 44(3), 415–444.

Kőszegi, B., and M. Rabin (2006): "A model of reference-dependent preferences," *Quarterly Journal of Economics*, 121(4), 1133–1165.

——— (2007): "Reference-dependent risk attitudes," *American Economic Review*, 97(4), 1047–1073.

——— (2009): "Reference-dependent consumption plans," *American Economic Review*, 99(3), 909–936.

Kreps, D., and E. Porteus (1978): "Temporal resolution of uncertainty and dynamic choice theory," *Econometrica*, 46(1), 185–200.

Machina, M. (1982): ""Expected Utility" Analysis without the Independence Axiom," *Econometrica*, 50(2), 277–323.

——— (1989): "Dynamic consistency and non-expected utility models of choice under uncertainty," *Journal of Economic Literature*, 27(4), 1622–1668.

Manzini, P., and M. Mariotti (2007): "Sequentially rationalizable choice," *American Economic Review*, 97(5), 1824–1839.

Masatlioglu, Y., and E. Ok (2005): "Rational choice with status quo bias," *Journal of Economic Theory*, 121(1), 1–29.

——— (2012): "A Canonical Model of Choice with Initial Endowments," Working Paper, University of Michigan.

Mason, W., and S. Suri (2011): "Conducting behavioral research on Amazon's Mechanical Turk," *Behavior Research Methods*, 44(1), 1–23.

McCord, M., and R. De Neufville (1986): ""Lottery Equivalents": Reduction of the Certainty Effect Problem in Utility Assessment," *Management Science*, 32(1), 56–60.

Munro, A., and R. Sugden (2003): "On the theory of reference-dependent preferences," *Journal of Economic Behavior & Organization*, 50(4), 407–428.

Neilson, W. S. (2001): "Calibration results for rank-dependent expected utility," *Economics Bulletin*, 4(10), 1–5.

Ok, E. (2012): "Elements of Order Theory," Book Draft, New York University.

Ok, E., P. Ortoleva, and G. Riella (2012): "Revealed (p)reference theory," Working Paper, California Institute of Technology.

Ortoleva, P. (2013): "The Price of Flexibility: Towards a Theory of Thinking Aversion," *Journal of Economic Theory*, 148(3), 903–934.

Paolacci, G., J. Chandler, and P. Ipeirotis (2010): "Running experiments on Amazon Mechanical Turk," *Judgment and Decision Making*, 5(5), 411–419.

Pope, D., and M. Schweitzer (2011): "Is Tiger Woods loss averse? Persistent bias in the face of experience, competition, and high stakes," *American Economic Review*, 101(1), 129–157.

Prelec, D. (1998): "The probability weighting function," *Econometrica*, 66(3), 497–527.

Quiggin, J. (1982): "A theory of anticipated utility," *Journal of Economic Behavior & Organization*, 3(4), 323–343.

Rabin, M. (2000): "Risk Aversion and Expected-utility Theory: A Calibration Theorem," *Econometrica*, 68(5), 1281–1292.

Richter, M. (1966): "Revealed preference theory," *Econometrica*, 34(3), 635–645.

Sadiraj, V. (2012): "Probabilistic Risk Attitudes and Local Risk Aversion: a Paradox," Working Paper, Georgia State University.

Safra, Z., and U. Segal (1998): "Constant risk aversion," *Journal of Economic Theory*, 83(1), 19–42.

Safra, Z., and U. Segal (2002): "On the Economic Meaning of Machina's Frechet Differentiability Assumption," *Journal of Economic Theory*, 104(2), 450–461.

Safra, Z., and U. Segal (2008): "Calibration Results for Non-Expected Utility Theories," *Econometrica*, 76(5), 1143–1166.

Safra, Z., U. Segal, and A. Spivak (1990): "Preference reversal and nonexpected utility behavior," *The American Economic Review*, 80(4), 922–930.

Sagi, J. (2006): "Anchored preference relations," *Journal of Economic Theory*, 130(1), 283–295.

Samuelson, P. A. (1952): "Probability, utility, and the independence axiom," *Econometrica*, 20(4), 670–678.

Samuelson, P. A. (1963): "Risk and uncertainty: A fallacy of large numbers," *Scientia*, 98(4-5), 108–113.

Segal, U. (1987a): "Some remarks on Quiggin's anticipated utility," *Journal of Economic Behavior & Organization*, 8(1), 145–154.

——— (1987b): "The Ellsberg paradox and risk aversion: An anticipated utility approach," *International Economic Review*, 28(1), 175–202.

——— (1988): "Does the preference reversal phenomenon necessarily contradict the independence axiom?," *American Economic Review*, 78(1), 233–236.

Segal, U. (1990): "Two-stage lotteries without the reduction axiom," *Econometrica*, 58(2), 349–377.

Segal, U., and A. Spivak (1990): "First order versus second order risk aversion," *Journal of Economic Theory*, 51(1), 111–125.

Sen, A. (1971): "Choice functions and revealed preference," *Review of Economic Studies*, 38(3), 307–317.

Starmer, C., and R. Sugden (1991): "Does the random-lottery incentive system elicit true preferences? An experimental investigation," *American Economic Review*, 81(4), 971–978.

Sydnor, J. (2010): "(Over) insuring modest risks," *American Economic Journal: Applied Economics*, 2(4), 177–199.

Tversky, A., and D. Kahneman (1991): "Loss aversion in riskless choice: a reference-dependent model," *Quarterly Journal of Economics*, 106(4), 1039–1061.

——— (1992): "Advances in prospect theory: cumulative representation of uncertainty," *Journal of Risk and Uncertainty*, 5(4), 297–323.

Wakker, P. P. (2010): *Prospect Theory: For Risk and Ambiguity*. Cambridge University Press, Cambridge, UK.

Wang, T. (1993): "Lp-Fréchet Differentiable Preference and "Local Utility" Analysis," *Journal of Economic Theory*, 61(1), 139–159.

Webb, C. S., and H. Zank (2011): "Accounting for optimism and pessimism in expected utility," *Journal of Mathematical Economics*, 47(6), 706–717.

Yaari, M. (1987): "The dual theory of choice under risk," *Econometrica*, 55(1), 95–115.

Appendix A

Proofs for Chapter 2

Lemma A.1. *For any two sets* D, D' *and any asymmetric binary relation* P*,* $m(D,\, P) \cup m(D',\, P) \supseteq m(D \cup D',\, P)$*.*

Proof. Suppose $p \in m(D \cup D', P) \cap D$.

$\implies \nexists q \in D \cup D'$ s.t. qPp.

$\implies \nexists q \in D$ s.t. qPp

$\implies p \in m(D, P)$.

If $p \in m(D \cup D', P) \cap D'$, an analogous result would follow.

Thus $p \in m(D \cup D', P)$ implies $p \in m(D, P) \cup m(D', P)$.

$\implies m(D,\, P) \cup m(D',\, P) \supseteq m(D \cup D',\, P)$ □

Results on IIA Independence and IIA Independence 2.

Lemma A.2. *Suppose Expansion and Weak RARP hold. If* $p \in c(D)$*,* $p \notin c(D \cup q) \ni r$*, and* $p\tilde{\tilde{W}}r$*, then* $\nexists D_{pq}$ *such that* $p \in c(D_{pq})$*.*

Proof. If $\exists D_{pq}$ such that $p \in c(D_{pq})$ then by Expansion, $p \in c(D \cup D_{pq})$. Since $D \cup q \subseteq D \cup D_{pq}$ and $r \in c(D \cup q)$ with $p\tilde{\tilde{W}}r$, it follows by Weak RARP that $p \in c(D \cup q)$, a contradiction. Thus no such D_{pq} can exist. □

Lemma A.3. *Suppose Expansion and Sen's α hold. If* $p \in c(D)$*,* $p \notin c(D \cup q)$ *, then* $\nexists D_{pq}$ *such that* $p \in c(D_{pq})$*.*

Proof. If $p \in c(D) \cap D_{pq}$ then by Expansion, $p \in c(D \cup D_{pq})$. Then by Sen's α, $p \in c(D \cup q)$. This proves the claim. □

Proof of Proposition 2.2.

(i) $\iff$ (iii)

Let P_1, P_2 denote the asymmetric part of relations $\bar{P}_1$, $\bar{P}_2$ that form a transitive short-listing representation. By definition, $m(D,\, P_i) = m(D,\, \bar{P}_i)$ for $i = 1, 2$ and for any D.

Necessity of Expansion. $p \in c(D)$ and $p \in c(D')$ implies:

(i) $p \in m(D, P_1)$ and $p \in m(D', P_1)$

$\implies \nexists q \in D$ s.t. qP_1p and $\nexists q \in D'$ s.t. qP_1p

$\implies \nexists q \in D \cup D'$ s.t. qP_1p

$\implies\ p \in m(D \cup D',\, P_1)$

(ii) $p \in m(m(D, P_1), P_2)$ and $p \in m(m(D', P_1), P_2)$

$\implies \nexists q \in m(D, P_1)$ s.t. qP_2p and $\nexists q \in m(D', P_1)$ s.t. qP_2p

$\implies \nexists q \in m(D, P_1) \cup m(D', P_1)$ s.t. qP_2p

by Lemma A.1,

$\implies \nexists q \in m(D \cup D', P_1)$ s.t. qP_2p

By (i),

$\implies\ p \in m(m(D \cup D', P_1), P_2) = c(D \cup D')$

This implies that Expansion holds.

Necessity of Weak RARP. Suppose $q\tilde{\bar{W}}p$, and there are D, D' such that: $\{p, q\} \subseteq D \subseteq \bar{D}$ and $p \in c(D)$, $q \in c^U(\bar{D})$.

By definition of $q\tilde{\bar{W}}p$, there is a chain $q = r^0, r^1, ..., r^{n-1}, r^n = p$ such that for each $i \in \{1, ..., n\}$, there are $D^i, \bar{D}^i$ such that $\{r^{i-1}, r^i\} \subseteq D^i \subseteq \bar{D}^i$, $r^i \in c^U(\bar{D}^i)$ and $r^{i-1} \in c(D^i)$, or (if not) there is a net $\{\bar{D}^{i,\epsilon},\, D^{i,\epsilon}\} \to \bar{D}^i,\, D^i$ for which $r^i \in c^U(\bar{D}^{i,,\epsilon})$ and $r^{i-1} \in c(D^{i,\epsilon})$ $\forall \epsilon > 0$.

For each i, from the representation, it follows that:

$\implies r^i \in m(D^i,\, P_1)$

$\implies$ not $r^iP_2r^{i-1}$.

Since the transitive completion of P_2 is transitive, it follows that not qP_2p.

Since $q \in c^U(\bar{D})$, by continuity of P_1, $q \in m(\bar{D},\, P_1)$.

Since $q \in D \subseteq \bar{D}$ as well, $q \in m(D,\, P_1)$.

Since $p \in m(m(D, P_1), P_2)$, not pP_2q, and P_2 has a transitive completion, it follows that not $rP_2q\ \forall r \in m(D,\, P_1)$.

Thus, $q \in m(m(D,\, P_1),\, P_2) = c(D)$.

<u>Sufficiency</u>. Part of the idea of the proof follows Manzini and Mariotti (2007). The two rationales constructed here are not unique.

Define P_1 by:

$$qP_1p \text{ if } \nexists D_{pq} \,\text{s.t.}\; p \in c^U(D_{pq})$$

Define $\bar{P}_2$ by:

$$\bar{P}_2 = \tilde{\bar{W}}$$

Define P_2 as the asymmetric part of P_2.

First, show that P_1 and P_2 are appropriately continuous.

If pP_1q, $\nexists$ a net $\{D_{p^\epsilon q^\epsilon}\}_\epsilon \to D_{pq}$ with $p^\epsilon \in c(D_{p^\epsilon q^\epsilon})$ and $\max\,[d(p^\epsilon, p), d(q^\epsilon, q)] < \epsilon$ for each $\epsilon > 0$, since then we would have $p \in c^U(D_{pq})$ for some D_{pq}. Thus, $\exists \bar{\epsilon} > 0$ such that $\forall p^\epsilon \in N_p^{\bar{\epsilon}}$, $\forall q^\epsilon \in N_q^{\bar{\epsilon}}$, $p^\epsilon P_1 q^\epsilon$. This implies that P_1 has open better and worse than sets.

P_2 is continuous by construction.

Second, show $c(D) \subseteq m(m(D, P_1), P_2)$.

By definition of P_1, $p \in c(D)$ implies $p \in m(D, P_1)$.

Take any $q \in m(D, P_1)$. By the definition of P_1, $\forall r \in D$, $\exists D_{qr}$ such that $q \in c(D_{qr})$. Successively applying Expansion implies that $q \in c(\bigcup\limits_{r\in D} D_{qr})$. Since $D \subseteq \bigcup\limits_{r\in D} D_{qr}$ and $p \in c(D)$, it follows that $p\tilde{\bar{W}}q$, thus $p\bar{P}_2q$. Since this implies not qP_2p for any arbitrary $q \in m(D, P_1)$, it further follows that $p \in m(m(D, P_1), P_2)$.

Third, show $m(m(D, P_1), P_2) \subseteq c(D)$

Suppose $p \in m(m(D, P_1), P_2)$.

Then, $\forall r \in D$, $\exists D_{pr} :\; p \in c(D_{pr})$. By Expansion, $p \in c(\bigcup\limits_{r\in D} D_{pr})$.

Since $p \in m(m(D, P_1), P_2)$, it $p\tilde{\bar{W}}q\ \forall q \in c(D)$ by the definition of $\tilde{\bar{W}}$.

Thus by Weak RARP, $p \in c(D)$.

(ii) $\iff$ (iii) Consider a continuous PPE representation v that represents c, and a continuous and transitive shortlisting representation P_1, P_2.

Map between v and P_1 by:

$$qP_1p \iff v(q|p) > v(p|p)$$

Map between v and P_2 by:

$$qP_2p \iff v(q|q) > v(p|p)$$

Joint continuity of v will map to continuity of P_1 and P_2.

Notice that the mapping from P_1 to v only specifies $v(\cdot|p)$ partially; the mapping from P_2 to v imposes an continuous additive normalization on v.

Consider the following construction of v from P_1, P_2:

Let $u : \Delta \to \Re$ be a continuous utility function that represents P_2. Define $v(p|p) = u(p)\ \forall p \in \Delta$. Let $I(p) = \{q \in \Delta :\ (q,p) \in \text{cl}\{(\hat{q},\hat{p}) :\ \hat{q}P_1\hat{p}\}\backslash\{(\hat{q},\hat{p}) :\ \hat{q}P_1\hat{p}\}\}$.

The following definition of v is consistent with the mapping proposed above:

$$v(q|p) = \begin{cases} u(p) + d^H(\{q\}, I(p)) & \text{if } qP_1p \\ u(p) - d^H(\{q\}, I(p)) & \text{otherwise} \end{cases}$$

It can be verified that continuity of P_1 and u imply that v so constructed satisfies joint continuity.

□

Proof of Theorem 2.1.

Notation.

Let for $p, q \in \Delta$, let $D_{pq} \in \mathcal{D}$ denote an arbitrary choice set that contains p and q.

Sufficiency: Lemmas.

In the lemmas in this section, assume that c satisfies Expansion, Weak RARP, IIA Independence, Induced Reference Lottery Bias, and Transitive Limit.

Lemma A.4. *$\bar{R}_p$ is complete, transitive, and if there exists a net $\{p^\epsilon, q^\epsilon, r^\epsilon\} \to p, q, r$ with $q^\epsilon \bar{R}_{p^\epsilon} r^\epsilon$ for each term in the net, then $q\bar{R}_p r$.*

Proof. Transitivity of $\bar{R}_p$ follows by Transitive Limit.

For any net $\{p^\epsilon, q^\epsilon, r^\epsilon\} \to p, q, r$, non-emptiness of c implies that the net either has a convergent subnet $p^\delta, q^\delta, r^\delta$ in which $(1-\delta)p^\delta+\delta q^\delta \in c(\{(1-\delta)p^\delta+\delta q^\delta,\ (1-\delta)p^\delta+\delta r^\delta\}$ or in which $(1-\delta)p^\delta + \delta r^\delta \in c(\{(1-\delta)p^\delta + \delta q^\delta,\ (1-\delta)p^\delta + \delta r^\delta\}$ for each term in the subnet. Thus $\bar{R}_p$ is complete.

Take a net $\{p^\epsilon, q^\epsilon, r^\epsilon\} \to p, q, r$, for which $q^\epsilon \bar{R}_{p^\epsilon} r^\epsilon$ for each term in the net. By the definition of $\bar{R}_{p^\epsilon}$, for each ϵ there is a net $\{p^{\epsilon,\delta}, q^{\epsilon,\delta}, r^{\epsilon,\delta}\}_\delta \to p^\epsilon, q^\epsilon, r^\epsilon$ such that $(1-\delta)p^{\epsilon,\delta} + \delta q^{\epsilon,\delta} \in c((1-\delta)p^{\epsilon,\delta} + \delta\{q^{\epsilon,\delta}, r^{\epsilon,\delta}\})$ for each term in the net. Let $\bar{\delta}_\epsilon$ denote the largest element in the index set for $\{p^{\epsilon,\delta}, q^{\epsilon,\delta}, r^{\epsilon,\delta}\}_\delta$ and $\bar{\epsilon}$ the largest element in the index set for $\{p^\epsilon, q^\epsilon, r^\epsilon\}$. Take $\bar{\delta} := \bar{\delta}_{\bar{\epsilon}}$. For each $\delta < \bar{\delta}$, define ϵ_δ as a decreasing net such that for each $\delta < \bar{\delta}_{\epsilon_\delta}$. Then define $\{\hat{p}^\delta, \hat{q}^\delta, \hat{r}^\delta\} := \{p^{\epsilon_\delta,\delta}, q^{\epsilon_\delta,\delta}, r^{\epsilon_\delta,\delta}\}_\delta$. By construction, $\{\hat{p}^\delta, \hat{q}^\delta, \hat{r}^\delta\}$ establishes that $q\bar{R}_p r$. □

Let R_p denote the strict part of $\bar{R}_p$. Lemma A.5 shows that R_p satisfies the Independence Axiom.

For a binary relation R, say that R *satisfies the Independence axiom* if $qRr \iff (1-\alpha)s + \alpha q R (1-\alpha)s + \alpha r\, \forall \alpha \in (0,1). \forall s \in \Delta$.

Lemma A.5. *R_p satisfies the Independence Axiom if $p \in int\Delta$.*

Proof. Part I: suppose qR_pr, and take a $\alpha \in (0,1)$ and $s \in \Delta$.

Then,$\exists \bar{\delta}, \bar{\epsilon} > 0$ such that $\forall \epsilon \in (0, \bar{\epsilon})$, $\hat{p}, \hat{q}, \hat{r} \in N_p^{\bar{\delta}} \times N_q^{\bar{\delta}} \times N_r^{\bar{\delta}}$, $\{(1-\epsilon)\hat{p} + \epsilon\hat{q}\} = c((1-\epsilon)\hat{p} + \epsilon\{\hat{q},\ \hat{r}\})$.

Define $\bar{\delta}_\alpha = \min\left[\alpha\bar{\delta},\ (1-\alpha)\bar{\delta}\right]$.

Let $\hat{p}, \hat{s} \in N_p^{\bar{\delta}_\alpha} \times N_s^{\bar{\delta}_\alpha}$. Since $d^E(p,q) \le 1$, it follows that $d^E((1-\beta)\hat{p}+\beta\hat{s},\ p) \le (1-\beta)\bar{\delta}_\alpha + \beta$ by the triangle inequality. Thus if $\beta \le \bar{\beta}_\alpha := \frac{\bar{\delta}-\bar{\delta}_\alpha}{1-\bar{\delta}_\alpha}$, then $(1-\beta)\hat{p} + \beta\hat{s} \in N_p^{\bar{\delta}}$.

Then for any $\hat{q}, \hat{r} \in N_q^{\bar{\delta}} \times N_r^{\bar{\delta}}$, $\epsilon \in (0, \bar{\epsilon})$, and $\beta \in (0, \bar{\beta})$, $\{(1-\epsilon)\left((1-\beta)\hat{p} + \beta\hat{s}\right) + \epsilon\hat{q}\} = c((1-\epsilon)\left((1-\beta)\hat{p} + \beta\hat{s}\right) + \epsilon\{\hat{q},\ \hat{r}\})$. Define $\hat{\epsilon} := \frac{\epsilon}{\alpha}$ and $\beta^{\epsilon,\alpha} := \frac{\epsilon}{\alpha}\frac{1-\alpha}{1-\epsilon}$. Then $\forall\hat{\epsilon}$ such that $\alpha\hat{\epsilon} \in (0, \bar{\epsilon})$ and $\hat{\epsilon}\frac{1-\alpha}{1-\alpha\hat{\epsilon}} \in (0, \bar{\beta})$, it follows that $\{(1-\hat{\epsilon})\hat{p} + \hat{\epsilon}\left((1-\alpha)\hat{s} + \alpha\hat{q}\right)\} = c((1-\hat{\epsilon})\hat{p} + \hat{\epsilon}\left((1-\alpha)\hat{s} + \alpha\{\hat{q},\ \hat{r}\}\right))$. Since $N^{\bar{\delta}_\alpha}_{(1-\alpha)s+\alpha q} \subset (1-\alpha)N_s^{\bar{\delta}} + \alpha N_q^{\bar{\delta}}$ and $N^{\bar{\delta}_\alpha}_{(1-\alpha)s+\alpha q} \subset (1-\alpha)N_s^{\bar{\delta}} + \alpha N_q^{\bar{\delta}}$

It follows that $(1-\alpha)s + \alpha q R_p (1-\alpha)s + \alpha r$.

Part II: suppose $(1-\alpha)s + \alpha q R_p (1-\alpha)s + \alpha r$.

Recall that $N^{\bar{\delta}}_{(1-\alpha)s+\alpha q} \subseteq (1-\alpha)N_s^{\bar{\delta}} + \alpha N_q^{\bar{\delta}}$.

Then,$\exists\bar{\delta},\bar{\epsilon} > 0$ such that $N_p^{\bar{\delta}} \subset \mathrm{int}\Delta$ and $\forall \epsilon \in (0,\bar{\epsilon})$, $\hat{p},\hat{q},\hat{r},\hat{s} \in N_p^{\bar{\delta}} \times N_q^{\bar{\delta}} \times N_r^{\bar{\delta}} \times N_s^{\bar{\delta}}$, $\{(1-\epsilon)\hat{p}+\epsilon((1-\alpha)\hat{s}+\alpha\hat{q})\} = c((1-\epsilon)\hat{p}+\epsilon((1-\alpha)\hat{s}+\alpha\{\hat{q},\,\hat{r}\}))$.

Fix $\kappa \in (0,1)$. Fix $\hat{p},\hat{q},\hat{r},\hat{s} \in N_p^{\kappa\bar{\delta}} \times N_q^{\kappa\bar{\delta}} \times N_r^{\kappa\bar{\delta}} \times N_s^{\kappa\bar{\delta}}$.

Given $\epsilon \in (0,\bar{\epsilon})$, take $\gamma^{\epsilon,\alpha} := \epsilon\frac{1-\alpha}{1-\epsilon}$. If $\gamma < (1-\kappa)\bar{\delta}$, then $\hat{p}+\gamma^{\epsilon,\alpha}(\hat{p}-\hat{s}) \in N_p^{\bar{\delta}} \subseteq \Delta$. Then,

$(1-\epsilon)\left(\hat{p}+\gamma^{\epsilon,\alpha}(\hat{p}-\hat{s})\right)+\epsilon((1-\alpha)\hat{s}+\alpha\hat{q}) = c((1-\epsilon)\left(\hat{p}+\gamma^{\epsilon,\alpha}(\hat{p}-\hat{s})\right)+\epsilon((1-\alpha)\hat{s}+\alpha\{\hat{q},\,\hat{r}\}))$

$\iff (1-\alpha\epsilon)\hat{p}+\alpha\epsilon\{\hat{q}\} = c((1-\alpha\epsilon)\hat{p}+\alpha\epsilon\{\hat{q},\,\hat{r}\})$

Since the above holds $\forall \hat{p},\hat{q},\hat{r},\hat{s},\epsilon \in N_p^{\kappa\bar{\delta}} \times N_q^{\kappa\bar{\delta}} \times N_r^{\kappa\bar{\delta}} \times N_s^{\kappa\bar{\delta}} \times (0,\bar{\epsilon})$ it follows that qR_pr. □

Lemma A.6. *$\bar{R}_p$ satisfies the Independence Axiom if $p \in int\Delta$.*

Proof. I already have a proof that R_p satisfies the Independence Axiom.

Suppose that $q\bar{R}_pr$ and take $(1-\alpha)s+\alpha q$ and $(1-\alpha)s+\alpha r$.

If it is not the case that $(1-\alpha)s+\alpha q\bar{R}_p(1-\alpha)s+\alpha r$, then $(1-\alpha)s+\alpha rR_p(1-\alpha)s+\alpha q$.

Then it follows by Lemma A.5 that rR_pq, which contradicts that $q\bar{R}_pr$. □

Define $q\bar{\bar{R}}_pr$ if either:

(i) $p \in \mathrm{int}\Delta$ and $q\bar{R}_pr$

(ii) $p \notin \mathrm{int}\Delta$, and $\exists \alpha, s \in (0,1)\times\Delta$ such that $(1-\alpha)s+\alpha q\bar{R}_p(1-\alpha)s+\alpha r$

(iii) $\exists \alpha, s, \hat{q}, \hat{r}$ such that $q = (1-\alpha)s+\alpha\hat{q}$, $r = (1-\alpha)s+\alpha\hat{r}$, and $\hat{q}\bar{R}_p\hat{r}$

The relation $\bar{\bar{R}}_p$ is the minimal extension of $\bar{R}_p$ that respects with the Independence Axiom for all $p \in \Delta$.

By construction, $\bar{\bar{R}}_p$ satisfies the joint continuity properties in Lemma A.4 as well.

Lemma A.7. *For each $p \in \Delta$, there exists a vector $\hat{u}^p \in \Re^N$ such that $q\bar{\bar{R}}_pr \iff q\cdot\hat{u}^p \geq r\cdot\hat{u}^p$.*

Proof. Lemma A.4. shows $\bar{R}_p$ is complete, and transitive. By construction, $\bar{\bar{R}}_p$ satisfies the Independence axiom. The joint continuity property on $\bar{R}_p$ in Lemma A.4 then implies the notion of mixture continuity required (condition 3) to apply Fishburn's (1970) Theorem 8.2. □

Say that a vector u^p is *flat* if $\max_i u_i^p = \min_i u_i^p$. Let $F := \{p \in \Delta : \ u^p \text{ is flat}\}$.

Lemma A.8. *Suppose u^p is not flat. Then, there is an ϵ neighbourhood N_p^ϵ of p such that $\forall \hat{p} \in N_p^\epsilon$, $u^{\hat{p}}$ is not flat.*

Proof. Suppose there is a net $\hat{p}^\epsilon$ such that $\hat{p}^\epsilon \in N_p^\epsilon$ and $u^{\hat{p}^\epsilon}$ is flat. Since $u^{\hat{p}^\epsilon}$ must represent $\bar{R}_{\hat{p}^\epsilon}$, it follows that $q\bar{R}_{\hat{p}^\epsilon}r \ \forall q, r \in \Delta$ and for each $\hat{p}^\epsilon$. By Lemma A.4, it follows that $q\bar{R}_p r$. It follows that u^p must be flat as well, a contradiction. □

Let $\hat{u}^p$ denote a vector that provides an EU representarion for $\bar{\bar{R}}_p$ (i.e. $q \cdot \hat{u}^p \geq r \cdot \hat{u}^p \iff q\bar{\bar{R}}_p r \ \forall q, r \in \Delta$). For all p such that p is non-flat, define:

$$u^p := \frac{d^H(\{p\},\, F)}{\max_i \left[\hat{u}_i^p - \sum_j \hat{u}_j^p \right]} \left(\hat{u}^p - \sum_j \hat{u}_j^p \right) \tag{A.1}$$

If $\hat{u}^p$ is flat, define u^p as the zero vector.

By Lemma A.8 and the EU theorem, u^p provides an EU representation for $\bar{\bar{R}}_p$.

Lemma A.9. *If $p^\epsilon \to p$, then $u^{p^\epsilon} \to u^p$*

Proof. If u^p is flat, then $d^H(\{p^\epsilon\},\, F) \to 0$ as $\epsilon \to 0$, thus $u^{p^\epsilon} \to u^p$.

Now suppose that u^p is non-flat. Suppose $p^\epsilon \to p$ but for convergent subnet $\{p^{\epsilon'}\}$ of $\{p^\epsilon\}$, $u^{p^{\epsilon'}} \to \bar{u}^p \neq u^p$. Since $u^{p^{\epsilon'}}$ represents $\bar{R}_{p^{\epsilon'}}$, by the joint continuity property in Lemma A.4., it follows that $\bar{u}^p$ ranks $q \sim r$ if and only if u^p ranks $q \sim r$. Since u^p and $\bar{u}^p$ must satisfy the same normalizations, they must coincide by the uniqueness result of the EU theorem.

□

Define $v : \Delta \times \Delta \to \Re$ by $v(q|p) := q \cdot u^p$

Lemma A.10. *v is jointly continuous.*

Proof. $v(q|p) = q \cdot u^p = \sum_i q_i u_i^p$ and u^p is continuous as a function of p, and joint continuity of the sum $\sum_i q_i u_i^p$ in q and u^p is a standard exercise. □

Lemma A.11 shows that Limit Consistency is implied by the axioms assumed in Theorem 2.1.

Lemma A.11. *The axioms in Theorem 2.1 imply Limit Consistency.*

Proof. Part 1. Suppose $\{q\} = m(D, R_p) \neq \{p\} \subseteq c(D)$.

That is, $qR_pr\ \forall r \in D$.

Then $\forall r \in D\ \exists \bar{\alpha}_r > 0$ such that $\forall \alpha \in (0, \bar{\alpha}_r)$, $(1-\alpha)p + \alpha q = c((1-\alpha) + \alpha\{q,\ r\})$.

Since D is finite, $\min\limits_{r\in D} \bar{\alpha}_r > 0$.

By Expansion, $\forall \alpha \in (0, \min\limits_{r\in D} \bar{\alpha}_r)$, $(1-\alpha)p + \alpha q \in c((1-\alpha) + \alpha D)$.

By Induced Reference Lottery Bias, $\forall \alpha \in (0, \min\limits_{r\in D} \bar{\alpha}_r)\ p \in c((1-\alpha) + \alpha D)$. Thus $p\tilde{\bar{W}}(1-\alpha)p + \alpha q$. Weak RARP then implies that $p \in c((1-\alpha)p + \alpha\{p,\ q\})\ \forall \alpha \in (0, \min\limits_{r\in D} \bar{\alpha}_r)$. This implies that $p\bar{R}_pq$, a contradiction.

Part 2. Suppose there are elements $q^1, ..., q^l \in D$ such that q^iR_pp for each $i = 1, ..., l$.

Suppose $q^i \in m(D, R_p)$, and let $\hat{D} := D \backslash m(D, R_p)$.

Then by the previous result $\forall i = 1, ..., l$, $\exists \bar{\alpha}_i > 0$ such that $\forall \alpha \in (0, \bar{\alpha}_i)$, $(1-\alpha)p + \alpha q^i \in c((1-\alpha)p + \alpha\left(\hat{D} \cup q^i\right))$.

Since $\{q^1, ..., q^l\}$ is finite and each $\bar{\alpha}_i > 0$, $\min\limits_i \bar{\alpha}_i > 0$.

For each $\alpha \in (0, 1)$, $c((1-\alpha)p + \alpha\{q^1,\ ...,\ q^l\})$ is non-empty.

For $\hat{q}$ such that $(1-\alpha)p + \alpha\hat{q} \in c((1-\alpha)p + \alpha\{q^1,\ ...,\ q^l\})$, Expansion implies that
$(1-\alpha)p + \alpha\hat{q} \in c(\left((1-\alpha)p + \alpha\{q^1,\ ...,\ q^l\}\right) \cup \left((1-\alpha)p + \alpha\left(\hat{D} \cup \hat{q}\right)\right))$
$= c((1-\alpha)p + \alpha D)$.

Thus $\forall \alpha \in (0, \min\limits_i \bar{\alpha}_i)$, $((1-\alpha)p + \alpha\{q^1, ..., q^l\}) \cap c((1-\alpha)p + \alpha D) \neq \emptyset$.

It follows that for at least one $\hat{q} \in \{q^1,\ ...,\ q^l\}$, $\forall \bar{\alpha} \in (0, \min\limits_i \bar{\alpha}_i)$, $\exists \alpha < (0, \bar{\alpha})$ such that $((1-\alpha)p + \alpha\hat{q} \in c((1-\alpha)p + \alpha D)$.

Since $p \in c((1-\alpha)p + \alpha D)\ \forall \alpha \in (0, 1)$ by Induced Reference Lottery Bias, it follows that $p\tilde{\bar{W}}(1-\alpha)p + \alpha\hat{q}$ whenever $(1-\alpha)p + \alpha\hat{q} \in c((1-\alpha)p + \alpha D)$. For such α, it further follows by Weak RARP that $p \in c((1-\alpha)p + \alpha\{p,\ \hat{q}\})$. This contradicts that $\hat{q}R_pp$. $\square$

Define $\hat{PE}(D) = \{p \in D :\ p\bar{R}_pq\ \forall q \in D\}$.

Define $P\hat{P}E(D) = \{p \in \hat{PE}(D) :\ \nexists q \in \hat{PE}(D)$ s.t. $q\tilde{\bar{W}}p\}$.

Lemma A.9 establishes that $p \in c(\{p,\ q\})$ implies $p \in P\hat{P}E(\{p,\ q\})$.

Lemma A.12. *If $q\bar{R}_q p$ and $p \in c(\{p,\, q\})$, then $p\tilde{\tilde{W}}q$.*

Proof. If $\exists D_{pq}$ such that $q \in c^U(D_{pq})$ then the result follows automatically. Similarly if there exists a chain $p = r^0, r^1, ..., r^n = q$ such that $r^{i-1}\tilde{\tilde{W}}r^i$ for $i = 1, ..., n$.

If $\exists p^\epsilon, q^\epsilon$ that establish $q\bar{R}_q p$, then if for some such sequence, $p^\epsilon \in c(\{p^\epsilon,\, q^\epsilon\})$ for a convergent subsequence of p^ϵ, q^ϵ, then $p^\epsilon \tilde{\tilde{W}} q^\epsilon$ for such pairs. Then, continuity of $\tilde{\tilde{W}}$ implies that $p\tilde{\tilde{W}}q$.

So suppose instead that for each sequence p^ϵ, q^ϵ that establishes that $q\bar{R}_q p$, $q^\epsilon = c(\{p^\epsilon,\, q^\epsilon\})$ except on a non-convergent subsequence of p^ϵ, q^ϵ. This implies that $q \in c^U(\{p,\, q\})$. Then by the definition of $\tilde{\tilde{W}}$, $p\tilde{\tilde{W}}q$. □

Lemma A.10 establishes that $p \in P\hat{P}E(\{p,\, q\})$ implies $p \in c(\{p,\, q\})$.

Lemma A.13. *If $p\bar{R}_p q$ and $p\tilde{\tilde{W}}q$, then $p \in c(\{p,\, q\})$.*

Proof. Since $\{p\} = c(\{p\})$, if $\{q\} = c(\{p,\, q\})$ and $p\tilde{\tilde{W}}q$ it would follow by IIA Independence and the definition of R_p that $qR_p p$. This would contradict the assumption that $p\bar{R}_p q$. Since $c(\{p,\, q\}) \neq \emptyset$, it then follows that $p \in c(\{p,\, q\})$. □

Lemmas A.11-A.12 establish that $P\hat{P}E(\{p,\, q,\, r\}) = c(\{p,\, q,\, r\})$ $\forall p, q, r \in D$.

Lemma A.14. *If $p \in c(\{p,\, q,\, r\})$ and $q\bar{R}_q p$ then $p\tilde{\tilde{W}}q$ or $r\bar{R}_q q$.*

Proof. Suppose $p \in c(\{p,\, q,\, r\})$ and $q\bar{R}_q p$.

If $p \in c(\{p,\, q\})$, then $p\tilde{\tilde{W}}q$ holds.

So suppose instead that $q = c(\{p,\, q\})$.

Then, if $q \in c(\{q,\, r\})$ it would follow by Expansion that $q \in c(\{p,\, q,\, r\})$. Since $p \in c(\{p,\, q,\, r\})$ as well, it follows that $p\tilde{\tilde{W}}q$; by Weak RARP, it follows that $p \in c(\{p,\, q\})$, a contradiction. Thus $r = c(\{q,\, r\})$.

By Lemma A.5, it follows that either $rR_q q$ or $r\tilde{\tilde{W}}q$; in the former case we're done, so suppose $r\tilde{\tilde{W}}q$ and that it is not the case that $rR_q q$.

If $p \in c(\{p,\, r\})$, then it follows that either $pR_r r$ or $p\tilde{\tilde{W}}r$. In the latter case, transitivity of $\tilde{\tilde{W}}$ implies $p\tilde{\tilde{W}}q$ and we're done, so suppose we have that $pR_r r$. Then by Limit Consistency, $p = c(\{p,\, r\})$.

To summarize, we now have that $q = c(\{p,\, q\})$, $p = c(\{p,\, r\}) = c(\{p,\, q,\, r\})$, and $r = c(\{q,\, r\})$. Then, by IIA Independence, it follows that $\exists\epsilon > 0 :\ \forall\alpha \in (0, 1)$, $\forall\hat{q} \in N_q^\epsilon$, $\forall\hat{r} \in N_r^\epsilon$, $\forall\hat{D} \supseteq \{\hat{q},\, (1-\alpha)\hat{q}+\alpha\hat{r}\}$, $\hat{q} \notin c(\hat{D})$. It follows that $rR_q q$, a contradiction.

It follows that either $rR_q q$ or $p\tilde{\bar{W}}q$. $\square$

Lemma A.15. *If $p \in P\hat{P}E(\{p,\, q,\, r\})$ then $p \in c(\{p,\, q,\, r\})$.*

Proof. Suppose $p \in P\hat{P}E(\{p,\, q,\, r\})$.

We know that $c(\{p,\, q,\, r\}) \neq \emptyset$. So it is sufficient to prove that $q \in c(\{p,\, q,\, r\}) \implies p \in c(\{p,\, q,\, r\})$ and similarly if $r \in c(\{p,\, q,\, r\})$.

Suppose $q \in c(\{p,\, q,\, r\})$; the argument starting from $r \in c(\{p,\, q,\, r\})$ is symmetric.

Then, $q\bar{R}_q p$ and $q\bar{R}_q r$ by Limit Consistency. Since $p \in P\hat{P}E(\{p,\, q,\, r\})$ and $q \in \hat{PE}(\{p,\, q,\, r\})$, it follows that $p\tilde{\bar{W}}q$. Then by Lemma A.10, since $p\bar{R}_p q$ as well, $p \in c(\{p,\, q\})$.

If $r \in c(\{p,\, q,\, r\})$ then a similar argument implies $p \in c(\{p,\, r\})$. Then by Expansion, $p \in c(\{p,\, q,\, r\})$.

If instead $r \notin c(\{p,\, q,\, r\})$, we have (recalling Lemma A.6) that either $p \in c(\{p,\, r\})$ or $r = c(\{p,\, r\})$. In the former case, Expansion implies $p \in c(\{p,\, q,\, r\})$. In the latter case, $r = c(\{p,\, r\})$. Recall that $p \in c(\{p,\, q\})$. If $p \notin c(\{p,\, q,\, r\})$ then $q = c(\{p,\, q,\, r\})$; by IIA Independence and the definition of R_p, it follows that $rR_p p$, a contradiction of the assumption that $p \in P\hat{P}E(\{p,\, q,\, r\})$.

It follows that $p \in P\hat{P}E(\{p,\, q,\, r\}) \implies p \in c(\{p,\, q,\, r\})$. $\square$

Remark. $P\hat{P}E(D) = P\hat{P}E(\hat{PE}(D))$

Lemma A.16. *Suppose we have established that $P\hat{P}E(D) = c(D)$ whenever $|D| < n$. If $\hat{PE}(D) = D$ and $|D| \leq n$, then $c(D) = P\hat{P}E(D)$.*

Proof. First, suppose $\hat{PE}(D) = D$.

Take $p \in P\hat{P}E(D)$. Then $p \in P\hat{P}E(D\backslash r)\ \forall \in D\backslash p$. Take any distinct $r, r' \in D\backslash p$, and then since $|D\backslash r| = |D\backslash r'| = n - 1 < n$, $p \in c(D\backslash r) \cap c(D\backslash r')$. By Expansion, it follows that $p \in c(D)$.

In the reverse, suppose $p \in c(D)$. Then if $q\tilde{\tilde{W}}r\ \forall r \in D$, since $\hat{PE}(D) = D$, it follows that $q \in c(\{q,\, r\})\ \forall r \in D$. By Expansion, it follows that $q \in c(D)$. Then since $p \in c(D)$ and $q \in c(D)$, $p\tilde{\tilde{W}}q$ by definition. Thus $p \in P\hat{P}E(D)$. □

Lemma A.14 establishes by induction that $c(D) = P\hat{P}E(D)$ for any $D \in \mathcal{D}$.

Lemma A.17. *Suppose $c(D) = P\hat{P}E(D)$ whenever $|D| < n$. Then, $c(D) = P\hat{P}E(D)$ whenever $|D| \leq n$ as well.*

Proof. Consider D with $|D| = n$ and $\hat{PE}(D) \neq D$. Partition D into $\hat{PE}(D)$ and $D \backslash \hat{PE}(D)$. The case where $\hat{PE}(D) = D$ was proven in Lemma A.9.

Since $|\hat{PE}(D)| \leq n - 1 < n$, $c(\hat{PE}(D)) = P\hat{P}E(\hat{PE}(D)) = P\hat{P}E(D)$.

Say that $q^0, q^1, ..., q^m$ form a *chain* if $q^i R_{q^{i-1}} q^{i-1}$ for $i = 1, ..., m$. Notice that if $q^0, ..., q^m$ form a chain, Limit Consistency implies that $q^m = c(\{q^0, ..., q^m\}) = \hat{PE}(D) = P\hat{P}E(D)$. So if the longest chain in D contains all elements of D, then $c(D) = P\hat{P}E(D)$.

Now suppose $p \in P\hat{P}E(D)$.

First, further suppose the longest chain in D has length $n - 1$; denote the chain $q^0, q^1, ..., q^{n-1}$. Then, $q^{n-1} = c(\{q^0, q^1, ..., q^{n-1}\})$ and since $q^0, q^1, ..., q^{n-1}$ is the longest chain in D and $p \in \hat{PE}(D)$, $\{p,\, q^{n-1}\} = \hat{PE}(D)$. Since $p \in P\hat{P}E(D)$, it follows that $p\tilde{\tilde{W}}q^{n-1}$; Lemmas A.8 and A.10, $p \in c(\{p,\, q^{n-1}\})$. Suppose $p \in c(\{p,\, q^k, ..., q^{n-1}\})$ for some $k \leq n-1$. Then, since if $p \notin c(\{p,\, q^{k-1}, ..., q^{n-1}\})$ it follows by IIA Independence and the definition of R_p that $q^{k-1}R_p p$, which contradicts that $p \in \hat{PE}(D)$. Thus it follows by induction that $p \in c(D)$.

Take an arbitrary chain $q^0, ..., q^m$ that cannot be extended further as a chain using elements of D. Since $q^0, ..., q^m$ cannot be extended, $q^m \in \hat{PE}(D)$. Since $p \in P\hat{P}E(D)$, $p\tilde{\tilde{W}}q^m$ and by Lemma A.8, $p \in c(\{p,\, q^m\})$. Suppose $p \in c(\{p,\, q^k, ..., q^m\})$ for some $k \leq m$. Then if $p \notin c(\{p,\, q^{k-1}, ..., q^m\})$ it follows by IIA Independence and the definition of R_p that $q^{k-1}R_p p$; this would which contradicts that $p \in \hat{PE}(D)$. Thus it follows by induction that $p \in c(\{p,\, q^0, ..., q^m\})$.

Notice that any element of $D \backslash \hat{PE}(D)$ is in a chain in D. Let $\hat{D}$ is the choice set formed by taking the union of $\{p\}$ and of the all of the choice sets formed by chains in D. Since for any chain $q^0, ..., q^m$ in D, $p \in c(\{p,\, q^0, ..., q^m\})$, $p \in c(\hat{D})$ follows by

Expansion. Since $p \in c(\hat{PE}(D))$ as well follows (because $|\hat{PE}(D)| < n$ or Lemma A.13 applies), it follows by Expansion that $p \in c(D)$. Thus $P\hat{P}E(D) \subseteq c(D)$.

In the reverse direction, now suppose $p \in c(D)$. By Limit Consistency, $p \in \hat{PE}(D)$. Since $c(D) \supseteq P\hat{P}E(D) = P\hat{P}E(\hat{PE}(D)) = c(\hat{PE}(D)) \neq \emptyset$, $\exists q \in c(D) \cap P\hat{P}E(D)$. Since $p, q \in c(D)$, $p\tilde{\bar{W}}q$. Thus $p \in P\hat{P}E(D)$. $\square$

Lemma A.15 relates the dislike of mixtures property to the Induced Reference Lottery Bias axiom.

Lemma A.18. *Induced Reference Lottery Bias implies that v dislikes mixtures.*

Proof. By the representation thus far, $c(D) = P\hat{P}E(D)$.

If $p \in P\hat{P}E(\{p, q\})$ then $v(p|p) \geq v(q|p)$ and either $v(p|p) \geq v(q|q)$ or $v(p|q) > v(q|q)$. Thus $v(p|p) \geq v(q|p)$ and $v(q|q) \leq \max[v(p|p), v(q|p)]$. Then the Induced Reference Lottery Bias axiom implies that then $p \in c((1-\alpha)p + \alpha D) = P\hat{P}E((1-\alpha)p + \alpha D)$, thus $v(p|p) \geq v((1-\alpha)p + \alpha q|p)$ and $v((1-\alpha)p + \alpha q|(1-\alpha)p + \alpha q) \leq \max[v(p|p), v((1-\alpha)p + \alpha q|p)]$. $\square$

Remark. $P\hat{P}E(D) = PPE_v(D)$

Necessity.

Proposition 2.2 implies that Expansion and Weak RARP are necessary conditions for any PPE representation.

Lemma A.19. *Suppose v represents c by a PPE representation. Then $p\tilde{\bar{W}}r$ implies that $v(p|p) \geq v(r|r)$.*

Proof. Suppose $p\tilde{\bar{W}}r$. If $\exists D, \bar{D}$ with $\{p, r\} \subseteq D \subseteq \bar{D}$ and $p \in c(D)$ and $r \in c(\bar{D})$ then it follows that $v(p|p) \geq v(r|r)$ since $r \in PE(\bar{D}) \cap D \subseteq PE(D)$ follows by the representation.

If instead there is a chain such that $p^{i-1}\tilde{\bar{W}}p^i$ for $i = 1, ..., n$ and $p^0 = p$, $p^n = r$, then it follows that $v(p^{i-1}|p^{i-1}) \geq v(p^i|p^i)$ for each i. Chaining these inequalities together, it follows that $v(p|p) \geq v(r|r)$. $\square$

Necessity of IIA Independence. Suppose $p\tilde{\tilde{W}}r$. Then by Lemma A.12, $v(p|p) \geq v(r|r)$. If $p \in PPE(D)$ and $p \notin PPE(D \cup q) \ni r$, then it follows that $v(q|p) > v(p|p)$. Since v is jointly continuous, $\exists \epsilon > 0$ such that $\forall \hat{p} \in N_p^\epsilon$, $\forall \hat{q} \in N_q^\epsilon$, $v(\hat{q}|\hat{p}) > v(\hat{q}|\hat{p})$. Since v is expected utility, it follows that for all such $\hat{p}, \hat{q}$ pairs and $\forall \alpha \in [0, 1)$, $v((1-\alpha)\hat{p} + \alpha\hat{q}|\hat{p}) > v(\hat{p}|\hat{p})$. It follows that for all such $\hat{p}, \hat{q}$ pairs and for any such $\alpha \in [0, 1)$, whenever $(1-\alpha)\hat{p} + \alpha\hat{q} \in \hat{D}$ it follows that $\hat{p} \notin PPE(\hat{D}) = c(\hat{D})$. Thus IIA Independence holds.

Necessity of Transitive Limit. First, I show that the antecedent of Transitive Limit has bite in the presence of, and only in the presence of, a strict preference. To be precise, suppose $(1-\epsilon)p^\delta + \epsilon q^\delta = c(\{(1-\epsilon)p^\delta + \epsilon q^\delta,\ (1-\epsilon)p^\delta + \epsilon r^\delta\})$ for all small ϵ, and $p^\delta, q^\delta, r^\delta$ sufficiently close to p, q, r. By the representation, this holds only if for all p^δ close to p, q^δ close to q, r^δ close to r, and ϵ close to zero, $v(q^\delta|(1-\epsilon)p^\delta + \epsilon q^\delta) \geq v(r^\delta|(1-\epsilon)p^\delta + \epsilon q^\delta)$, thus $v(q^\delta|p^\delta) \geq v(r^\delta|p^\delta)$ for all $p^\delta, q^\delta, r^\delta$. If $v(q|p) = v(r|p)$, then for every q^δ near q, $v(q^\delta|p) \geq v(q|p)$ and for every r^δ near r, $v(r|p) \geq v(r^\delta|p)$; this contradicts local strictness of $v(\cdot|p)$ in the representation. Thus when the antecedent of Transitive Limit holds, $v(q|p) > v(r|p)$ must hold.

Now take a continuous EU-PE representation and suppose $v(q|p) > v(r|p)$. Then, joint continuity implies that $v((1-\lambda)s + \lambda q^\delta|p^\delta) > v((1-\lambda)s + \lambda r^\delta|p^\delta)$ for any $s \in \Delta$, $\lambda > 0$, and δ close to zero. It follows that $v((1-\epsilon)p^\delta + \epsilon q^\delta|(1-\epsilon)p^\delta + \epsilon r^\delta) > v((1-\epsilon)p^\delta + \epsilon r^\delta|(1-\epsilon)p^\delta + \epsilon r^\delta)$ for all δ, ϵ sufficiently small. Thus for sufficiently small δ, ϵ, $(1-\epsilon)p^\delta + \epsilon q^\delta = c(\{(1-\epsilon)p^\delta + \epsilon q^\delta,\ (1-\epsilon)p^\delta + \epsilon r^\delta\})$. Thus the antecedent of Transitive Limit holds when $v(q|p) > v(r|p)$.

Since $v(q|p) > v(r|p)$ and $v(r|p) > v(s|p)$ imply $v(q|p) > v(s|p)$, the analysis above implies that qR_pr and rR_ps implies qR_ps, so Transitive Limit must hold.

Necessity of Induced Reference Lottery Bias. In the representation, $v(p|p) \geq v(q|p)$ and $v(q|q) \leq \max[v(p|p), v(p|q)]$ imply that $\forall \alpha \in (0, 1)$, $v((1-\alpha)p + \alpha q|(1-\alpha)p + \alpha q) \leq \max[v(p|p), v(p|(1-\alpha)p + \alpha q)]$.

Suppose $p \in c(D)$. Then, $v(p|p) \geq v(q|p)$ $\forall q \in D$, and $v(p|p) \geq v(q|q)$ $\forall q \in PE(D)$. It follows that $v(p|p) \geq v(q|p)$ and $v(q|q) \leq \max[v(p|p), v(p|q)]$. Since $v(\cdot|p)$ satisfies expected utility, $p \in PE((1-\alpha)p + \alpha D)$ $\forall \alpha \in (0, 1)$. Since $v((1-$

$\alpha)p + \alpha q|(1-\alpha)p + \alpha q) \leq \max[v(p|p), v(p|(1-\alpha)p + \alpha q)]\ \forall q \in D$, it follows that $v(p|p) \geq v((1-\alpha)p + \alpha q|(1-\alpha)p + \alpha q)\ \forall q : (1-\alpha)p + \alpha q \in PE((1-\alpha)p + \alpha D)$. Thus $p \in PPE((1-\alpha)p + \alpha D) = c((1-\alpha)p + \alpha D)\ \forall \alpha \in (0,1)$. Thus Induced Reference Lottery Bias holds.

□

Proof of Proposition 2.3.

Suppose that $v(\cdot|p)$ and $v(\cdot|q)$ are not ordinally equivalent. Then $\exists \bar{r}, \bar{s} \in \Delta$ such that $v(\bar{r}|p) > v(\bar{s}|p)$ but $v(\bar{r}|q) \leq v(\bar{s}|q)$. By local strictness, $\exists r, s \in \Delta$ that are close to $\bar{r}, \bar{s}$ such that $v(r|p) > v(s|p)$ but $v(r|q) < v(s|q)$. By EU of $v(\cdot|p)$ and continuity of v, this implies that $\exists \bar{\delta}, \bar{\epsilon} > 0$ such that $\forall \epsilon \in (0, \bar{\epsilon})$, $\forall r^\delta \in N_r^\delta$, $\forall s^\delta \in N_s^\delta$, $v((1-\epsilon)p + \epsilon r^\delta|(1-\epsilon)p + \epsilon s^\delta) > v((1-\epsilon)p + \epsilon s^\delta|(1-\epsilon)p + \epsilon s^\delta)$ but $v((1-\epsilon)q + \epsilon s^\delta|(1-\epsilon)q + \epsilon r^\delta) > v((1-\epsilon)q + \epsilon r^\delta|(1-\epsilon)q + \epsilon r^\delta)$. By the representation, this implies that for such $\epsilon, r^\delta, s^\delta$, (a) $(1-\epsilon)p + \epsilon r^\delta = c(\{(1-\epsilon)p + \epsilon r^\delta, (1-\epsilon)p + \epsilon s^\delta\})$ and (b) $(1-\epsilon)q + \epsilon s^\delta = c(\{(1-\epsilon)q + \epsilon r^\delta, (1-\epsilon)q + \epsilon s^\delta\})$. Thus if $v(\cdot|p)$ and $v(\cdot|q)$ are not ordinally equivalent, c strictly exhibits expectations-dependence.

Now suppose that c exhibits expectations-dependence at D, α, p, q, r. That is, $\exists \bar{\epsilon} > 0$ such that $\forall r^\epsilon \in N_r^\epsilon$, $\forall D^\epsilon \ni r^\epsilon$ such that $d^H(D^\epsilon, D) < \epsilon$, $(1-\alpha)p + \alpha r^\epsilon \in c((1-\alpha)p + \alpha D^\epsilon)$ but $(1-\alpha)q + \alpha r^\epsilon \notin c((1-\alpha)q + \alpha D^\epsilon)$. Since $(1-\alpha)q + \alpha r^\epsilon \notin c((1-\alpha)q + \alpha D^\epsilon)$, it follows that for each D^ϵ, $\exists \bar{s}^\epsilon \in D^\epsilon$, $v(\bar{s}^\epsilon|(1-\alpha)p + \alpha \bar{s}^\epsilon) \geq v(r^\epsilon|(1-\alpha)p + \alpha \bar{s}^\epsilon)$. Local strictness then implies that for each such $\bar{s}^\epsilon, r^\epsilon$ pair, there is an arbitrarily close pair $\hat{s}^\epsilon, \hat{r}^\epsilon$ such that $v(\hat{s}^\epsilon|(1-\alpha)p + \alpha \bar{s}^\epsilon) > v(\hat{r}^\epsilon|(1-\alpha)p + \alpha \bar{s}^\epsilon)$. By the representation, $(1-\alpha)p + \alpha r^\epsilon \in c((1-\alpha)p + \alpha D^\epsilon)$ implies that for each r^ϵ, $\forall s^\epsilon \in D^\epsilon$, $v(r^\epsilon|(1-\alpha)p + \alpha r^\epsilon) \geq v(s^\epsilon|(1-\alpha)p + \alpha r^\epsilon)$; thus $v(\hat{r}^\epsilon|(1-\alpha)p + \alpha \hat{r}^\epsilon) \geq v(\hat{s}^\epsilon|(1-\alpha)p + \alpha \hat{r}^\epsilon)$. Thus v exhibits strict expectations-dependence. This proves the first part of the proposition.

Now suppose c violates IIA. Then there are D, D' such that $D' \subset D$ and $c(D) \cap D' \neq \emptyset$ but $c(D') \neq c(D) \cap D'$. This implies that either (a) or (b) holds:

(a) $\exists p \in c(D')$ such that $p \notin c(D)$. Then by the representation, this implies that $v(p|p) = v(q|q)$ for $q \in c(D')$, so for some $r \in D$, $v(r|p) > v(p|p) \geq v(q|p)$ but $v(q|q) \geq v(r|q)$

(b) $\exists p \in c(D) \cap D'$ with $p \notin c(D')$. Since $PE(D) \cap D' \subset PE(D')$, this implies that there is a $q \in c(D')$ with $v(q|q) > v(p|p)$. Thus $q \notin c(D) \implies q \notin PE(D)$, which implies that $\exists r \in D \backslash D'$ such that $v(r|q) > v(q|q) \geq v(p|q)$ but $v(p|p) \geq v(r|p)$.

In either case (a) or (b), by the first part of the proposition, c exhibits strict expectations-dependence.

□

Proof of Proposition 2.5

First prove that Kőszegi-Rabin preferences with linear loss aversion satisfy the limited-cycle inequalities.

Start with a finite set X with $|X| = n+1$ and assume (for now) that there is a single hedonic dimension. Without loss of generality, assume $m(x_1) > m(x_2) > ... > m(x_{n+1})$

Define the matrix V according to:

$$[V]_{ij} = m(x_i) + \eta[m(x_i) - m(x_j)] + \eta[\lambda - 1]\min[0,\ m(x_i) - m(x_j)] \tag{A.2}$$

Observe that $v(p|r) = p^T V r$. Let $\delta, \epsilon \in \Re^{n+1}$ denote vectors with $\sum_{i=1}^{n+1} \delta_i = \sum_{i=1}^{n+1} \epsilon_i = 0$. By matrix multiplication,

$$\begin{gathered} \delta^T V \epsilon = \eta[\lambda - 1]\times \\ [(m(x_1) - m(x_2))\delta_1\epsilon_1 + (m(x_2) - m(x_3))(\delta_1 + \delta_2)(\epsilon_1 + \epsilon_2) + \\ ... + (m(x_n) - m(x_{n+1}))(\sum_{i=1}^{n} \delta_i)(\sum_{i=1}^{n} \epsilon_i)] \end{gathered} \tag{A.3}$$

Take a cycle $p^{i+1} = p^i + \epsilon^i$ with $v(p^{i+1}|p^i) > v(p^i|p^i)$ for $i = 0, ..., m$. Then:
$v(p^m|p^m) - v(p^0|p^m) = (p + \sum_{l=1}^m \epsilon^l)^T V (p + \sum_{l=1}^m \epsilon^l) - p^T V (p + \sum_{l=1}^m \epsilon^l)$
$= (\sum_{l=1}^m \epsilon^l)^T V (\sum_{l=1}^m \epsilon^l) + (\sum_{l=1}^m \epsilon^l)^T V p$
Rearranging the second term,
$= (\sum_{l=1}^m \epsilon^l)^T V (\sum_{l=1}^m \epsilon^l) + (\sum_{l=1}^{m-1} \epsilon^l)^T V p + (\epsilon^m)^T V (p + \sum_{l=1}^{m-1} \epsilon^l) - (\epsilon^m)^T V (\sum_{l=1}^{m-1} \epsilon^l)$

$= (\sum_{l=1}^{m} \epsilon^l)^T V (\sum_{l=1}^{m} \epsilon^l) + (\sum_{l=1}^{m-2} \epsilon^l)^T V p + (\epsilon^{m-1})^T V (p + \sum_{l=1}^{m-2} \epsilon^l) - (\epsilon^{m-1})^T V (\sum_{l=1}^{m-2} \epsilon^l) + (\epsilon^m)^T V (p + \sum_{l=1}^{m-1} \epsilon^l) - (\epsilon^m)^T V (\sum_{l=1}^{m-1} \epsilon^l)$

$= ... = (\sum_{l=1}^{m} \epsilon^l)^T V (\sum_{l=1}^{m} \epsilon^l) + \sum_i (\epsilon^i)^T V (p + \sum_{l=1}^{i-1} \epsilon^l) - \sum_{i=2}^{m} \epsilon^i V (\sum_{l=1}^{i-1} \epsilon^l)$

By the definition of the cycle, $(\epsilon^i)^T V (p + \sum_{l=1}^{i-1} \epsilon^l) > 0$ for each i, thus:

$> (\sum_{l=1}^{m} \epsilon^l)^T V (\sum_{l=1}^{m} \epsilon^l) - \sum_{i=2}^{m} \epsilon^i V (\sum_{l=1}^{i-1} \epsilon^l)$

By symmetry with respect to δ and ϵ in (A.3), it can be shown that $\sum_{i=2}^{m} \sum_{l=1}^{i-1} (\epsilon^i)^T V \epsilon^l = \sum_{j=1}^{m-1} \sum_{l=j+1}^{m} (\epsilon^j)^T V \epsilon^l$. Returning to the previous expression, more algebra establishes:

$= \sum_{l=1}^{m} (\epsilon^l)^T V \epsilon^l + \sum_{i=2}^{m} \sum_{l=1}^{i-1} (\epsilon^i)^T V \epsilon^l$

$= \frac{1}{2} \sum_{l=1}^{m} (\epsilon^l)^T V \epsilon^l + \frac{1}{2} (\sum_{l=1}^{m} \epsilon^l)^T V (\sum_{l=1}^{m} \epsilon^l)$

> 0

This completes the proof for the case with the case of one hedonic dimension.

To extend the argument to $K > 1$, break up a lottery p into marginals $\underset{k}{p}$ in each dimenion k, and define the matrix $\underset{k}{V}$ as the utility matrix corresponding to V in dimension k. we can write $v^{KR}(p|r) = \sum_k \underset{k}{p}^T \underset{k}{V} \underset{k}{r}$. Notice that all of the previously-proven properties of V apply to $\underset{k}{V}$; following through the previous steps yields the desired result.

Second prove that Kőszegi-Rabin preferences with linear loss aversion dislike mixtures.

Suppose $v(p|p) \geq v(q|p)$ and $v(q|q) \leq \max[v(p|p), v(p|q)]$.

Then,

$v((1-\alpha)p + \alpha q | (1-\alpha)p + \alpha q)$

$$= (1-\alpha)^2 v(p|p) + \alpha(1-\alpha) v(p|q) + \alpha(1-\alpha) v(q|p) + \alpha^2 v(q|q) \qquad \text{(A.4)}$$

by bilinearity of v under (2.3) and linear loss aversion.

If $v(p|p) \leq v(p|q)$, then two substitutions to (A.4) yield

$\leq (1-\alpha)^2 v(p|p) + \alpha(1-\alpha) v(p|q) + \alpha(1-\alpha) v(p|p) + \alpha^2 v(p|q)$

$= v(p|(1-\alpha)p + \alpha q)$ by bilinearity of v

$= \max[v(p|(1-\alpha)p + \alpha q), v(p|p)]$

If instead $v(p|q) \leq v(p|p)$, then two different substitutions to (A.4) yield

$\leq (1-\alpha)^2 v(p|p) + \alpha(1-\alpha) v(p|p) + \alpha(1-\alpha) v(p|p) + \alpha^2 v(p|p)$

$= v(p|p)$
$= \max\left[v(p|(1-\alpha)p + \alpha q),\, v(p|p)\right]$
This proves that v dislikes mixtures.
□

Proof of Proposition 2.6

Gul and Pesendorfer (2008) prove that on a finite set X there is an assignment of hedonic dimensions such that any reference-dependent utility function $\hat{v}(x|y)$ can be written as a Kőszegi-Rabin preference as in (2.3). Extend $\hat{v}(x|y)$ to lotteries by setting $v(p|q) = \sum_i \sum_j p_i q_j \hat{v}(x|y)$. The resulting representation over Δ is thus consistent with (2.3).

Kőszegi (2010, Example 3 and footnote 6) provides an example of $v : \Delta \times \Delta \to \Re$ in which the only personal equilibrium involves randomization among elements of a choice set. Mapping the v from Kőszegi's example to a Kőszegi-Rabin preference as described provides an example of a Kőszegi-Rabin preference that does not satisfy the limited-cycle inequalities.

□

Proof of Proposition 2.7

Take a continuous PPE representation corresponding to $\succeq_L, \{\succeq_p\}_{p\in\Delta}$. Take $p \in D$. Reference Lottery Bias implies that if $p \succeq_L q\ \forall q \in D$ then $p \succeq_p q\ \forall q \in D$; thus, $p \in m(D, \succeq_L) \implies p \in PE(D)$, which jointly imply $p \in PPE(D) = c(D)$. Since $\succeq_L$ is continuous and D is finite, it has a maximizer in D, thus there is a $p \in m(D, \succeq_L)$; by the previous argument, for any other $q \in c(D)$ it follows from the representation that $q \succeq_L p$ thus $q \in m(D, \succeq_L)$ as well. It follows that if $\succeq_L, \{\succeq_p\}_{p\in\Delta}$ satisfies Reference Lottery Bias, that $c(D) = m(D, \succeq_L)$.

□

Proof of Proposition 2.8.

(i) $\iff$ (iii)

Let suppose c is induced by the continuous binary relation P.

Necessity of Expansion. $p \in c(D) \iff \nexists q \in D$ such that qPp.

Thus, $p \in c(D)$ and $p \in c(D')$

$\iff$ both $\nexists q \in D$ such that qPp and $\nexists r \in D'$ such that rPp.

$\iff \nexists q \in D \cup D'$ such that qPp

$\iff p \in m(D \cup D', P)$

$\iff p \in c(D \cup D')$

Necessity of Sen's α. $p \in c(D) = m(D, P) \iff \nexists q \in D$ such that qPp

$\implies$ if $D' \subset D$, then $\nexists q \in D'$ such that qPp

$\iff p \in m(D', P) = c(D')$

Necessity of UHC. By contradiction.

Suppose $p^\epsilon \in c(D^\epsilon) = m(D^\epsilon, P)$ for a sequence $D^\epsilon \to D$ such that $d^H(D^\epsilon, D) < \epsilon$.

If $p \notin c(D)$, then $\exists q \in D$ such that qPp.

Then, since q has open better than and worse than sets, $\exists \epsilon$ such that $\forall p^\epsilon \in N_p^\epsilon$, $\forall q^\epsilon \in N_q^\epsilon$, $q^\epsilon P p^\epsilon$.

Since $d^H(D^\epsilon, D) < \epsilon$, it follows that $\forall D^\epsilon$ in the sequence, $\exists q^\epsilon \in D^\epsilon$ such that $d^E(q^\epsilon, q) < \epsilon$. Thus, $\exists \bar{\epsilon} > 0$ such that $\forall \epsilon < \bar{\epsilon}$, $q^\epsilon P p^\epsilon$. This contradicts that $p^\epsilon \in m(D^\epsilon, P)$ $\forall D^\epsilon$. $\lozenge$

Sufficiency. Construct $\bar{P}$ by:

$$p\bar{P}q \text{ if } \exists D_{pq} \text{ such that } p \in c(D_{pq})$$

Define P as the asymmetric part of $\bar{P}$.

(I) show $c(D) \subseteq m(D, P)$

If $p \in c(D)$, then $p \in m(D, P)$ by the definition of P.

(II) show $m(D, P) \subseteq c(D)$

Suppose $p \in m(D, P)$. Then, $\forall r \in D$, $\exists D_{pr} :\ p \in c(D_{pr})$.

By Expansion, $p \in c(\underset{r\in D}{\cup} D_{pr})$.

Since $D \subseteq \underset{r\in D}{\cup} D_{pr}$, by Sen's α, $p \in c(D)$ as well.

(III) show P is continuous.

If $p^\epsilon \bar{P} q^\epsilon$ for a sequence $p^\epsilon, q^\epsilon \to p, q$ then by steps (I) and (II), $p^\epsilon \in c(\{p^\epsilon, q^\epsilon\})$. By UHC, this implies $p \in c(\{p, q\})$ thus $p\bar{P}q$. Thus, $\bar{P}$ has closed better and worse than sets. Thus P has strictly open better and worse than sets.

□

Proof of Theorem 2.3.

Necessity. Necessity of Expansion, Sen's α, and UHC follows from Proposition 2.8.

Necessity of IIA Independence 2 and Transitive Limit are similar to Theorem 2.1.

To prove the necessity of Induced Reference Lottery Bias,

$p \in c(D) = PE(D)$

$\iff v(p|p) \geq v(q|p)\ \forall q \in D$

$\iff v(p|p) \geq v((1-\alpha)p + \alpha q|p)\ \forall q \in D$ since $v(\cdot|p)$ satisfies EU

$\iff p \in PE((1-\alpha)p + \alpha D) = c((1-\alpha)p + \alpha D)$

Thus the representation implies Induced Reference Lottery Bias.

Sufficiency.

Lemma A.20. *IIA Independence 2 implies Limit Consistency.*

Proof. Suppose qR_pp. Then $\exists \bar{\alpha} > 0$ such that $\forall \alpha \in (0, \bar{\alpha})$, $\{(1-\alpha) + \alpha q\} = c((1-\alpha)p + \alpha\{p,\ q\})$. By IIA Independence 2, it follows that $\forall \alpha \in (0, 1]$, $\forall D_{p,\,(1-\alpha)p+\alpha q}$ that $p \notin c(D_{p,\,(1-\alpha)p+\alpha q})$. Thus Limit Consistency holds. □

Take v from Lemma A.7 (from the proof of Theorem 2.1).

Define $PE(D) := \{p \in D :\ v(p|p) \geq v(q|p)\, \forall q \in D\}$.

By Lemma A.13, the axioms for Theorem 2.3 imply Limit Consistency. Since $v(\cdot|p)$ represents $\bar{R}_p$, Limit Consistency implies that $c(D) \subseteq PE(D)$.

Suppose $p \notin c(D)$ - I will show that $p \notin PE(D)$.

If $\forall q \in D$, $\exists D_{pq}$ such that $p \in c(D_{pq})$, then by Expansion, $p \in c(\underset{q\in D}{\cup} D_{pq})$; by Sen's α, it follows that $p \in c(D)$, a contradiction.

Thus $\exists q \in D$ such that $p \notin c(D_{pq})$ for any $D_{pq} \supseteq \{p, q\}$. It follows by IIA Independence 2 that $\exists \epsilon > 0$ such that $\forall \alpha \in (0, 1)$, $D_{p,\,(1-\alpha)p+\alpha q}$, and $\forall(\hat{p}, \hat{q}) \in N_p^\epsilon \times N_q^\epsilon$, $p \notin c(D_{\hat{p},\,(1-\alpha)\hat{p}+\alpha\hat{q}})$. This implies qR_pp. Thus $p \notin PE(D)$. It follows that $D\backslash c(D) \subseteq D\backslash PE(D)$, thus $PE(D) \subseteq D$.

This establishes that $PE(D) = c(D)$.

□

Remark. The proof of Theorem 2.3 makes no use of Induced Reference Lottery Bias. It follows that Induced Reference Lottery Bias is not independent of the remaining axioms.

Proof of Theorem 2.4.

Ok (2012, Chapters 5 and 9) proves that IIA and UHC hold if and only if c is induced by a continuous preference relation, if and only if c has a utility representation (since Δ is a separable metric space).[27]

For any continuous $u : \Delta \to \Re$, we can take any v that satisfies $v(p|p) = u(p)$; conversely, for any v we can define u by $u(p) := v(p|p)$. Under this mapping $CPE(D) = \max\limits_{p\in D} v(p|p) = \max\limits_{p\in D} u(p)$.

□

Proof of Proposition 2.9.

(i) $\iff$ (ii)

Assuming IRLB:

$p \in c(\{p, q\})$

$\iff p \succeq q$

$\implies p \in c((1-\alpha)p + \alpha\{p, q\})$ by IRLB

$\iff p \succeq (1-\alpha)p + \alpha q$

which proves that IRLB implies quasiconvexity of $\succeq$

[27] Arrow (1959) shows that IIA holds if and only if there exists a complete and transitive binary relation R such that c is induced by R.

Now assume quasiconvexity of $\succeq$:

$p \in c(D)$

$\iff p \succeq q \; \forall q \in D$

$\implies p \succeq (1-\alpha)p + \alpha q \; \forall q \in D$ by quasiconvexity

$\iff p \in c((1-\alpha)p + \alpha D)$.

(ii) $\iff$ (iii)

comparing the CPE and preference maximization representations, we see that:

$p \succeq q \iff v(p|p) \geq v(q|q)$.

Thus the statement "$p \succeq q \implies p \succeq (1-\alpha)p + \alpha q$" holds if and only if the statement "$v(p|p) \geq v(q|q) \implies v(p|p) \geq v((1-\alpha)p + \alpha q|(1-\alpha)p + \alpha q)$" holds.

$\square$

Appendix B

Proofs for Chapter 3

Proposition B.1. *The following are sufficient conditions for c to satisfy dual differentiability at w and weak dual differentiability at $\tilde{w}$: (i) c is RDU with u twice differentiable at w, (ii) c is semi-weighted utility with twice differentiable upper and lower weighting functions and u, (iii) c is Frechet differentiable.*

Proof. Case (i): $c(w+\hat{p}^t) = u^{-1}\left(\sum_{i=1}^{n}[g(\sum_{j=1}^{i}p_j) - g(\sum_{j=1}^{i-1}p_j)]u(w+ty_i)\right)$; define $\hat{c}(\hat{p}) = \sum_{i=1}^{n}[g(\sum_{j=1}^{i}p_j) - g(\sum_{j=1}^{i-1}p_j)]y_i$; if u is twice differentiable then $u(w+ty_i) = ty_i u'(w) + o(t)$ follows by Taylor's theorem. Similarly, $c(\tilde{w}+ty) = \left(\frac{1}{u'(c(\tilde{w}))}\int u'(w)dg(F_{\tilde{w}}(w))\right)ty + c(\tilde{w}) + o(t)$.

Case (ii): Suppose V has a semi-weighted utility functional form with upper and lower weighting functions $\underline{\omega}, \bar{\omega}$ and which need not coincide. For small t,

$V(w+\hat{p}^t) = \frac{\sum_{i:y_i\le 0} p_i\underline{\omega}(w+ty_i)u(w+ty_i) + \sum_{i:y_i>0} p_i\bar{\omega}(w+ty_i)u(w+ty_i)}{\sum_{i:y_i\le 0} p_i\underline{\omega}(w+ty_i) + \sum_{i:y_i>0} p_i\bar{\omega}(w+ty_i)}$

If $\bar{\omega}, \underline{\omega}, u$ are all twice differentiable around w, then so is $V(w+\hat{p}^t)$ as a function of t. This implies $c(w+\hat{p}^t) - w = \frac{1}{u'(w)}\frac{dV(w+\hat{p}^t)}{dt}|_{t=0^+} + o(t)$. Taking the derivative and reorganizing, $\frac{dV(w+\hat{p}^t)}{dt}|_{t=0^+} = \frac{\sum_{i:y_i\le 0} p_i\underline{\omega}(w)y_i + \sum_{i:y_i>0} p_i\bar{\omega}(w)y_i}{\sum_{i:y_i\le 0} p_i\underline{\omega}(w) + \sum_{i:y_i>0} p_i\bar{\omega}(w)}u'(w)$ and define $\hat{c}(\hat{p}) = \frac{\sum_{i:y_i\le 0} p_i\underline{\omega}(w)y_i + \sum_{i:y_i>0} p_i\bar{\omega}(w)y_i}{\sum_{i:y_i\le 0} p_i\underline{\omega}(w) + \sum_{i:y_i>0} p_i\bar{\omega}(w)}$.

Similarly,

$V(\tilde{w}+ty)$

$= \left(\int \omega_{\text{sign}(w-c(\tilde{w}))}(w+ty)dF_{\tilde{w}}(w)\right)^{-1}\left(\int \omega^{\text{sign}(w-c(\tilde{w}))}(w+ty)u(w+ty)dF_{\tilde{w}}(w)\right)$.

Define $\bar{c}_{y}\text{sign}_{(y)} \equiv \lim_{t\to 0^+}\frac{1}{ty}[V(\tilde{w}+ty) - V(\tilde{w})]$. Algebra yields:

$c_{y}\text{sign}_{(y)} = \left(\int \omega^{\text{sign}(w-c(\tilde{w}))}(w)dF_{\tilde{w}}(w)\right)^{-1}\left(\int \omega_{\text{sign}(w-c(\tilde{w}))}(w)u'(w)dF_{\tilde{w}}(w)\right)$

Note that if $\tilde{w}$ has positive mass at its certainty equivalent, then $\tilde{w}+ty$ may or may not have positive mass on its certainty equivalent. However, we only need to take a one-sided derivative and can choose whether to assign the upper or lower weight function at $w = c(\tilde{w})$ in the limit as appropriate, and it follows that $c(\tilde{w} + ty) = c(\tilde{w}) + tyc_{y\,\mathrm{sign}(y)}(\tilde{w}) + o(t)$. When $\tilde{w}$ has mass on its $c(\tilde{w})$, $c_{y\,\mathrm{sign}(y)}$ may indeed depend on $\mathrm{sign}(y)$ but otherwise will not.

(iii) To illustrate the link between dual differentiability and variations on Fréchet-differentiability considered in Wang (1993), suppose that V is differentiable with respect to the L^ρ norm for some $\rho \geq 1$. Notice that $\parallel [w,1] - [w+\hat{p}^t] \parallel_\rho = (\sum_i p_i^\rho |ty_i|)^{\frac{1}{\rho}} = t^{\frac{1}{\rho}} (\sum_i p_i^\rho |y_i|)^{\frac{1}{\rho}} = o(t^{\frac{1}{\rho}})$ while similarly $\parallel \tilde{w} - [\tilde{w} + ty] \parallel_\rho = t^{\frac{1}{\rho}}$. Applying Wang's local utility approximation obtained by taking the derivative using the L^ρ norm yields a $o(t^{\frac{1}{\rho}})$ approximation of preference, which proves the desired result. Safra and Segal (2002) additionally prove that L^1 differentiability implies expected value maximization. $\square$

Remark B.1. The non-expected utility literature frequently makes use of a notion of Gateaux-differentiability relative to mixtures of lotteries in probability space which allow for local expected utility approximations of preferences; for example, see Chew and Safra (1987). The notion of dual differentiability used in this paper is a special case of a notion of Gateaux-differentiability relative to (comonotonic) mixtures of lotteries in outcome space in the sense of Yaari (1987).

Proof of Theorem 3.2.

Suppose c is weakly dually differentiable at $\tilde{w}$. Then $c(\tilde{w}+ty) = c(\tilde{w}) + ty\bar{c}_{y\,\mathrm{sign}(y)}(\tilde{w}) + o(t)$.

Then:

$U(\hat{p}^t \oplus \tilde{w})$

$= c([c(\tilde{w} + ty_1), p_1; ...; c(\tilde{w} + ty_n), p_n])$

$= c([c(\tilde{w}) + \bar{c}_{y\,\mathrm{sign}(y_1)}(\tilde{w})ty_1 + o(t), p_1; ...; c(\tilde{w}) + \bar{c}_{y\,\mathrm{sign}(y_n)}(\tilde{w})ty_n + o(t), p_n])$

$= c(c(\tilde{w}) + (\bar{c}_{y\,\mathrm{sign}(y)}(\tilde{w})ty_1 + o(t), p_1; ...; \bar{c}_{y\,\mathrm{sign}(y)}(\tilde{w})ty_n + o(t), p_n))$

Since c is dually differentiable at $c(\tilde{w})$, there is a $\check{c}$ such that $c(c(\tilde{w}) + \hat{p}^t + tx)) = c(\tilde{w}) + t\check{c}(\hat{p}) + tx + o(t)$.

Define $\hat{c}(\hat{p}) = \check{c}((\bar{c}_{y\,\mathrm{sign}(y_1)}(\tilde{w})y_1, p_1; ...; \bar{c}_{y\,\mathrm{sign}(y_n)}(\tilde{w})y_n, p_n))$. Since $\check{c}$ is a linear in money, so is $\hat{c}$. By linearity in money, we can take the sup and inf of the $o(t)$

terms in $\hat{c}(ty_1 + o(t), p_1; ...; ty_n + o(t), p_n)$, and bound from above and below by $\hat{c}(ty_1, p_1; ...; ty_n, p_n) \pm o(t)$. Thus $U(\hat{p}^t \oplus \tilde{w}) = c(\tilde{w}) + \hat{c}(\hat{p}^t) + o(t) = c(\tilde{w}) + t\hat{c}(\hat{p} - a) + at + o(t)$ for any $a \in \Re$.

First-order risk-aversion in c directly implies first-order risk aversion in $\check{c}$. Take j so that $y_j < 0$ for $j < i$ and $y_j \geq 0$ for $j \geq i$. Form a new gamble from $\hat{p}^t$, call it $\check{p}^t$, that has the same outcomes and probabilities of each as $\hat{p}^t$ except that all gain outcomes are multiplied by $\frac{\bar{c}_{y^+}(\tilde{w})}{\bar{c}_{y^-}(\tilde{w})}$. Notice that $\hat{c}(\hat{p}^t) = \check{c}(\check{p}^{t\bar{c}_y(\tilde{w})})$. Since $\frac{\bar{c}_{y^+}(\tilde{w})}{\bar{c}_{y^-}(\tilde{w})} \leq 1$, $\check{c}(\check{p}^{t\bar{c}_y(\tilde{w})}) \leq \check{c}(\hat{p}^{t\bar{c}_y(\tilde{w})})$. Thus $\hat{c}(\hat{p}^t) - tx \leq \check{c}(\hat{p}^t) - tx = t\,[\check{c}(\hat{p}) - x]$. Thus first-order risk aversion in $\check{c}$ implies first-order risk aversion in $\hat{c}$.

□

Proof of Theorem 3.3.

Completing the proof of Theorem 3.3 requires showing that if DM is risk averse and V is RDU/DA, then $V(\tilde{w} + x)$ is concave in x and $\hat{U}(\hat{p}^t \oplus \tilde{w}) \equiv (u \circ U)(\hat{p}^t \oplus \tilde{w})$ is concave in t.

Under RDU, $V(\tilde{w} + x) = \int u(w + x) dg(F_{\tilde{w}}(w))$ where $F_{\tilde{w}}$ is the CDF associated with $\tilde{w}$. Since u is concave, $V(\tilde{w} + x)$ is concave in x.

Now under RDU, $\hat{U}(\hat{p}^t \oplus \tilde{w}) = \int V(\tilde{w} + ty) dg(F_{\hat{p}}(y))$. Since $V(\tilde{w}+x)$ is concave in x, $\hat{U}(\hat{p}^t \oplus \tilde{w})$ is concave in t. This implies that $\frac{1}{t}(\hat{U}(\hat{p}^t \oplus \tilde{w}) - V(\tilde{w})) \leq \hat{U}(\hat{p} \oplus \tilde{w}) - V(\tilde{w}) < 0$ for $t > 1$, so DM will also turn down $\hat{p}^t$.

Under DA, the same argument applies, except that it is messier to prove that $V(\tilde{w} + x)$ is concave in x and $U(\hat{p}^t \oplus \tilde{w})$ is concave in t - this is proven below.

Proof that $V(\tilde{w} + x)$ is concave in x under DA A DA DM is globally risk averse in the sense of weakly not preferring mean-preserving spreads if and only if u is concave and $\beta \geq 0$.

I will show that $(1 + \beta)[V(\tilde{w} + x) - V(\tilde{w})] \leq (1 + \beta)[V(\tilde{w}) - V(\tilde{w} - x)]$ to prove concavity.

We can write out the left- and right- hand sides of the above equation as:

$$(1+\beta)[V(\tilde{w}+x)-V(\tilde{w})]$$
$$= \int \{u(w+x)-u(w)+\beta\min[u(w+x),V(\tilde{w}+x)]-\beta\min[u(w),V(\tilde{w})]\}dF_{\tilde{w}}(w) \tag{B.1}$$

$$(1+\beta)[V(\tilde{w})-V(\tilde{w}-x)]$$
$$= \int \{u(w)-u(w-x)+\beta\min[u(w),V(\tilde{w})]-\beta\min[u(w-x),V(\tilde{w}-x)]\}dF_{\tilde{w}}(w) \tag{B.2}$$

To compare (B.1) and (B.2), compare the term inside the integral for each w. First note $u(w+x)-u(w) \leq u(w)-u(w-x)$. Second, compare the remaining parts of the integrals by working with four different regions/cases that depend on w and V.

Case A $\min[V(\tilde{w}+x),u(w+x)]-\min[V(\tilde{w}),u(w)] = u(w+x)-u(w)$ and $\min[V(\tilde{w}),u(w)]-\min[V(\tilde{w}-x),u(w-x)] = u(w)-u(w-x)$. By concavity of u for these w, the (B.1) term is smaller than the (B.2) term.

Case B $\min[V(\tilde{w}+x),u(w+x)]-\min[V(\tilde{w}),u(w)] = V(\tilde{w}+x)-V(\tilde{w})$ and $\min[V(\tilde{w}),u(w)]-\min[V(\tilde{w}-x),u(w-x)] = V(\tilde{w})-V(\tilde{w}-x)$. We can cancel these terms from $(1+\beta)[V(\tilde{w}+x)-V(\tilde{w})]$ and $(1+\beta)[V(\tilde{w})-V(\tilde{w}-x)]$.

Case C Suppose neither of the above two cases applies and $u(w) \leq V(\tilde{w})$. Then, applying concavity of u,

$$\begin{aligned}\min[u(w+x),V(\tilde{w}+x)]-\min[u(w),V(\tilde{w})]\} &= \min[u(w+x),V(\tilde{w}+x)]-u(w)\\ &\leq u(w+x)-u(w)\\ &\leq u(w)-u(w-x)\\ \leq u(w)-\min[u(w-x),V(\tilde{w}-x)] &= \min[u(w),V(\tilde{w})]-\min[u(w-x),V(\tilde{w}-x)]\}\end{aligned}$$

so in this case, this term of the (B.1) is smaller than the corresponding term in (B.2).

Case D Suppose neither of the above three cases applies, so $u(w) > V(\tilde{w})$. Then,

$$\begin{aligned}\min[u(w+x), V(\tilde{w}+x)] - \min[u(w), V(\tilde{w})]\} &= \min[u(w+x), V(\tilde{w}+x)] - V(\tilde{w}) \\ &\leq V(\tilde{w}+x) - V(\tilde{w})\end{aligned}$$

and

$$\begin{aligned}\min[u(w), V(\tilde{w})] - \min[u(w-x), V(\tilde{w}-x)]\} &= V(\tilde{w}) - \min[u(w-x), V(\tilde{w}-x)]\} \\ &\geq V(\tilde{w}) - V(\tilde{w}-x)\end{aligned}$$

Plugging in these terms and cancelling out as case B establishes the desired inequality.

Proof that $\hat{U}(\hat{p}^t \oplus \tilde{w})$ is concave in t under DA Define $I_V(y) = 1$ if $V(\tilde{w}+ty) - V(\tilde{w}) < \hat{U}(\hat{p}^t \oplus \tilde{w}) - V(\tilde{w})]$ and zero otherwise, and $I_U(y) = 1$ if $V(\tilde{w}+ty) - V(\tilde{w}) \geq \hat{U}(\hat{p}^t \oplus \tilde{w}) - V(\tilde{w})]$ and zero otherwise. For $t > 1$,

$\frac{1+\beta}{t}\{\hat{U}(\hat{p}^t \oplus \tilde{w}) - V(\tilde{w})\}$

$= \frac{1}{t}\int\{\hat{U}(\tilde{w}+ty) - V(\tilde{w}) + \beta \min[V(\tilde{w}+ty) - V(\tilde{w}), \hat{U}(\hat{p}^t \oplus \tilde{w}) - V(\tilde{w})]\}dF_{\hat{p}}(y)$

$\leq \frac{1}{t}\int\{V(\tilde{w}+ty) - V(\tilde{w}) + \beta I_V(y)[V(\tilde{w}+ty) - V(\tilde{w})] + \beta I_U(y)[\hat{U}(\hat{p}^t \oplus \tilde{w}) - V(\tilde{w})]\}dF_{\hat{p}}(y)$

Let $\bar{p} = 1 - \int I_U(y)dF_{\hat{p}}(y)$. Then, rearranging the above expression yields the inequality:

$\frac{1+\beta\bar{p}}{t}\{\hat{U}(\hat{p}^t \oplus \tilde{w}) - V(\tilde{w})\} \leq \frac{1}{t}\int\{V(\tilde{w}+ty) - V(\tilde{w}) + \beta I_V(y)[V(\tilde{w}+ty) - V(\tilde{w})]\}dF_{\hat{p}}(y)$

$\leq \int\{V(\tilde{w}+y) - V(\tilde{w}) + \beta I_V(y)[V(\tilde{w}+y) - V(\tilde{w})]\}dF_{\hat{p}}(y)$

$= \frac{1+\beta\bar{p}}{t}\{\hat{U}(\hat{p} \oplus \tilde{w}) - V(\tilde{w})\}$

□

Appendix C

Examples that rationalize our results

C.0.4 Rationalizing our data using ROCL

In the analysis below, we show by way of a simple example that if a subject views a list as a compound lottery, and satisfies reduction of compound lotteries, then even if she would exhibit a certainty effect when given binary choices, she will behave as if she satisfies expected utility when her preferences over lotteries are elicited using a probability list under the (mistaken) assumption of compound independence.[28] This example shows the potential of the Karni and Safra (1987) approach to rationalize our main findings.

For simplicity, approximate the 'discrete' price list with a smooth one. Suppose that we have smooth lists for Q1 and Q2 that ask for an indifference points $q \in [\frac{1}{2}, 1]$ and $r \in [\frac{1}{4}, \frac{1}{2}]$, and a number is drawn from $U[\frac{1}{2}, 1]$ for Q1 and from $U[\frac{1}{4}, \frac{1}{2}]$ for Q2 to determine which "line" is paid out.

Suppose the subject views Q1 and Q2 as compound lotteries in which the external randomizing device picks a line for payment at the first stage, and she gets her chosen lottery from that line at the second stage, and suppose further that the subject's preferences have a rank-dependent utility representation (Quiggin, 1982; Yaari, 1987). If the subject applies Compound Independence to evaluate this compound lottery, she would choose switch points q in Q1 and r in Q2 to satisfy:

$$u(3) = f(q)u(4) \text{ (Q1)}$$

$$f(.5)u(3) = f(r)u(4) \text{ (Q2)}$$

[28] The results below do not require the full force of reduction: if subjects' evaluation of a compound lottery was a weighted average of its value when reduced and when evaluated recursively, similar results would apply.

Suppose instead, following Karni and Safra (1987), the subject satisfies the Reduction of Compound Lotteries axiom. Then, she views her choice of q in Q1 as giving her the reduced lottery $(0,\ (1-q)^2;\ 3,\ 2q-1;\ 4,\ 1-q^2)$. Similarly, she evaluates her choice of r in Q2 based on the reduced lottery $(0,\ 1-2r+2r^2;\ 3,\ 2r-\frac{1}{2};\ 4,\ \frac{1}{2}-2r^2)$.

In the smooth approximation to the list, a subject would choose her switch points to solve:

$$\max_{q\in[\frac{1}{2},1]} f(1-q^2)\,[u(4)-u(3)] + f(2q-q^2)u(3)\ \text{(Q1)}$$

$$\max_{r\in[\frac{1}{4},\frac{1}{2}]} f(\tfrac{1}{2}-2r^2)\,[u(4)-u(3)] + f(2r-2r^2)u(3)\ \text{(Q2)}$$

First, we can see that under EU ($f(p)=p$) that subjects will switch at the same line ($q=2r$) in Q1 and Q2. If f is non-linear, then the above analysis suggests that subjects may switch at different points, depending on the shape of their probability weighting function. The neo-additive probability weighting function (Chateauneuf, Eichberger, and Grant, 2007; Webb and Zank, 2011) is closest to the spirit of the certainty and possibility effects motivating probability weighting, and accomodates them by a piece-wise linear f.

For $a\in(0,1)$, and $a+b<1$ this probability weighting function demonstrates a certainty effect, and when $b>0$ generates a possibility effect. Thus, for normal parameter values for a and b, f will generate a standard common ratio effect. This form for f can alternatively be motivated as a piece-wise linear approximation to more popular, smooth, probability weighting functions.

With a neoadditive f, when responding to the list elicitation, the FOCs for Q1 and Q2 reduce to:

$$q\,[u(4)-u(3)] = [1-q]\,u(3)\ \text{(Q1)}$$

$$2r\,[u(4)-u(3)] = [1-2r]\,u(3)\ \text{(Q2)}$$

so $q=2r$, and subjects behave as if they were expected utility maximizers. This is because some risks in both lists are going to be too attractive for subjects to pass up under reasonable parameter values. Thus, subjects will always face some risk in both lists. A certainty effect would bias subjects towards the certain option if any of the lines of Q1 were presented independently as a single binary choice question. But, when facing the list subjects will always choose some risk, so the certainty effect will not bias subjects towards A in Q1.

When f is a power function, then when responding to the list elicitation, the FOCs for Q1 and Q2 reduce to:

$$q(1-q^2)^{\beta-1}\left[u(4)-u(3)\right] = [1-q]\,(2q-q^2)^{\beta-1}u(3) \text{ (Q1)}$$
$$2r(\tfrac{1}{2}-2r^2)^{\beta-1}\left[u(4)-u(3)\right] = [1-2r]\,(2r-2r^2)^{\beta-1}u(3) \text{ (Q2)}$$

which implies that $q = 2r$, that is, the subject's behaviour in the list elicitation experiment is indistinguishable from expected utility maximization. Numerical simulations or algebra can be used to verify that for reasonable parameters with $\beta > 1$, a subject who switches at q would rank (\$3, 1) $\succ$ (\$4, q) in a binary choice task.

Now suppose that instead of satisfying RDU, U falls in the class of NCI-satisfying preferences studied by Cerreia-Vioglio, Dillenberger, and Ortoleva (2013). Then, their NCI axiom directly requires that if (\$3, 1) $\prec$ (\$4, .8) that the subject will rank(\$3, 1) $\prec$ (\$4, .8) in line 11 of the list. This follows from the fact that a subject who decides between switching on either line 11 of line 12 faces the compound lottery in (4.1). Applying reduction as in (4.2), we this choice can be express as a choice between $\frac{25}{26}\left[\$4,\ \frac{9.1}{25};\ \$3,\ \frac{15}{25}\right] + \frac{1}{26}\left[\$3,\ 1\right]$ and $\frac{25}{26}\left[\$4,\ \frac{9.1}{25};\ \$3,\ \frac{15}{25}\right] + \frac{1}{26}\left[\$4,\ .8\right]$, and the subject must prefer the latter lottery. Thus NCI predicts more risk aversion under list elicitation. However, NCI does not have any direct implications for Q2.

One limitation of the above analysis is that even when considering reduction of compound lotteries within a list, we assume that (consistent with our results) compound independence holds at the previous stage of the compound lottery at which Q1 or Q2 is selected to determine payment. One possible explanation is that subjects violate both compound independence and reduction, but the violation of compound independence is much less severe, and the violation of reduction much more severe, when different lotteries that form branches of a compound lottery are never displayed on the same visual screen in a way that would facilitate reduction.[29] Another limitation of the above explanation is that rank-dependent utility with the neo-additive

[29] If instead subjects reduced the compound lottery formed by the entire experiment, then if they correctly anticipated the second experimental question subject responses in the two-list treatments (L and S) would be consistent with EU regardless of their weighting function.

weighting function can only explain the observed reversal between choices in binary choice and list elicitation treatments in Q1, not the comparable (but statistically weaker) finding in Q2.

C.0.5 Non-standard application of Compound Independence

In the analysis below, we show by way of a simple example using RDU with the power weighting function that if a subject views a list according to a non-standard compound lottery as in 4.3, and and evaluates this compound lottery recursively, then the subject's behaviour under list elicitation. This example shows the potential of the Segal (1988) approach to rationalize our main findings.

Applying RDU with the power weighting function to (4.3) yields the conditions:

$$\max_{i\in\{1,2,\ldots,26\}}\left[\left(\tfrac{i}{26}\right)^{\beta}\left(\sum_{j=1}^{i}\left[\left(\tfrac{j}{i}\right)^{\beta}-\left(\tfrac{j-1}{i}\right)^{\beta}\right](1.02-.02j)^{\beta}\right)u(4)+\left[1-\left(\tfrac{i}{26}\right)^{\beta}\right]u(3)\right] \tag{Q1}$$

$$\max_{i\in\{1,2,\ldots,26\}}\left[\left(\tfrac{i}{26}\right)^{\beta}\left(\sum_{j=1}^{i}\left[\left(\tfrac{j}{i}\right)^{\beta}-\left(\tfrac{j-1}{i}\right)^{\beta}\right](.51-.01j)^{\beta}\right)u(4)+\left[1-\left(\tfrac{i}{26}\right)^{\beta}\right](.5)^{\beta}\,u(3)\right] \tag{Q2}$$

With some algebra, we can see that the maximand in (Q2) is just $.5^{\beta}$ times the maximand in (Q1). It follows that the subject would switch on the same line in both questions. That is, her behaviour would be indistinguishable from expected utility maximization in spite of her non-expected utility preferences.

Using a smooth approximation to the compound lottery formed by list along, it is possible to solve out continuous analogues of (Q1) and (Q2). With the power weighting function, the subject would switch at a $q^{\star}$ that satisfies $q^{\star}=1-\left[1-\frac{u(3)}{u(4)}\right]^{\frac{1}{\beta}}<\frac{u(3)}{u(4)}$ for $\beta>1$; in contrast, the same subject would require $\hat{q}=\left(\frac{u(3)}{u(4)}\right)^{\frac{1}{\beta}}>\frac{u(3)}{u(4)}$ to be indifferent between (\$3, 1) and (\$4, q). That is, the subject would display more risk aversion in the binary choice task than under list elicitation.