Spring 5-17-2010

Essays on the Information Acquisition of Doubt-Prone Decision Makers

larbi alaoui
University of Pennsylvania, larbi@sas.upenn.edu

Essays on the Information Acquisition of Doubt-Prone Decision Makers

Abstract
There are many situations in which individuals have a choice of whether or not to observe the eventual outcome. In these instances, individuals often prefer to avoid observing the outcome. The standard von Neumann-Morgenstern (vNM) Expected Utility model cannot accommodate these cases, since it does not distinguish between lotteries for which outcomes are observed by the agent and lotteries for which they are not. I develop an axiomatic model that admits preferences for observing the outcome or remaining in doubt. I then use this model to analyze the connection between the agent's attitude towards risk, doubt, and what I refer to as 'optimism'. This framework accommodates a wide array of field and experimental observations that violate the vNM model, and that may not seem related, prima facie. For instance, this framework accommodates self-handicapping, in which an agent chooses to impair his own performance. Unlike other frameworks, this model accommodates self-handicapping without using notions of self-deception, cognitive dissonance and belief manipulation. It also admits a status quo bias without having recourse to framing effects or reference points. Furthermore, this framework accommodates behavior associated with anticipated regret, the Allais paradox and preferences for smaller menus, which are all difficult to reconcile with the vNM framework. In financial settings, this model accommodates a safe allocation bias, in which agents choose neither to buy nor short sell an asset for an interval of prices; this behavior has so far been explained using ambiguity aversion, which this model does not allow. Recently, experiments have been conducted in which dictators in dictator games who seem to exhibit preferences for fairness often switch to the selfish choice if they can avoid observing the recipients allocation. While the empirical findings of these experiments are difficult to reconcile either with models of Expected Utility or models of fairness, they fit the predictions of this model well. This framework accommodates all the well-known observations mentioned here and others described in the papers with a single, natural extension of the standard vNM model, and using the same assumption on preferences throughout.

Degree Type
Dissertation

Degree Name
Doctor of Philosophy (PhD)

Graduate Group
Economics

First Advisor
Alvaro Sandroni

Second Advisor
Andrew Postlewaite

Third Advisor
David Dillenberger

Keywords
Value of information, uncertainty, doubt, unobserved outcomes, unresolved lotteries, regret

Subject Categories
Behavioral Economics | Economic Theory | Political Economy

ESSAYS ON THE INFORMATION ACQUISITION OF DOUBT-PRONE DECISION MAKERS

Larbi Alaoui

A DISSERTATION

in

Economics

Presented to the Faculties of the University of Pennsylvania in Partial Fulfillment of the Requirements for the Degree of Doctor of Philosophy

2010

Alvaro Sandroni, Professor of Economics
Supervisor of Dissertation

Dirk Krueger, Professor of Economics
Graduate Group Chairperson

Dissertation Committee

Alvaro Sandroni, Professor of Economics

Andrew Postlewaite, Professor of Economics

David Dillenberger, Assistant Professor of Economics

Acknowledgements

I will begin by thanking my main advisor, Alvaro Sandroni. He has endorsed this project from the very beginning, long before I could express my ideas in a coherent manner. He has patiently mentored me through the last few years without ever showing a trace of frustration, and has served as an example of the kind of economist I aspire to become. Above all, he matches the ideals I grew up having of what a thinker ought to be, and I recall him saying "I don't see why I would ever retire", which I think captures perfectly his joy of doing research. I also would like to thank Andrew Postlewaite, who has the rare ability to instantly shed away all layers of a model other than its essence, and to then rebuild a much more sound and interesting construction. I'm confident that I speak for many people when I say that having a conversation with him often feels like thinking in letters while he thinks in words. David Dillenberger, who is the youngest member of my committee and whose field of research is closest to mine, is simply a delight to talk to and I'm very fortunate that he joined the Penn faculty when he did.

I of course thank my parents, always supportive, who are unique in that they never questioned the idea of wanting to do research. Quite the contrary, they were more concerned when I pivoted from the 'sciences' to working in the real world, before pivoting from the real world to the 'humanities'. My mom's relief was palpable when she looked at these pages and saw equations that could not possibly make any sense, just as she had secretly hoped. Her own intellectual curiosity, unbounded culture and love for imparting knowledge have always been a source of inspiration for me. As for my dad, who remains the most important professor to me, and whose unconditional trust in my choices I am grateful for, I often found myself wondering whether his days as a graduate student were similar to mine, and it is gratifying that we share this bond. I'd like to thank my brother Hachem, who truly understood and felt the pain of the Wimbledon 2008 finals and to whom I'm close in many other ways, my sister Abla, who I see as my kindred spirit with a contrasting personality, my sister Mhani, with whom every discussion is thought-provoking, and my numerous family members, with a special thought for Amine, La Mama, Hmed and Meriem, to whom this also belongs.

Finally, I'd like to thank my close friends here, who have made all the difference. I could not possibly count the number of economics conversations and accompanying beers I have had with Antonio, or the number of shadows of ideas that have ensued. I'm grateful to Gil, one of the most loyal and sincere people I know, to Deniz, for always being world-thirsty and engaging, and for all his reassuring advice, and to Kirsten, who often served as my non-economist anchor.

ABSTRACT

ESSAYS ON THE INFORMATION ACQUISITION OF DOUBT-PRONE DECISION MAKERS

Larbi Alaoui

Alvaro Sandroni

There are many situations in which individuals have a choice of whether or not to observe the eventual outcome. In these instances, individuals often prefer to avoid observing the outcome. The standard von Neumann-Morgenstern (vNM) Expected Utility model cannot accommodate these cases, since it does not distinguish between lotteries for which outcomes are observed by the agent and lotteries for which they are not. I develop an axiomatic model that admits preferences for observing the outcome or remaining in doubt. I then use this model to analyze the connection between the agent's attitude towards risk, doubt, and what I refer to as 'optimism'. This framework accommodates a wide array of field and experimental observations that violate the vNM model, and that may not seem related, prima facie. For instance, this framework accommodates self-handicapping, in which an agent chooses to impair his own performance. Unlike other frameworks, this model accommodates self-handicapping without using notions of self-deception, cognitive dissonance and belief manipulation. It also admits a status quo bias without having recourse to framing effects or reference points. Furthermore, this framework accommodates behavior associated with anticipated regret, the Allais paradox and preferences for smaller menus, which are all difficult to reconcile with the vNM framework. In financial settings, this model accommodates a safe allocation bias, in which agents choose neither to buy nor short sell an asset for an interval of prices; this behavior has so far been explained using ambiguity aversion, which this model does not allow. Recently, experiments have been conducted in which dictators in dictator games who seem to exhibit preferences for fairness often switch to the selfish choice if they can avoid observing the recipients allocation. While the empirical findings of

these experiments are difficult to reconcile either with models of Expected Utility or models of fairness, they fit the predictions of this model well. This framework accommodates all the well-known observations mentioned here and others described in the papers with a single, natural extension of the standard vNM model, and using the same assumption on preferences throughout.

Contents

List of Tables

List of Figures

Preface

This dissertation consists of two related chapters that are designed to stand alone. In theory, they could be read independently, in either order. The first chapter, titled 'The value of useless information', has a more theoretical emphasis and delves deeper into the origin of the decision maker's preferences, although it does consider applications. The second chapter, titled 'Information aversion and the preservation of self-image', assumes the same type of individual preferences, and focuses on a specific set of applications, namely those in which the agent's self-image plays an important role in his choices. There is a theoretical component to this second chapter as well, even though it is mainly applications based.

This project began with my interest in a set of economically relevant situations that had not been formally analyzed. In particular, there are a number of situations in which individuals care about outcomes that they might *never* actually observe. This notion of having preferences over something that might *never* be observed is not one that is usually discussed in standard economics. In fact economic models, including the standard von Neumann-Morgenstern (vNM) Expected Utility model, usually do not distinguish between preferences when outcomes are eventually observed and when they are not. It did not appear that there were tools for analyzing this type of scenario. This does not mean that economic models do not *contain* setups in which agents care about outcomes they might never see; for instance, a very common Overlapping Generations Model with altruism effectively assumes that agents care about their descendants, whose consumption occurs after their death.

Rather, my motivation was that agents are restricted in the type of preferences they are allowed to express in these cases. That is, the standard vNM model (for example) does not distinguish between lotteries for which the outcomes are observed by the agent and lotteries for which they are not. Hence, in the case where agents have a choice between observing and *never* observing an outcome, they must be indifferent. Here, I take a formal axiomatic approach that builds on the vNM model, with the intention of characterizing more general preferences that admit preferences for observing an outcome or for remaining in doubt.

The original title of the project, now the title of the first chapter, 'The value of useless information', comes from the type of leading examples I initially considered. My leading examples were that patrons of local restaurant often do not wish to enter the kitchen, or that a significant percentage of people do not want to know whether the clothes they buy have been made by children. For those who are unwilling to change their behavior conditional on the information they receive, it may seem that this information is non-instrumental, and in that sense useless. But this information could in fact influence some individuals' behavior, and in most of the cases that I consider, choices can change radically depending on the information people receive. This makes the individual choice to *avoid* ever observing an outcome even more puzzling, as this information is instrumental.

After analyzing this model (henceforth VUI, for 'Value of Useless Information') in more depth, it became clear that it could accommodate many field and experimental observations that appear to violate the standard vNM model, and that might not appear related. Distinct models have been developed to accommodate some of these different empirical results that seem inconsistent with the standard vNM framework, but these models differ significantly from the vNM framework and from each other. Here, I show that a simple and natural extension of the vNM model can accommodate many behavioral patterns in what I think is a plausible way, using axioms that are very close to the standard vNM axioms. In addition, I show that these patterns are consistent with doubt-proneness. What I mean by this term is that agents exhibit a preference for not observing the resolution of uncertainty, and in that sense, they have a preference for remaining in doubt. To give a sense of the type of the seemingly disparate behavioral patterns that this model can accommodate, I list a few examples below, namely self-deception, the status quo bias, and anticipated regret:

(i) *Self-deception.* Economic models of self-image typically assume a technology for belief manipulation or temporal inconsistency. The notion of self-deception requires the agent to have a paradoxical ability to lie to himself. However, I show that an agent who is not temporally inconsistent and does not have manipulable beliefs may still act as if he were deceiving himself. One instance is self-handicapping, in which an individual deliberately reduces his chances of succeeding at a task. This result is consistent with doubt-proneness; an agent makes less effort to obtain a coarser signal of his decision making ability. This reasoning is essentially a formalization of the colloquial 'fear of failure'.

(ii) *Status quo bias.* The status quo bias refers to individuals' tendency to prefer their current endowment to other alternatives. This bias cannot be explained using the vNM model, and is one of the reasons cited by Kahneman, Knetch and Thaler (1991) for suggesting "a revised version of preference theory that would assign a special role to the status quo." I show that the VUI framework also admits a status quo bias, without having recourse to a notion of reference points, gains or losses. This bias occurs when inaction is a less informative indicator of the agent's ability than other actions. In some settings, however, the VUI framework predicts a bias towards the safe allocation, rather than a bias towards the status quo allocation. For example, I consider a financial market setting, and show that an agent will be biased towards the safe 'zero position', as implied by models of ambiguity aversion.

(iii) *Anticipated regret.* In Zeelenberg's (1999) Dutch lottery example, subjects in two groups have the choice between buying a lottery ticket and keeping their money. Agents in the 'feedback' group observe the result of the lottery ticket even if they do not buy it, and agents in the 'no-feedback' group only observe the result if they buy the ticket. Studies show that a higher percentage of individuals in the feedback group purchase the lottery than in the no-feedback group. This evidence is difficult to reconcile with utility theory, and is generally associated with anticipated regret. I show that this behavior is consistent with the VUI model.

The VUI model also accommodates other results that have motivated models of anticipated regret, disappointment aversion and prospect theory. A well-known instance is the variant of the Allais paradox known as the common ratio effect.[1] The same individuals who prefer \$200 with probability 1 to \$300 with probability 4/5 also often prefer \$300 with probability 1/2 to \$200 with probability 3/5. While this empirical finding is a clear violation of the standard vNM model, it is consistent with the VUI framework.

In addition to these examples, the VUI model also applies to other economic environments. In a political economy setting, voters deliberately remain ignorant, and as the importance of the relevant issue increases, their incentive to acquire information decreases. This model therefore serves as a unifying axiomatic foundation for a wide array of observed patterns of behavior. Of course, my aim is not to have this model replace all other explanations. There is, for example, much evidence and little debate that people do lie to themselves, and have myriad ways of holding contradictory ideas at once. There are models that capture this behavior well. My intention, rather, is to demonstrate that a wide range of behavior that may have seemed inconsistent with the standard vNM model can be accommodated with a simple extension of the standard vNM Expected Utility framework.

[1]This example is introduced in Kahneman and Tversky (1979). The numbers I use are from Gul (1998).

Chapter 1

The value of useless information

Abstract

There are many situations in which individuals have a choice of whether or not to observe the eventual outcome. In these instances, individuals often prefer to *avoid* observing the outcome. The standard von Neumann-Morgenstern (vNM) Expected Utility model cannot accommodate these cases, since it does not distinguish between lotteries for which outcomes are observed by the agent and lotteries for which they are not. I develop an axiomatic model that admits preferences for observing the outcome or remaining in doubt. I then use this model to analyze the connection between the agent's attitude towards risk, doubt, and what I refer to as 'optimism'.

This framework accommodates a wide array of field and experimental observations that violate the vNM model, and that may not seem related, prima facie. For instance, this framework accommodates self-handicapping, in which an agent chooses to impair his own performance. It also admits a status quo bias, without having recourse to framing effects. In a political economy setting, a voter avoids free information if he believes other voters will do the same.

1.1 Introduction

Models of decision making under uncertainty usually assume that the agents expect to eventually observe the resolution of uncertainty. However, there are many situations in which individuals can choose to avoid finding out which outcome has occurred. In these cases, individuals often decide *not* to observe the resolution of uncertainty. Consider the classic example of genetic diseases. As Pinker (2007) discusses, "the children of parents with Huntington's disease [HD] usually refuse to take the test that would tell them whether they carry the gene for it." HD is a neurodegenerative disease with severe physical and cognitive symptoms. It reduces life expectancy significantly, and there is currently no known cure. A person can take a predictive test to determine whether he himself will develop HD. A prenatal test can also be done to determine whether his unborn child will have the disease as well.[1] In an experimental study, Adam et al. (1993) find low demand for prenatal testing for HD. This is supported by a number of other studies as well, and Simpson et al. (2002) find that the demand for prenatal testing is significantly lower than the demand for predictive tests. That is, individuals who are willing to know their own HD status are often unwilling to find out their unborn child's status. Observing the result is an important decision, since the prenatal test is done at a stage in which parents can still terminate the pregnancy. As for parents who do not consider pregnancy termination to be an option, the information could still impact the way they decide to raise their child. For example, if they know that their child will develop HD, they might choose to prepare him psychologically for the difficult choices he will have to make in the future.

It may seem puzzling that some parents prefer to avoid the test. It may appear particularly surprising that a person who prefers to be certain of his own HD status now rather than later would also choose not to find out whether his unborn child will

[1] An affected individual has a 50% chance of passing the disease to each child. The average age of onsets varies between ages 35 and 55. See Tyler et al. (1990) for details.

develop the disease.[2] But note that the average age of onset for HD is high enough that the subjects who do not see the result of the prenatal test may *never* find out whether their children are affected. That is, while choosing the predictive test mostly reveals a preference for early resolution of uncertainty, choosing (or refusing) the prenatal test mainly reveals a preference for never observing the outcome of a lottery. It is precisely this type of preference on which this paper focuses.[3]

The standard von Neumann-Morgenstern (vNM) Expected Utility model cannot accommodate preferences for remaining in doubt, since it does not make a distinction between lotteries for which the final outcomes are observed and lotteries for which they are not.[4] Redefining the outcome space to include whether the prize is observed does not resolve the issue.[5] In this paper, I modify the basic axioms of the vNM framework to develop a model that admits strict preferences for remaining in doubt or for observing the outcome. This model is a natural extension of the vNM framework, but it can accommodate a wide array of field and experimental observations that are considered incompatible with the vNM model, including self-handicapping and the status quo bias.

1.1.1 Framework

An agent has primitive preferences over general lotteries that lead either to outcomes that he observes or to lotteries that never resolve, from his frame of reference.[6] This

[2]The prenatal test is not costless, as the procedure does involve a small chance of miscarriage. However, this cost appears small compared to the severity of the disease.

[3]In particular, this paper does not consider other factors that are present in the HD example, such as parents' concern that their child will be treated differently if it is known that he has HD, as discussed in Simpson (2002).

[4]The term observation is defined as learning what the outcome is. This model does not take into account a possible disutility from the graphical nature of the observation itself.

[5]See appendix for a discussion on the problem with redefining the outcome space to include the observation.

[6]Throughout this paper, probabilities are taken to be objective. With subjective probabilities, there are cases in which it may seem more natural to interpret the preferences as state-dependent. For instance, if a person has an intrinsic preference over his ability but is unsure of his type, it is unclear whether ability is better viewed as a state of the world or a consequence.

is a richer domain of lotteries than in the standard vNM case. If the agent receives a lottery that never resolves then he knows that he will not observe the outcome, and his terminal prize is the lottery itself. I apply the three standard vNM axioms on this expanded domain; that is, weak order, continuity and independence hold. I also assume that the agent is indifferent between observing a specific outcome and receiving an unresolved lottery that places probability one on that same outcome, since he is certain of the outcome's occurrence. The observation itself has no effect on the value of the outcome in this model. This property restricts the agent's allowable preferences over unresolved lotteries, as I demonstrate in section 2.

I obtain a representation theorem that separates the agent's risk-attitude over lotteries whose outcomes he observes from his risk attitude over unresolved lotteries. While this representation theorem suffices for most of the analysis, I also consider a second representation in a two-period setting in which the agent may learn 'early' or 'late' whether or not a lottery will resolve. His preferences over unresolved lotteries are allowed to change over time. In contrast, his preferences over lotteries that resolve do *not* change over time, as this model does not aim to capture a notion of anxiety.

Using the first (static) representation, I explore the connection between risk-aversion, doubt-proneness (a preference to avoid observing the outcome), and a new notion of optimism over unresolved lotteries, which I formally define. Intuitively, an optimistic agent prefers more 'scrambled' information. I show that an agent who is both doubt-prone and risk-averse over the unresolved lotteries can be neither optimistic nor pessimistic . In addition, his utility function associated with unresolved lotteries must be more concave than his utility function associated with lotteries whose outcome he observes. If an agent exhibits optimism over unresolved lotteries has the same utility function for both lotteries that resolve and lotteries that do not, then he must be doubt-prone.

Restricting attention only to preferences over purely unresolved lotteries, this model does not assume that these preferences obey the independence axiom. Instead, I assume the Rank-Dependent Utility (RDU) axioms, for reasons discussed in section 3. As there exists an accepted notion of optimism (Quiggin (1982)) in an RDU setting, it is of interest to formally relate RDU optimism to this paper's definition of optimism. RDU optimism essentially corresponds to a notion of overweighing the probabilities over the better outcomes. I show that my definition of optimism is equivalent to RDU optimism, if it holds everywhere. In that sense, it serves as a new axiom for RDU optimism.

1.1.2 Applications

The model presented here can accommodate seemingly unrelated behavioral patterns that are inconsistent with the standard vNM model, and that have motivated frameworks that are significantly different. Two important examples are self-handicapping and the status quo bias. In this analysis, I assume throughout that the agent is doubt-prone, but I do not allow him to be optimistic (or pessimistic) in his beliefs.

Consider first self-handicapping, in which individuals choose to reduce their chances of succeeding at a task. As discussed in Benabou and Tirole (2002), people may "choose to remain ignorant about their own abilities, and [...] they sometimes deliberately impair their own performance or choose overambitious tasks in which they are sure to fail (self-handicapping)." This behavior has been studied extensively, and seems difficult to reconcile with the standard Expected Utility theory.[7] For that reason, models that study self-handicapping make a substantial departure from the standard vNM assumptions. A number of models follow Akerlof and Dickens' (1982) approach of endowing agents with manipulable beliefs or selective memory. Alternatively, Carillo and Mariotti (2000) consider a model of temporal-inconsistency, in which a game is played between the selves, and Benabou and Tirole (2002) use both

[7] Berglass and Jones (1978) conduct an experiment in which they find that males take performance-inhibiting drugs, and argue that they do so precisely because it interferes with their performance.

manipulable beliefs and time-inconsistent agents.[8]

The frameworks mentioned above capture a notion of self-deception, which involves either a hard-wired form of selective memory (or perhaps a rule of thumb), or some form of conflict between distinct selves. These models are often not axiomatized. In contrast, this model simply extends the vNM framework and does not allow agents to manipulate their beliefs or to have access to any other means for deceiving themselves.[9] Yet it still accommodates the decision to self-handicap, as is shown in section 4. Intuitively, a doubt-prone agent prefers doing worse in a task if this allows him to avoid information concerning his own ability. This is essentially a formalization of the colloquial 'fear of failure'; an agent exerts less effort so as to obtain a coarser signal.

This model can also accommodate a status quo bias. The status quo bias refers to the well-known tendency people have for preferring their current endowment to other alternatives. This phenomenon is often seen as a behavioral anomaly that cannot be explained using the vNM model. On the other hand, it can be accommodated using loss aversion, which refers to the agent being more averse to avoiding a loss than to making a gain (Kahneman, Knetch and Thaler (1991)). The status quo bias is therefore an immediate consequence of the agent taking the status quo to be the reference point for gains versus losses. The vNM model does not allow an agent to evaluate a bundle differently based on whether it is a gain or a loss, and hence cannot accommodate a status quo bias. Arguably, this is an important systematic

[8]See also Compte and Postlewaite (2004), who focus on the positive welfare implications of having a degree of selective memory (assuming such technology exists) in the case where performance depends on emotions. Benabou (2008) and Benabou and Tirole (2006a, 2007) explore further implications of belief manipulation, particularly in political economy settings, in which multiple equilibria emerge. Brunnermeier and Parker (2005) treat a general-equilbrium model in which beliefs are essentially choice variables in the first period; an agent manipulates his beliefs about the future to maximize his felicity, which depends on future utility flow. Caplin and Leahy (2001) present an axiomatic model where agents have 'anticipatory feelings' prior to resolution of uncertainty, which may lead to time inconsistency. Koszegi (2006) considers an application of Caplin and Leahy (2001). Wu (1999) presents a model of anxiety.

[9]The notion of optimism can perhaps be seen as a form of belief manipulation, which is why I do not allow agents to be optimistic in this part of the analysis.

violation of the vNM model, and is one of the reasons cited by Kahneman, Knetch and Thaler (1991) for suggesting "a revised version of preference theory that would assign a special role to the status quo."

This model does not make use of a notion of reference points or of relative gains and losses.[10] In the cases where the choices also have an informational component on the agent's ability to perform a task well, a doubt-prone agent has incentive to choose the bundle that is less informative. This leads to a status quo bias when it is reasonable to assume that maintaining the status quo is a less informative indicator of the agent's ability than other actions. Since this model does not resort to reference points, there is no arbitrariness in defining what constitutes a gain and what constitutes a loss. The bias of a doubt-prone agent is always towards the least-informative signal of his ability. In instances where the status quo provides the most informative signal, the bias would be *against* the status quo. For example, an individual could have incentive to change activities frequently rather than obtaining a sharp signal of his ability in one particular field.

This framework admits other instances of seemingly paradoxical behavior. In one example, an individual pays a firm to invest for him even though he does not expect that firm to have superior expertise. In other words, the agent's utility not only depends on the outcome, but also on who makes the decision. This result is not due to a cost of effort, but rather to the amount of information acquired by the decision maker. This framework can also be used in a political economy setting, as there are many government decisions that are never observed by voters. As shown in section 3, voters may have strong incentives to remain ignorant over these issues, even if information is free. This is in line with the well-known observation that there has been a consistently high level of political ignorance amongst voters in the U.S. (see Bartels (1996) for details). This model suggests that if voters care more about policies that they may never observe, then they have *less* incentive to acquire

[10]There are, however, examples of the status quo bias for which this model does not seem to provide as natural an explanation as does loss-aversion.

information. Finally, this framework also accommodates behavior associated with anticipated regret, including preferences for smaller menus and the Allais paradox. This analysis is outside the scope of the paper, and is instead conducted in chapter 2 of the dissertation.[11]

1.1.3 Relation to the literature

The approach used in this paper is related to, but distinct from, the recursive expected Utility (REU) framework introduced by Kreps and Porteus (1978), and extended by Epstein and Zin (1989), Segal (1990) and Grant, Kajii and Polak (1998, 2000).[12] These earlier contribution address the issue of temporal resolution, in which an agent has a preference for knowing now versus knowing later. While the REU framework treats the issue of the timing of the resolution, this paper treats the case of *no* resolution. Simply adding a 'never' stage to the REU space does not yield an equivalent representation. To demonstrate this point, I place the agent in a two-stage model (in section 5), but do not allow the agent to have preferences over temporal resolution. The agent may, however, change his preferences over unresolved lotteries over time. For instance, he may prefer to avoid information in the early stage, but be curious in the later stage. In addition to the formal differences between the two frameworks, there are also interpretational ones. The REU model captures a notion of 'anxiety' (wanting to know sooner or later) which is distinct from the notion of doubt-proneness (not wanting to know at all) addressed here.

This paper is structured as follows. Section 2 introduces the model and derives the representation theorem. Section 3 defines optimism and doubt-proneness, and discusses the connection between these two properties and risk-aversion. Section

[11] See Bell (1982), Loomes and Sugden (1982), and Sarver (2008) for theoretical models of anticipated regret. See Zeelenberg (1999) for a review.

[12] Grant, Kajii and Polak (1998) focus on preferences for early resolution of uncertainty, and Dillenberger (2008) considers preferences for one-shot resolution of uncertainty. Selden's (1978) framework is also closely related to the REU model.

4 presents applications of this model. Section 5 relaxes the main independence axiom of the framework and introduces an axiom that allows different classes of models to incorporate outcomes that are never observed. In addition, it presents a representation theorem for a two-period setting. Section 6 concludes. All proofs are in the appendix.

1.2 Model and Representation Theorem

This section derives a representation theorem, which I then extend in section 5. I use the following objects. Let $\mathbf{Z} = [\underline{z}, \bar{z}] \subset \Re$ be the outcome space, an let $\mathfrak{L}_0$ be the set of simple probability measures on $\mathbf{Z}$. For $f = (z_1, p_1; z_2, p_2; ...; z_m, p_m) \in \mathfrak{L}_0$, z_i occurs with probability p_i. I use the notation $f(z_i)$ to mean the probability p_i (in lottery f) that z_i occurs. Let $\mathfrak{L}_1$ be the set of simple lotteries over $\mathbf{Z} \cup \mathfrak{L}_0$. For $X \in \mathfrak{L}_1$, I use the notation $X = (z_1, q_1^I; ...; z_n, q_n^I;\ f_1, q_1^N; ...; f_m, q_m^N)$. Here, z_i occurs with probability q_i^I, and lottery f_j occurs with probability q_j^N. Note that $\sum_{i=1}^{n} q_i^I + \sum_{i=1}^{m} q_i^N = 1$. The reason for using this notation, rather than the simpler enumeration $q_1, q_2, ..., q_n$ is explained shortly. Let $\succeq$ denotes the agent's preferences over $\mathfrak{L}_1$, and $\succ$, $\sim$ are defined in the usual manner. Assume the agent's preferences are monotone.

For any $X = (z_1, q_1^I; z_2, q_2^I; ...; z_n, q_n^I;\ f_1, q_1^N; f_2, q_2^N; ...; f_m, q_m^N)$, the agent expects to observe the outcome of the first-stage lottery. He knows, for instance, that with probability q_i^I, outcome z_i occurs, and furthermore he knows that he will observe it. Similarly, he knows that with probability q_i^N, lottery f_i occurs. But while he does observe that he is now faced with lottery f_i, he does *not* observe the outcome of f_i. I refer to lottery f_i as an 'unresolved' lottery. I also use the notation q_i^I and q_i^N to distinguish between the probabilities that lead to prizes where the agent is informed of the outcome (since he directly observes which z occurs), and the probabilities that lead to prizes where he is not (since he only observes the ensuing lottery). The

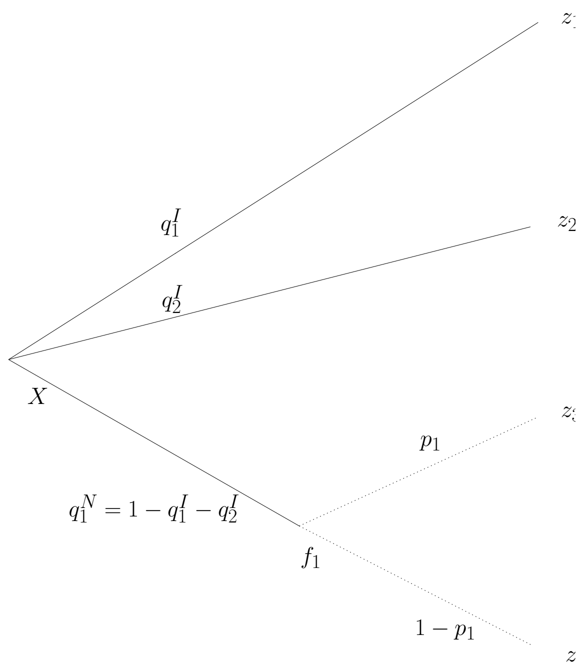

Figure 1.1: Lottery $X = (z_1, q_1^I; z_2, q_2^I; f_1, q_1^N)$, where $f_1 = (z_3, p_1; z_4, 1 - p_1)$

superscript I in q_i^I stands for 'Informed', and N in q_i^N for 'Not informed' (see figure 1).

Denote the degenerate one-stage lottery that leads to $z_i \in \mathbf{Z}$ with certainty $\delta_{z_i} = (z_i, 1) \in \mathfrak{L}_0$. The degenerate lottery that leads to $f_i \in \mathfrak{L}_0$ with certainty is denoted $\delta_{f_i} = (f_i, 1) \in \mathfrak{L}_1$. Note that all lotteries of form $X = f$, where $f \in \mathfrak{L}_0$, are purely resolved (or 'informed') lotteries, in the sense that the agent expects to observe whatever outcome occurs. Similarly, all lotteries of form $X = \delta_f$, where $f \in \mathfrak{L}_0$, are purely unresolved lotteries. With slight abuse, the notation $f \succeq f'$ (or $\delta_f \succeq \delta_{f'}$) is used, where $f, f' \in \mathfrak{L}_0$. In addition, $f \succeq \delta_f$ (or $\delta_f \succeq f$) indicates that the agent prefers (not) to observe the outcome of lottery f than to remain in doubt.

1.2.1 General axioms

The following certainty axiom **A.1** is assumed throughout:

AXIOM A.1 (Certainty): Take any $z_i \in \mathbf{Z}$, and let $X = \delta_{z_i} = (z_i, 1)$ and $X' = (\delta_{z_i}, 1)$. Then $X \sim X'$.

The certainty axiom **A.1** concerns the case in which an agent is certain that an outcome z_i occurs. In that case, it makes no difference whether he is presented with a resolved lottery that leads to z_i for sure or an unresolved lottery that leads to z_i for sure. He is indifferent between the two lotteries. Hence axiom **A.1** does not allow the agent to have a preference for being informed of something that he already knows for sure. This simple axiom provides a formal link between the agent's preferences over resolved lotteries and his preferences over unresolved lotteries. The following three axioms are standard.

AXIOM A.2 (Weak Order): $\succeq$ is complete and transitive.

AXIOM A.3 (Continuity): $\succeq$ is continuous in the weak convergence topology. That is, for each $X \in \mathfrak{L}_1$, the sets $\{X' \in \mathfrak{L}_1 : X' \succeq X\}$ and $\{X' \in \mathfrak{L}_1 : X \succeq X'\}$ are both closed in the weak convergence topology.

AXIOM A.4 (Independence): For all $X, Y, Z \in \mathfrak{L}_1$ and $\alpha \in (0, 1]$, $X \succ Y$ implies $\alpha X + (1-\alpha)Z \succ \alpha Y + (1-\alpha)Z$.

Focusing on axiom **A.4**, it is noteworthy that the agent's preferences $\succeq$ are on a richer space than in the standard framework. The independence axiom in the standard vNM model is taken on preferences over lotteries over outcomes, since all lotteries lead to outcomes that are eventually observed. In this paper, the agent's prize is not always an outcome z_i, and can instead be an unresolved lottery f_i. By assumption **A.4**, however, there is no axiomatic difference between receiving an outcome z_i as a prize and obtaining an unresolved lottery f_i as a prize. Under this approach, the rationale for using the independence axiom in the standard model holds in this case as well. Since this section aims to depart as little as possible from the vNM Expected Utility model, I assume the independence axiom **A.4** throughout. I relax this assumption in section 5 and replace it with a weaker axiom.

Axioms **A.1** through **A.4** suffice for this model to subsume the standard vNM representation for preferences over outcomes that the agent eventually observes. That is, suppose we focus on lotteries of form $X = f$, i.e. lotteries that lead to outcomes. Then all the standard vNM axioms over these lotteries hold, and the EU representation follows directly. These axioms are not sufficient, however, to characterize the agent's preferences over lotteries that do not resolve. If, for instance, the agent receives a lottery $X = \delta_f$, it is unclear what his 'perception' of unresolved lottery f is. The next step, therefore, is to consider axioms that allow us to characterize the agent's preferences over these 'purely' unresolved lotteries of form $X = \delta_f$. As there is a natural isomorphism between lotteries of form $X = \delta_f \in \mathfrak{L}_1$ and one-stage lotteries in $\mathfrak{L}_0$, define the preference relation $\succeq_N$ in the following way:

Definition of $\succeq_N$. For any $f^N, f'^N \in \mathfrak{L}_0$, $f \succeq_N f'$ if $\delta_f \succeq \delta_{f'}$.

Define $\succ_N$ and $\sim_N$ in the usual way. I do *not* assume independence over the preference relation $\succeq_N$, for the following reason. Suppose that an agent is given a choice between three lottery tickets. The first ticket consists of a lottery $f = (\$1000, 1/3; \$400, 1/3; \$0, 1/3)$. With probability 1/3, the ticket yields \$1000, with probability 1/3 it yields \$400, and it yields 0 otherwise. The second ticket consists of lottery $f' = (\$1000, 1/2; \$0, 1/2)$ and the third ticket consists of $f'' = (\$400, 1) = \delta_{400}$, which yields \$400 for certain. In addition, suppose that the agent does not purchase the ticket for himself, but for a charitable organization that he holds in high esteem.

It is plausible that a risk-averse agent prefers the safe lottery δ_{400} to lottery f', *if* he expects to observe the outcome of the lotteries (for instance, if the charity thanks him for his contribution of the quantity it receives). But it may also be the case that the same agent has different preferences and choose risky lottery f' over the safe lottery δ_{400} ($f' \succ_N \delta_{400}$), if he donates the unresolved ticket to the charity and does not expect to observe which outcome occurs. There is a 1/2 chance that the charity has received \$1000, and he does not expect to ever find out if it has received \$0. These preferences may be driven by a notion of 'optimism'.

Now compare lotteries f to f', still for the case in which the agent does not expect to observe the resolution of uncertainty. It is also plausible that the agent prefers lottery f to f' ($f \succ_N f'$): lottery f is less risky than lottery f', and at the same time he still does not find out whether the charity has received \$0:

$$(\$1000, 1/3; \$400, 1/3; \$0, 1/3) \succ_N (\$1000, 1/2; \$0, 1/2) \succ_N \delta_{400}.$$

These preferences appear reasonable, but they violate independence. In fact, they violate the stronger axiom of betweenness, and so do not fall in the Dekel (1986) class of preferences.[13]

This example illustrates that there are two distinct notions that play a role in the agent's preference over unresolved lotteries. The agent may be risk-averse over unresolved lotteries, and this risk-aversion manifests itself in his comparison between lottery f and the more risky lottery f'. At the same time, he may be 'optimistic' that the good outcome has occurred if he does not observe the lottery, which affects his assessment of lottery f', compared to the safe lottery δ_{400}. A single utility function v cannot capture both these notions, since risk-aversion and optimism do not necessarily coincide, as in the previous example. However, both risk-aversion and optimism are contributing factors to the agent's preferences to remain in doubt.

I now assume the Rank-Dependent Utility (RDU) axioms, which are general enough to allow the previous example. The RDU representation allows for two

[13]Note that $f = \frac{2}{3}f' + \frac{1}{3}\delta_{400}$. This is a violation of independence (and betweenness) because the following does not hold: $f' \succ_N \frac{2}{3}f' + \frac{1}{3}\delta_{400} \succ_N \delta_{400}$. More specifically, this violates quasi-convexity.

functions, v and w, the first that reweighs the outcomes (identically to the vNM model), and the second reweighs the probabilities. I show, in the following section, that an RDU representation captures a notion of risk and optimism that are suitable to this model, even though my formal definition of optimism will be different from the accepted RDU definition. I later consider conditions which force the function w to be linear, essentially reducing the representation of $\succeq_N$ to a vNM representation.[14]

1.2.2 RDU representation for $\succeq_N$.

The following notation is convenient for the RDU representation. For lottery $f = (z_1, p_1; ; ...; z_m, p_m) \in \mathfrak{L}_\mathfrak{o}$, the $z_i's$ are ordered from smallest to highest, i.e. $z_m > ... > z_1$. Recall that the agent's preferences are monotone, which implies that $\delta_{z_m} \succ_N ... \succ_N \delta_{z_1}$. In addition, p_i^* denotes the probability of reaching outcome z_i or an outcome that is weakly preferred to z_i. That is, $p_i^* = \sum_{j=i}^{m} p_j$. Note that for the least-preferred outcome z_1, $p_1^* = 1$. Probabilities p_i^* are referred to here as 'decumulative' probabilities. The RDU form, introduced by Quiggin (1982), is defined in the following manner:[15]

Definition (RDU) Rank-dependent utility (RDU) holds if there exists a strictly increasing continuous probability weighting function $w : [0,1] \to [0,1]$ with $w(0) = 0$ and $w(1) = 1$ and a strictly increasing utility function $v : \mathbf{Z} \to \Re$ such that for all $f, f' \in \mathfrak{L}_\mathfrak{o}$,

$$f \succ_N f' \text{ if and only if } V_{RDU}(f) > V_{RDU}(f')$$

[14] The notion of 'optimism' may seem at odds with the previous claim that an agent who is not allowed to manipulate his beliefs may still choose to 'self-handicapping'. That is, one interpretation of a rank-dependent utility representation is that the agent distorts the actual probability. For this reason, In the analysis of self-handicapping (section 4), I do *not* allow the agent to be either optimistic or pessimistic.

[15] See also Yaari (1987), and Diecidue and Wakker (2001) for a thorough discussion of RDU.

where V_{RDU} is defined to be: for all $f = (z_1, p_1; z_2, p_2; ...; z_m, p_m)$,

$$V_{RDU}(f) = v(z_1) + \sum_{i=2}^{m}[v(z_i) - v(z_{i-1})]w(p_i^*)$$

Moreover, v is unique up to positive affine transformation.

Note that if the weighting function w is linear, then V_{RDU} reduces to the standard EU form.[16] I now briefly discuss the axiomatic foundation of the RDU representation, in the context of this model. Suppose that

$$f_\alpha = (z_1, p_1; ...; \alpha, p_i; ...; z_m, p_m) \succeq_N (z'_1, p_1; ...; \beta, p_i; ...; z'_m, p_m) = f'_\beta$$
$$f'_\kappa = (z'_1, p_1; ...; \kappa, p_i; ...; z'_m, p_m) \succeq_N (z_1, p_1; ...; \gamma, p_i; ...; z_m, p_m) = f_\gamma$$

where $\alpha, \beta, \gamma, \kappa \in \mathbf{Z}$. Comparing lotteries f_α and f_γ, the only difference is in whether α or γ is reached with probability p_i. Since all the other outcomes are the same in both lotteries and are reached with the same probabilities, the difference is in the value of outcome α compared to the value of outcome γ (and similarly for f'_β,f'_κ and β, κ). In the comparison of $f_\alpha \succeq_N f'_\beta$ and $f'_\kappa \succeq_N f_\gamma$, all the probabilities of reaching the (rank-preserved) outcomes are the same. For that reason, this model assumes that the switch in preference is due to a difference in the value of outcomes α and β relative to γ and κ, and not in the way the probabilities are aggregated. It is precisely this property that RDU provides: if $f_\alpha \succeq_N f'_\beta$ and $f'_\kappa \succeq_N, f_\gamma$, and if $\succeq_N$ is of the RDU form, then $v(\alpha) - v(\beta) \geq v(\gamma) - v(\kappa)$. Note that this does not depend on the choice of $z's$ and $p's$, and so the following axiom, adapted from Wakker (1994), must hold:

[16] This is not the most common form of RDU; this notation is taken from Abdellaoui (2002). Given the rank-ordering above, the typical form would be $V_{RDU} = \sum_{i=1}^{n-1}[w(p_i^*) - w(p_{i+1})^*]v(z_i) + w(p_n)v(z_n^*)$. It is easy to check that the two representations are identical.

AXIOM N.RDU (Wakker tradeoff consistency for $\succeq_N$):

Let $f_\alpha = (z_1, p_1; ...; \alpha, p_i; ...; z_m, p_m)$, $f_\gamma = (z_1, p_1; ...; \gamma, p_i; ...; z_m, p_m)$, $f'_\beta = (z'_1, p_1; ...; \beta, p_i; ...; z'_m, p_m)$ and $f'_\kappa = (z'_1, p_1; ...; \kappa, p_i; ...; z'_m, p_m)$. If:

$$f_\alpha \succeq_N f'_\beta$$
$$f'_\kappa \succeq_N f_\gamma$$

then for any lotteries $g_\alpha = (\hat{z}_1, \hat{p}_1; ...; \alpha, \hat{p}_i; ...; \hat{z}_{\hat{m}}, \hat{p}_{\hat{m}})$, $g_\gamma = (\hat{z}_1, \hat{p}_1; ...; \gamma, \hat{p}_i; ...; \hat{z}_{\hat{m}}, \hat{p}_{\hat{m}})$, $g'_\beta = (\hat{z}'_1, \hat{p}_1; ...; \beta, \hat{p}_i; ...; \hat{z}'_{\hat{m}}, \hat{p}_{\hat{m}})$, $g'_\kappa = (\hat{z}'_1, \hat{p}_1; ...; \kappa, p_i; ...; \hat{z}'_{\hat{m}}, \hat{p}_{\hat{m}})$ such that $g_\gamma \succeq_N g'_\kappa$, it must be that $g_\alpha \succeq_N g'_\beta$.

Under this axiom, only the values of α,β,γ and κ are relevant to the ordering of the agent's preferences when all the probabilities of reaching all other outcomes are the same across the four lotteries. In fact, as shown in Wakker (1994), this axiom is sufficient, along with stochastic dominance and continuity, for the RDU representation to hold. Using this result, the general representation theorem for $\succeq$ is as follows:

Main Representation Theorem. *Suppose axioms **A.1** through **A.4** and axiom **N.RDU** hold. In addition, suppose stochastic dominance holds for $\succeq_N$. Then there exist strictly increasing, continuous and bounded functions $u : \mathbf{Z} \to \mathbb{R}$, $v : \mathbf{Z} \to \mathbb{R}$, $w : [0,1] \to [0,1]$ with $w(0) = 0$ and $w(1) = 1$, such that for all $X, Y \in \mathfrak{L}_1$,*

$$X \succ Y \text{ if and only if } W(X) > W(Y)$$

where W is defined to be: for all $X = ((z_1, q_1^I; ...; z_n, q_n^I;\ f_1, q_1^N; ...; f_m, q_m^N) \in \mathfrak{L}_1$,

$$W(X) = \sum_{i=1}^{n} q_i^I u(z_i) + \sum_{j=1}^{m} q_j^N u\left(v^{-1}\left(V_{RDU}(f_j^N)\right)\right)$$

and

$$V_{RDU}(f) = v(z_1) + \sum_{h=2}^{m}[v(z_h) - v(z_{h-1})]w(p_i^*).$$

Moreover, u and v are unique up to positive affine transformation.

Note that u remains the utility function associated with the general lotteries (and final outcomes). In addition, v is the utility function associated with unresolved lotteries, and w is the probability weighting function associated with unresolved lotteries. It is not immediately clear from this representation what doubt-proneness implies, in terms of the shapes of the functions. The next section defines optimism, and formally relates it to the accepted notion of optimism in an RDU setting. I then connect doubt-proneness, risk-aversion, and this new notion of optimism.

1.3 Risk-aversion, doubt-proneness and optimism

In this section, I focus on the relationship between doubt-proneness and the shapes of the functions u, v and w. I first define formally what optimism means in this context. Returning to the charity example from the previous section, recall that lottery $f = (\$1000, 1/3; \$400, 1/3; \$0, 1/3)$, lottery $f' = (\$1000, 1/2; \$0, 1/2)$ and lottery $\delta_{400} = (\$400, 1)$. While $f' \succ_N \delta_{400}$, it is not the case that $f' \succ_N af' + (1 + a)\delta_{400} \succ_N \delta_{400}$ for all a, which the independence axiom would imply. In this example, $f = \frac{2}{3}f' + \frac{1}{3}\delta_{400} \succ_N f'$.

The notion of optimism over unresolved lotteries I aim to capture allows the agent prefer more 'scrambled' information, since it essentially allows him to form a better assessment of these unresolved lotteries. Consider lottery δ_{400}, in which the agent is certain that the outcome is \$400. Now suppose that it is mixed with a lottery $\tilde{f}' = (\$400 + \delta, 1/2; \$400 - \epsilon)$, where $\tilde{f}'$ is chosen such that $\tilde{f}' \sim_N f'$, and ϵ is close to 0.[17] Specifically, consider the mixture $\tilde{f} = 2/3\ f' + 1/3\ \delta_{400} = (\$400 + \delta, 1/3; \$400, 1/3; \$400 - \epsilon, 1/3)$ (see figure 2). If independence were to hold,

[17] For δ to also be close to 0, \$400 would have to be close to the certainty equivalent of the unresolved lottery $f' = (\$1000, 1/2; \$0, 1/2)$.

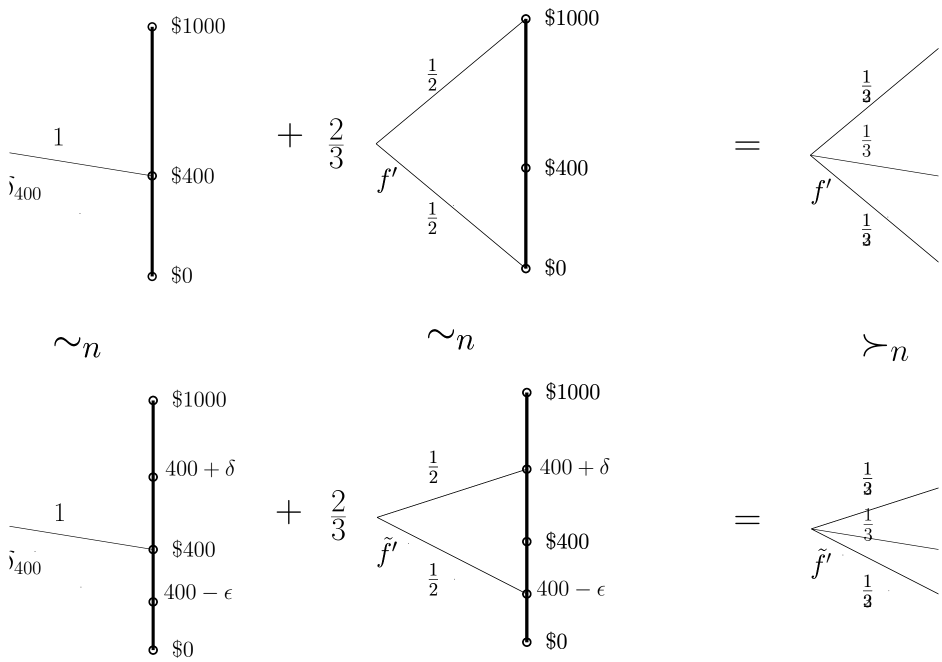

Figure 1.2: Optimism.

then $f \sim \tilde{f}$. But I also allow $f \succ_N \tilde{f}$, with the reasoning that the optimist agent prefers knowing as little as as possible about the unresolved lottery. With lottery f, the optimist can form a more reassuring perception of the outcome, as it could be much higher ($ 1000). With lottery $\tilde{f}$, however, as ϵ becomes smaller and smaller, it becomes less and less attractive to the optimist agent, as he is more and more certain of the vicinity of the outcome. In brief, an optimist has a preference for more 'scrambled' information. A pessimistic agent, on the other hand, prefers less scrambled information, since knowing less would lead him to form a more negative perception. I allow the agent to be optimist, pessimism or neutral (i.e. independence may hold), but I assume that his preferences are preserved, given a specific mixture a and specific probabilities. That is, if the agent prefers unresolved lottery f to $\tilde{f}$, as in the example above, then this preference is preserved as ϵ becomes smaller. I refer to this property, which I now generalize, as 'information scrambling consistency' (ISC).

Definition (ISC) $\succeq_N$ satisfies information scrambling consistency (ISC) if:
let $f = (z_1, p_1; ...; z_i; p_i; z_{i+1}, p_{i+1}; ...; z_n, p_n)$, $f' = (z_1, p_1; ...; z'_i; p_i; z'_{i+1}, p_{i+1}; ...; z_n, p_n)$ $\in \mathfrak{L}_o$ such that $f \sim_N f'$, and *case 1*: $(z'_i, z'_{i+1}) \subset (z_i, z_{i+1})$ (*case 2*: $(z_i, z_{i+1}) \subset (z'_i, z'_{i+1})$). If, for some $a \in (0, 1)$ and some $z \in (z'_i, z'_{i+1})$:

$$af + (1-a)\delta_z \succeq_N af' + (1-a)\delta_z,$$

then it must also be that:

$$a\tilde{f} + (1-a)\delta_{\tilde{z}} \succeq_N a\tilde{f}' + (1-a)\delta_{\tilde{z}}$$

for any
$\tilde{f} = (\tilde{z}_1, p_1; ...\tilde{z}_i; p_i; \tilde{z}_{i+1}, p_{i+1}; ...; \tilde{z}_n, p_n)$, $\tilde{f}' = (\tilde{z}_1, p_1; ...\tilde{z}'_i; p_i; \tilde{z}'_{i+1}, p_{i+1}; ...; \tilde{z}_n, p_n)$ and $\tilde{z}$ such that $\tilde{z} \in (\tilde{z}'_i, \tilde{z}'_{i+1}) \subset (\tilde{z}_i, \tilde{z}_{i+1})$ (*case 2*: $\tilde{z} \in (\tilde{z}_i, \tilde{z}_{i+1}) \subset (\tilde{z}'_i, \tilde{z}'_{i+1})$).

A preference for more scrambled information (optimism) corresponds to case 1, i.e. preferring $af + (1-a)\delta_z \succ af' + (1-a)\delta_z$ when $(z'_i, z'_{i+1}) \subset (z_i, z_{i+1})$. Similarly, a preference for less scrambled information (pessimism) corresponds to case 2. The appeal of the RDU representation is that it satisfies the ISC property:

Theorem 2. *Suppose that RDU holds for $\succeq_N$. Then $\succeq_N$ satisfies ISC.*

A local preference for more scrambled information, which I refer to as local optimism, does *not* correspond to the accepted RDU notion of optimism, analyzed by Wakker (1994). I prove, however, that an agent has a *global* preference for more scrambled information if and only if the weighting function w is concave, and therefore corresponds to the Wakker notion of (global) optimism. Defining (global) optimism:

Definition (Optimism) The preference relation $\succeq_N$ exhibits optimism if and only if $\succeq_N$ always exhibits a preference for more scrambled information. That is, for any

$f = (z_1, p_1; ...z_i; p_i; z_{i+1}, p_{i+1}; ...; z_n, p_n)$, $f' = (z_1, p_1; ...z_i'; p_i; z_{i+1}', p_{i+1}; ...; z_n, p_n) \in \mathfrak{L}_o$ such that $f \sim_N f'$, and $(z_i', z_{i+1}') \subset (z_i, z_{i+1})$, and for all $a \in (0,1)$ and $z \in (z_i, z_{i+1})$,

$$af + (1-a)\delta_z \succeq_N af' + (1-a)\delta_z.$$

The next theorem demonstrates that this definition of optimism corresponds to the accepted RDU definition.

Theorem 3. *Suppose that $\succeq_N$ satisfies RDU, and let w be the associated weighting function. Then w is concave (convex) if and only if $\succeq_N$ exhibits optimism.*

I now define doubt-proneness in the natural way.

Definition (Doubt-proneness)

- An agent is doubt-prone *somewhere*if there exists some f such that $\delta_f \succ f$.
- An agent is doubt-prone *everywhere* if: (i) there exists no $f \in \mathfrak{L}_o$ such that $f \succ \delta_f$ and (ii) there exists some f such that $\delta_f \succ f$.

An agent who prefers not to observe the resolution of some lottery than to observe it is doubt-prone somewhere. An agent who (weakly, and strictly for one lottery) prefers not to observe the outcome of *any* lottery is doubt-prone everywhere. Doubt-aversion is defined in a similar manner. The next result below connects doubt-proneness, the properties of the utility functions, and the properties of the probability weighting function $w(p)$. A similar result hold for doubt-aversion, and is deferred to the appendix.

Theorem 4. *Suppose that axioms **A.1** through **A.4** and the RDU axioms hold, and let u and v be the utility functions associated with the resolved and unresolved lotteries, respectively, and w be the decision weight associated with the unresolved lotteries. In addition, suppose that u and v are both differentiable. Then:*

(i) If there exists a $p \in (0,1)$ such that $p < w(p)$, then the agent is doubt-prone somewhere. Similarly, if there exists $p' \in (0,1)$ such that $p' > w(p')$, then the agent is doubt-averse somewhere.
(ii) If the agent is doubt-prone everywhere, then $p \leq w(p)$ for all $p \in (0,1)$. Moreover, if v exhibits stronger diminishing marginal utility than u, then $\succeq_N$ violates quasi-convexity (that is, there exists some $f', f'' \in \mathfrak{L}_0$, and $\alpha \in (0,1)$ such that $f' \succ f''$ and $\alpha f' + (1-\alpha)f'' \succ_N f'$).

The differentiability assumption, though common, may seem bothersome as it is not taken over the primitives. Alternatively, we could make an assumption over the primitives that guarantees (for instance) strict concavity of u and v, which would in fact be sufficient for the result.[18] Given the results above, an assumption or deduction over the agent's doubt-attitude has testable implications concerning his aggregation of probabilities (w) for unresolved lottery, and vice-versa. In addition, these implications can be disentangled from the agent's diminishing marginal utility. Since it is not necessary that w satisfies the same empirical properties as the typical case considered under rank-dependent utility, an experimental study would be useful for a better understanding of the shape of w. If, in addition to doubt-proneness, mean-preserving risk-aversion (in the standard sense) of $\succeq_N$ is assumed, then the RDU representation collapses to the recursive EU representation:

Corollary 4.1. *Suppose that the conditions of theorem 4 all hold. Then the following two statements are equivalent:*

(i) Preference $\succeq$ displays doubt-proneness everywhere and $\succeq_N$ displays mean preserving risk-aversion.

(ii) Function V_{RDU} is of the EU form (i.e. $w(p) = p$ for all $p \in [0,1]$), both u and v are concave, and $u = \lambda \circ v$ for some continuous, concave, and increasing λ.

[18] For a discussion of the differentiability assumption, see Chew, Karni and Safra (1987).

This result further shows that attitude toward risk and attitude towards doubt constrain the probability weighting function, and can in fact completely characterize it.[19] But note that in an RDU setting, mean-preserving risk aversion is not identical to diminishing marginal utility. That is, the previous result does not imply that a doubt-prone agent who obeys risk aversion cannot have a concave utility function v. I now focus a counterexample for which doubt-proneness is entirely due to the weighting factor w, and *not* the difference in concavity between u and v.

Consider an agent for whom functions u and v are identical. It is already immediate from theorem 4 that for a doubt-prone agent, it is necessary that $p \leq w(p)$ for all p. In fact, this condition is sufficient.[20] The following result does not require differentiability.

Theorem 5. *Suppose that the conditions of theorem 4 all hold. Furthermore, suppose that $u(z) = v(z)$ for all $z \in \mathbf{Z}$ (or, more generally, $u = \lambda \circ v$ for some continuous, weakly concave, and increasing λ). Then the agent is doubt-prone everywhere if and only if $p \leq w(p)$ (with $p < w(p)$ for some $p \in (0, 1)$ if $u(z) = v(z)$ for all $z \in \mathbf{Z}$).*

It follows that an optimistic agent for whom u is identical v (or for whom u is more concave than v) must be doubt-prone. These results therefore connect optimism, doubt-proneness, and risk-aversion (in the standard sense). Before concluding this section, note that extensive research has been conducted on the shape of w in the usual RDU setting, in which uncertainty eventually resolves.[21] As this a different setting, I have not made similar assumptions over the shape of w. Instead, I have shown that the induced preferences to remain in doubt have strong implications on the weighting function w. Consider, for example, the common assumption that w is

[19]This last corollary is similar to a result in Grant, Kajii and Polak (2000) but with a notion of doubt-proneness that is weaker than the preference for late-resolution that would be required in the framework they use; the difference in assumptions is due to the difference in settings. It is also of note that under Grant, Kajii and Polak (2000)'s restriction, there is no need to assume differentiability, as it is in fact implied.

[20]It is clear that if $p = w(p)$ for all $p \in (0, 1)$ and if $u(z) = v(z)$ for all $z \in \mathbf{Z}$, then the agent is doubt-neutral.

[21]See, for instance, Karni and Safra (1990), and Prelec (1998) for an axiomatic treatment of w.

S-shaped (concave on the initial interval and convex beyond). In that case, it must be that the agent is doubt-prone for some lotteries and doubt-averse for others. But an empirical discussion of whether w is S-shaped in this setting is outside the scope of this paper. I now turn to the applications.

1.4 Applications

I consider two applications in this section. In the first, an agent's utility depends directly on his ability, since it is related to his self-image. He may never fully observe his ability, but his success at performing tasks provides him with an imperfect signal. How well he performs a task also depends on his effort. Performing a task better provides him with a reward, and so in the standard EU setting, he would always put in as much effort as he can if effort is costless. In this setting, however, there is a tradeoff between obtaining a better reward by putting in more effort and obtaining a coarser signal of ability by putting in less effort. Under some conditions, the agent has an incentive to self-handicap, as is shown below. This setup also accommodates other well known behavioral patterns. Under one version of this setup, an agent has an incentive to remain with the status quo. In another version of this setup, a risk-neutral agent prefers less risky bonds with a lower expected return to more risky stocks with a higher expected return. This agent is also willing to pay a firm to invest for him, even if he knows that the firm does not have superior expertise.

In the second application, voters all have the same preferences, but they do not know who the better candidate is. However, they can acquire this information at no cost. I demonstrate that there are equilibria in which they choose to remain ignorant, and the wrong candidate is as likely to win as the right candidate.

1.4.1 Preservation of self-image

I first introduce a general setup, before analyzing the implications of the results in different contexts. I assume that the agent places direct value on his ability, independently of the effect it has on his monetary reward. Arguably, individuals care about their self-image, and would rather think of themselves as being of higher ability than lower ability. Their success at achieving their goals, given how much effort they put in, provides them with imperfect signals of their ability.

Suppose then that the agent is endowed with ability (or type) $t \in [\underline{t}, \overline{t}] \in \mathbb{R}$. He does not know what his ability is, but his prior probability of having ability t is $p(t)$. The agent chooses effort $e \in [\underline{e}, \overline{e}] \in \mathbb{R}$, to obtain a reward $m \in [\underline{m}, \overline{m}] \in \mathbb{R}$. Although the agent may never observe his ability, he does observe m. The reward depends on his ability, the effort he puts in, and an intrinsic uncertainty. Let $p(m|e,t)$ denote his probability of receiving reward m given his effort e and his ability t. Since he does not know what his ability is ex-ante, his prior probability of receiving m given effort e is $p(m|e) = \sum_{t \in [\underline{t}, \overline{t}]} p(m|e,t)p(t)$. Assume that the expected reward is higher if he puts in more effort for any given ability, and it is higher if he is of higher ability at any given effort level: $Em(e,t) > Em(e,t') \Leftrightarrow t > t'$, and $Em(e,t) > Em(e',t) \Leftrightarrow e > e'$.[22]

The agent's value function W depends on both his reward m and on his intrinsic ability t. Assume that his utility for m is linear; more precisely, his expected utility over m is $Em(e)$. In addition, it is linearly separable from his utility over t. He is weakly risk-averse over t (for both resolved and unresolved lotteries) as well as doubt-prone.[23] As in the theory section, let u be his resolved utility, and let v be his unresolved utility. Notice that with these assumptions, the agent's preferences over his ability reduce to a two-period Kreps-Porteus (KP) representation.

In the standard case in which the agent expects to observe both his ability t and

[22] All the probability distributions in this section have finite support.

[23] Note that by corollary 4.1, the weighting function here is linear, w(p) = p. In addition, since the agent is doubt-prone and risk-averse in the unresolved lotteries, corollary 4.1 also implies that that he is risk-averse in the resolved lotteries.

his reward m, then his value function is:

$$W(e) = Em(e) + Eu(t)$$

Since effort is costless, it is immediate that he should put in the highest level of effort, $e = \overline{e}$. But now suppose that he does *not* necessarily observe his ability ex-post. In this case, when he receives his monetary reward, he simply updates his probability on his ability, given m and his chosen effort level e. His value function is therefore:

$$W(e) = Em(e) + \sum_m p(m|e)u\left(v^{-1}(Ev(t|m,e))\right)$$

Depending on the functional form, the agent might not put in effort $e = \overline{e}$. His effort level also depends on his incentive to obtain the least information concerning his ability, since he is doubt-prone. In other words, he takes into account what the combination of his effort and the reward he obtains allow him to deduce about his ability. Suppose that there is a unique effort level e_o (the 'ostrich' effort) that is entirely uninformative, i.e. $p(t|m, e_o) = p(t)$ for all $t \in [\underline{t}, \overline{t}]$ and for all $m \in [\underline{m}, \overline{m}]$. Note that e_o provides the agent with the highest expected utility over his ability. That is, define

$$C(e) \equiv u\left(v^{-1}(Ev(t))\right) - \sum_m p(m|e)u\left(v^{-1}(Ev(t|m,e))\right)$$

As shown in the appendix, it is always the case that $C(e) > 0$ (for $e \neq e_o$) for a doubt-prone agent, with $C(e_o) = 0$. Redefining the value function to be $\tilde{W}(e) = W(e) - u\left(v^{-1}(Ev(t))\right)$, the agent maximizes

$$\tilde{W}(e) = Em(e) - C(e)$$

Hence $C(e)$ is effectively the 'shadow' cost of effort due to acquiring information that he would rather ignore. The optimal effort level depends on the importance of the expected reward $Em(e)$ relative to the agent's disutility of acquiring information concerning his ability, as is captured by $C(e)$. Suppose now that $e_0 = \underline{e}$, and that the agent obtains a more informative signal (in the Blackwell sense) for a higher effort e. Then $C(\underline{e}) = 0$, and $C(e)$ is strictly increasing, so that the 'shadow' cost is increasing in effort level. The following simple example serves as an illustration.

Numerical Example

Let $\underline{e} = \underline{t} = 0$, $\overline{e} = \overline{t} = 1$, $p(t = 0) = \frac{1}{2}$ and $p(t = 1) = \frac{1}{2}$. The agent's reward m only takes value \$0 and \$100. The probability of obtaining reward $m = \$100$ given e and t are:

$$p(m = \$100 | t = 1, e) = e$$
$$p(m = \$100 | t = 0, e) = 0$$

and $p(m = \$0 | t, e) = 1 - p(m = \$100 | t, e)$. The utility functions are $u = a\sqrt{t}$ for some $a > 0$, and $v = t$.

Note that in this example, the completely uninformative effort e_o is equal to 0. At effort $e = 0$, he is sure to obtain \$0, and his posterior on his ability is the same as his prior. As he puts in more effort, he obtains a sharper signal of his ability. If he puts in maximum effort $e = 1$, then he will fully deduce his ability ex-post: if he obtains \$100 then he knows he has ability $t = 1$, and if he obtains \$0 then he knows he has ability $t = 0$. His value function is now:

$$\tilde{W}(e) = 50 - C(e)$$

where $C(e) = \frac{a}{2}(\sqrt{2} - e - \sqrt{2 - 3e + e^2})$.

The optimal level of effort e^* is in the full range $[0, 1]$, depending on a. More precisely, for interior solutions, e^* is the smaller root of the equation $e^2 - 3e + \frac{2d-9}{d-4} =$

0, where $d = \left(\frac{200}{a} + 2\right)^2$. As a increases, the monetary reward m becomes less significant, and e^* decreases. As a decreases, the agent's utility over his ability becomes less significant, and the effort level increases (see appendix for details).

Self-handicapping

The setup presented here can be applied to several different contexts, the most immediate of which is self-handicapping. There is strong anecdotal evidence that people are sometimes restrained by a 'fear of failure', and will not put in as much effort as they could. Berglas and Jones (1978) find in an experiment that individuals deliberately impede their own chances of success, and attribute this behavior to people's desire to protect the image of the self.[24] The amount of optimal self-handicapping depends on the doubt-proneness of the agent, and how good of a signal he expects to obtain. As discussed above, choosing a higher effort level leads to a tradeoff between the improved reward $Em(e)$ and the incurred cost $C(e)$ of learning more about one's actual ability. This model also confirms Berglas and Jones' intuition that those who are more likely to self-handicap are not the most successful or the least successful, but rather those who are uncertain about their own competence. Akerlof and Dickens' (1982) observation that people will remain ignorant so as to protect their ego is also in agreement with the implications of this framework. But notice that here, self-handicapping follows from the agent's doubt-proneness over his decision making ability, and not from an ability to lie to himself or to manipulate his beliefs in any way.

Status quo bias

The endowment effect and status quo bias are analyzed by Kahneman, Knetch and Thaler (1991), and are explained using framing effects and loss aversion. The agent's preference for avoiding a loss is taken to be stronger than his preference for making

[24]See Benabou and Tirole (2002) for an explanation that uses manipulable beliefs.

a gain, and the reference point for what constitutes a gain or a loss is assumed to be the status quo. However, Samuelson and Zeckhauser (1988) do not view the status quo bias to be solely a consequence of loss-aversion: "Our results show the presence of status quo bias even when there are no explicit gain/loss framing effects.... Thus, we conclude that status quo bias is a general experimental finding – consistent with, but not solely prompted by, loss aversion." The framework discussed here can be applied to some settings in which a status quo bias is present.

Suppose that e now represents a choice over different bundles rather than effort. For instance, suppose that the agent only places probability on $\underline{e}$ and $\overline{e}$, and that $\underline{e}$ corresponds to keeping the current allocation, while $\overline{e}$ corresponds to switching to another bundle. In addition, suppose that acquiring a bundle also carries information on the agent's decision making ability. In this case, rather than representing a cost of effort, $C(e)$ represents the cost of deviating from the bundle that is least informative of the agent's decision making ability. Suppose that $e_0 = \underline{e}$, so that keeping the same bundle is uninformative. Then the agent exhibits a status quo bias, since inaction (keeping the same bundle) has information cost $C(e_0) = 0$. Note, however, that when keeping the status quo bundle is more informative than obtaining other bundles, then a doubt-prone agent would be biased *against* the status quo.

The key difference between the model presented here and the standard vNM model is that this model allows for an asymmetry in the value of acquiring a bundle compared to losing that bundle. The bundle itself does not change value based on whether the agent is endowed with it or not, and in that sense there is no framing effect. Instead, acquiring a new bundle *in itself* has different informational implications than selling it. In the case where the unobserved prize is the agent's ability, then acquiring a new bundle may provide him with more information on his ability than keeping his current allocation. A more thorough explanation can be found in chapter 2 of the dissertation.

Bonds, stocks and paternity

Consider the case in which e represents an investment decision rather than effort. A higher e represents a more risky investment, but in expectation it leads to a higher monetary reward. As before, t corresponds to a notion of ability. An individual who is of higher decision-making ability makes a wiser investment choice and therefore obtains a higher expected monetary reward, given the chosen risk level. For instance, $\underline{e}$ might be a portfolio consisting solely of bonds, while $\overline{e}$ consists solely of higher-risk stocks. Maintain the assumption that $e_o = \underline{e}$. In other words, the riskless option is also least informative concerning the agent's potential as an investor.

In this setting, although the agent is risk-neutral in money, his chosen bundle e^* may still consist of more bonds than it would if the reward were purely monetary, as there is a bias towards $\underline{e}$.[25] In addition, suppose that a firm exists which offers to invest the agent's money in his place. Even if the agent puts the same prior on his ability in investing as he does on the firm's, he still agrees to pay. Since the optimal level of risk in this case is $\overline{e}$, he is willing to pay up to $Em(\overline{e}) - Em(e^*) + C(e^*)$. In fact, even if the firm were to choose the suboptimal level e^*, he would be willing to pay up to $C(e^*)$.

In the standard EU model, the agent's choice would only depend on the monetary reward he expects to obtain. In contrast, the framework presented here allows the agent's choice to depend on the decision making process as well as on the reward he expects to receive. That is, the agent bases his choice on the *manner* in which he expects to obtain the monetary reward.

1.4.2 Political Ignorance

The high degree of political ignorance of voters has been thoroughly researched, particularly in the US (see Bartels (1996)). Given the length of electoral campaigns in American politics, the amount of media coverage and the accessibility of

[25] Of course, no claim is made concerning the empirical significance of this effect.

informational sources, it seems that the cost of acquiring information should not be prohibitive for voters. Note that there are political issues whose resolution the voters may never observe. For instance, the voters may choose not to observe the amount of foreign aid given, the degree of lobbying or nepotism, or the government stance on interrogation methods. For those issues, a doubt-prone agent may have incentive to ignore information even if it is free. In other words, making information more accessible would not necessarily have a strong impact on the individual's informativeness on these issues. Since voters affect the election result as a group, each individual's decision to acquire information has an externality on other voters and on *their* decision to acquire information. This section discusses a very simple example in which voters' information acquisition plays a dominant role on the other voters' decision to acquire information. Although voting is sincere, there is a strategic aspect to the decision to acquire information.

Consider an economy in which N citizens care about issue $\gamma \in [0,1]$, which is determined by a politician that they vote for. They can choose not to observe what the politician does. Suppose that there are two candidates, A and B. One of the two will choose policy $\gamma = 0$ if elected, and the other will choose $\gamma = 1$. The voters do not know which one is which, and place probability $1/2$ that A will choose $\gamma = 0$, and $1/2$ that A will choose $\gamma = 1$ (and similarly for B). But they can acquire that information at no cost, if they choose to do so. Let p_i be the ex-post probability that the ith agent places on the winner being the candidate who implements $\gamma = 1$, where $i \in \{1, .., N\}$. The timing is as follows:

1) Each voter decides whether or not to observe where candidates A and B stand. A voter cannot force another voter to acquire information.

2) Each voter votes sincerely, i.e. he votes for the candidate on whom he places a higher probability of implementing policy γ that he prefers. If he is indifferent or if he places equal probability on either candidate implementing his preferred policy, then he tosses a fair coin and votes accordingly.

3) The candidate who obtains the majority wins the election. In case of a tie, a coin toss determines the winner. The winner then implements the policy he prefers, and there is no possibility of reelection.

Now suppose that every voter prefers γ to be higher. In addition, every voter is also strictly doubt-prone. Let his value function be W_i^I if he acquires information and W_i^N if he does not. Even though every voter prefers the candidate who implements $\gamma = 1$, and even though information is free, there is still an equilibrium in which no one acquires information, and the candidate who implements $\gamma = 0$ wins with probability $\frac{1}{2}$. This equilibrium is Pareto-dominated (in expectation) by the other equilibria, in which at least a strict majority of agents acquires information, and the candidate who implements $\gamma = 1$ wins with probability 1. This is briefly shown below.

1) Equilibrium in which no voter is informed. If no other voter is informed, then voter i does not acquire information either. Since $p_i \in (0, 1)$ if no one else is informed, it follows that $W_i^I < W_i^N$ (on his own he cannot force $p_i \in \{0, 1\}$). Unless agent i is certain that either the right candidate or the wrong candidate always wins the election, i.e. that $p_i = 1$ or that $p_i = 0$, he does not acquire information.
Note that there is no equilibrium in which a minority of voters acquires information, since each voter in the minority has incentive to deviate. Note also that the difference between W_i^I and W_i^N for a given $p_i \in (0, 1)$ is higher if the difference between the agent's utility of $\gamma = 1$ and $\gamma = 0$ is larger.

2) Equilibrium in which at least a strict majority is informed. If at least a strict majority is informed, then the right candidate wins with probability 1. Hence $p_i = 1$ for each agent i, and so he is indifferent, since $W_i^I = W_i^N$. Note, however, that this equilibrium does not survive if each voter i places an arbitrarily small probability $\delta > 0$ that each of the other voters does not acquire information.

The externality of information plays an excessive role in this simple example, however it may still have an impact in a more realistic model. In particular, this example suggests that as the difference between the agent's utility of the good policy and his utility of the bad policy increases, a doubt-prone agent has *less* incentive to acquire information. In addition, a Pareto gain would be achieved if enough voters were 'forced' to acquire information on the candidates' policies.

1.5 Extensions and relation to the KP representation

In this section, I first analyze the relation between this model and the Kreps-Porteus (KP) representation (and, more generally, REU), and I show that the models are formally distinct, even if independence axioms hold at every stage. This last result may appear counterintuitive, since it may appear that a 'never' stage is formally equivalent to a 'much later' stage, but with a different interpretation. I discuss the reasons for the distinction between the two frameworks. The second part of this section presents a general methodology for extending other models to incorporate preferences over unresolved lotteries.

1.5.1 Relation to the KP representation

Suppose now, for simplicity, that there are 2 stages of resolution (early and late) in a KP setup. Assume, however, that the agent is indifferent between early and late resolution of uncertainty, so that there is a single utility function u associated with lotteries that resolve. It is clear that in this case, the KP representation is identical to an Expected Utility representation. But now, suppose that we include preferences over unresolved lotteries. That is, let $\mathfrak{L}_2$ is the set of simple lotteries over $\mathfrak{L}_1 \cup \mathfrak{L}_0$. For $\mathbf{X} \in \mathfrak{L}_2$, the notation $\mathbf{X} = (X_1, q^I_{1,e}; ...; X_{n_e}, q^I_{n_e,e};\ f_{1,e}, q^N_{1,e}; ...; f_{m_e,e}, q^N_{m_e,e}) \in \mathfrak{L}_2$,

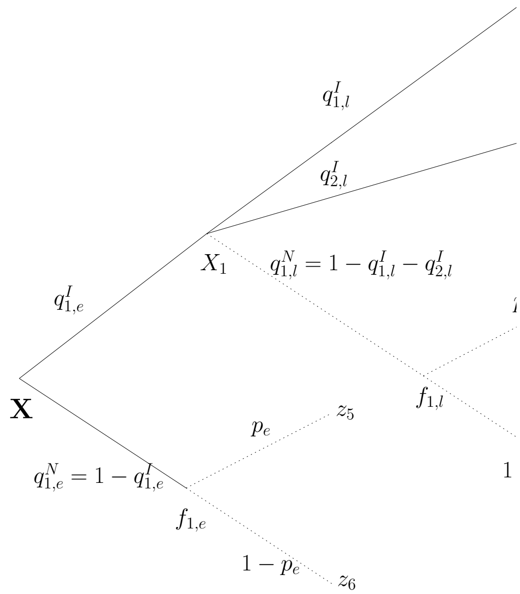

Figure 1.3: Lottery $\mathbf{X} = (X_1, q^I_{1,e}; f_{1,e}, q^N_{1,e} = 1 - q^I_{1,e})$.

where $X_{i,e} \in \mathfrak{L}_1$, and $f_{j,e} \in \mathfrak{L}_0$. The subscript 'e' denotes the early stage. The agent's preferences $\succeq$ are now over $\mathfrak{L}_2$, rather than $\mathfrak{L}_1$ (see figure 3).

The timing is as follows. The agent first observes the outcome of the first stage lottery (the early stage). For instance, with probability $q^I_{i,e}$, he receives a second lottery $X_i \in \mathfrak{L}_1$. The superscript I ('Informed') denotes that the agent expects to observe the outcome of lottery X_i. With probability $q^N_{j,e}$, the agent receives a lottery $f^N_{j,e} \in \mathfrak{L}_0$, which does *not* resolve. Here, the superscript N ('Not informed' denotes that the agent never observes the resolution of $f^N_{j,e}$. A lottery $f^N_{j,e}$ (henceforth 'early unresolved lottery') is a terminal node, in the sense that the agent does not expect it to lead to a second stage. Now suppose that the first (early) stage lottery leads to a second (late) stage lottery $X_i = (z_1, q^I_{1,l}; z_2, q^I_{2,l}; ...; z_n, q^I_{n_l,l};\ f^N_{1,l}, q^N_{1,l}; f^N_{2,l}, q^N_{2,l}; ...; f^N_{m,l}, q^N_{m,l})$. This second lottery always resolves. With probability $q^I_{h,l}$, the agent receives a final outcome $z^I_{h,l}$, which he observes. With probability $q^N_{k,l}$, he receives a lottery $f^N_{k,l} \in \mathfrak{L}_0$ which never resolves (henceforth 'late unresolved lottery'). The difference between a lottery f^N_e and a lottery f^N_l is that f^N_e knows after the early stage that f^N_e never resolves, while he knows after the late stage that f_l^N does not resolve. As before, the q^I's and q^N's are used to distinguish between the probabilities that lead to prizes where he is fully informed of the outcome (since he directly observes which z occurs), and the probabilities that lead to prizes where he is *not* informed (since he only observes the ensuing lottery).

Suppose now that an independence axiom for unresolved lottery holds at every stage. That is, define $\succeq_{N,e}$ and $\succeq_{N,l}$ in the natural way, and let an independence axiom hold for each of these preferences. In this case, there are unresolved utility functions v_e, v_l associated with $\succeq_{N,e}$ and $\succeq_{N,l}$, respectively:

$$\mathbf{W}(\mathbf{X}) = \sum q^I(z)u(z) + \sum q^I_{i,e}(z)\left(\sum q^N_{i,l}u\left(v_l^{-1}\left(Ev_l(z|f^N_{i,l})\right)\right)\right) + \sum q^N_{i,e}u\left(v_e^{-1}\left(Ev_e(z|f^N_{i,e})\right)\right)$$

Note that v_e and v_l need not be the same, since $\succeq_{N,e}$ and $\succeq_{N,l}$ are separate. Hence, there are three utility functions in this setting: utility u is associated with lotteries that eventually resolve, while functions v_e and v_l are associated with early and late unresolved lotteries. It is immediate, therefore, that having a KP model that accommodates unresolved lotteries is formally distinct from simply adding a 'never' stage, as this can only account for one additional utility function. The reason for this distinction is that the agent's perception of the unresolved lotteries need not be the same in the early stage as it is in the second stage.

There is another, and perhaps more fundamental, difference between temporal resolution and lack of resolution. While the early stage leads to the eventual occurrence of the late stage, there is no notion of sequence for unresolved lotteries. That is, the first unresolved lottery cannot lead to a second lottery; each unresolved lottery is a final prize, and hence a terminal node. For that reason, while the KP representation will have terms such as $u_e(u_l^{-1}(\cdot))$, there cannot be an equivalent unresolved term, $v_e(v_l^{-1}(\cdot))$. In this representation, both utility functions v_e and v_l are terminal, in the sense that the expectations are over outcomes, and not over any further lotteries. While the notation is cumbersome, this representation demonstrates that each unresolved lottery is essentially a final prize, and its value depends on whether it is obtained early or late. The agent's preferences over unresolved lotteries are allowed to vary in time, even when he has neutral preferences over the timing of resolution of uncertainty. The distinction between the KP representation and a representation that takes into account preferences for unresolved lotteries holds if the independence axioms over $\succeq_{N,e}$ and $\succeq_{N,l}$ are relaxed. In other words, this distinction carries through to more general REU representations.

1.5.2 General Methodology

This paper has extended the vNM EU model to allow for the distinction between lotteries that lead to observed outcomes and lotteries that never resolve, from the agent's viewpoint. I now present a simple methodology for extending other models to make this distinction as well. These models do not need to satisfy the general independence axiom **A.4**. I introduce another axiom instead. This axiom is weak enough to accommodate a broad class of continuous preferences, including a strict preference for randomization.

Suppose that an agent is indifferent between receiving an outcome $\tilde{z}$ as a final prize and an unresolved lottery f. It is now assumed that the agent is also indifferent between receiving unresolved lottery f and prize $\tilde{z}$ with the same probability. In other words, I assume that the agent's valuation, or perception, of unresolved lottery f is independent of the probability with which he receives it, and it is independent of the probability of receiving any other prize. The value placed on unresolved lottery f and the value placed on outcome $\tilde{z}$ are always the same.

AXIOM E.1 (Unresolved lottery equivalent): For all $f \in \mathfrak{L}_0$, $\tilde{z} \in \mathbf{Z}$ such that $\delta_f \sim \delta_{\tilde{z}}$, and for all $X, \tilde{X} \in \mathfrak{L}_1$ such that $X = (z_1, q_1^I; ...; z_n, q_n^I;\ f, q; f_2, q_2^N; ...; f_m, q_m^N)$

and $\tilde{X} = (z_1, q_1^I; ...; z_n, q_n^I; \tilde{z}, q;\ f_2, q_2^N; ...; f_m, q_m^N)$, the following holds: $X \sim \tilde{X}$.

Note, however, that the existence of a $\tilde{z}$ for which $\delta_f \sim \delta_{\tilde{z}}$ is at the moment not guaranteed. The following lemma presents conditions for which this is the case:

Lemma 1 (**Certainty equivalent**). *Suppose axioms **A.1** through **A.3** hold. In addition, suppose that $\succeq_N$ obeys stochastic dominance. Then there exists an H: $\mathfrak{L}_0 \rightarrow \mathbf{Z}$ such that for all $f \in \mathfrak{L}_0$, $\delta_{H(f)} \sim \delta_f$.*

That is, for any unresolved lottery δ_f, there exists a certainty equivalent $H(f)$ for which the agent is indifferent between receiving unresolved lottery δ_f and outcome $H(f)$ (or degenerate lottery $\delta_{H(f)}$) for sure. For any lottery f, therefore, $\tilde{z}$ in axiom **E.1** is equal to the certainty equivalent $H(f)$. Note that the main representation theorem in the paper makes no mention of axiom **E.1**; this is because it is trivially implied if the independence axiom **A.4** holds.

Lemma 2. *Suppose axioms **A.1** through **A.4** hold. Then axiom **E.1** holds.*

Without the independence axiom **A.4**, however, it is no longer the case that **E.1** necessarily holds. If it is explicitly assumed, if axioms **A.1** through **A.3** hold, and if $\succeq_N$ obeys stochastic dominance, then any lottery $X = (z_1, q_1^I; ...; z_n, q_n^I;\ f_1, q_1^N; ...; f_m, q_m^N) \in \mathfrak{L}_1$ can be replaced with a lottery $\hat{X} = (z_1, q_1^I; ...; z_n, q_n^I;\ H(f_1), q_1^N; ...; H(f_m), q_m^N) \in \mathfrak{L}_0$. Note that $X \sim \hat{X}$, by a repeated application of axiom **E.1**. This property essentially reduces two-stage lotteries to one-stage lotteries. It therefore allows a straightforward extension of different types of frameworks, so as to distinguish between resolved and unresolved lotteries. To emphasize this point, suppose that a 'simple model' is loosely defined as follows:

Definition (Simple Model) A simple model $\langle \hat{\succeq}, W, \mathcal{T} \rangle$ consists of :

- A preference relation $\hat{\succeq}$ over one-stage lotteries in $\mathfrak{L}_0$.
- A representation $W : \mathfrak{L}_0 \to \mathbb{R}$ for which $f \hat{\succeq} f' \Leftrightarrow W(f) \geq W(f')$ for all $f, f' \in \mathfrak{L}_0$.
- A set of axioms $\mathcal{T}$ that allow $\hat{\succeq}$ to be closed in the weak convergence topology, and that are sufficient for representation W to hold.

Then, any simple model can be expanded to accommodate the distinction between resolved and unresolved lotteries, in the following way. Take a simple model $\langle \hat{\succeq}, W, \mathcal{T} \rangle$. Since it is usually implicitly assumed that the agent will observe the outcome of a lottery, suppose that for all $f, f' \in \mathfrak{L}_0$, $f \hat{\succeq} f' \Leftrightarrow f \succeq f'$. That is, the set of axioms $\mathcal{T}$ is taken to hold for all resolved lotteries. If in addition, axioms **A.1** through **A.3** and axiom **E.1** hold, then $\succeq$ is represented as follows: for any $X, X' \in \mathfrak{L}_1$, $X \succeq X' \Leftrightarrow W(\hat{X}) \geq W(\hat{X}')$.[26] As for a representation of H, note that the set of axioms for unresolved lotteries considered in the paper can also be replaced by a second simple model $\langle \hat{\succeq}_N, W_N, \mathcal{T}_N \rangle$.

I now provide conditions for obtaining doubt-neutrality (indifference between observing and not observing the outcome) for preferences that satisfy **A.1** through **A.3** and stochastic dominance.[27] This simple result demonstrates that assuming doubt-neutrality has strong implications on the agent's allowable preferences, independently of the independence axiom **A.4**. Recall that for lotteries $f, f' \in \mathfrak{L}_0$, the notation $f \succ f'$ denotes a comparison between lotteries that the agent expects to observe; while $\delta_f \succ \delta_{f'}$ denotes a comparison between the same lotteries, but they remain unresolved.

[26] Where, as before, for $X = (z_1, q_1^I; ...; z_n, q_n^I;\ f_1, q_1^N; ...; f_m, q_m^N)$, $\hat{X} = (z_1, q_1^I; ...; z_n, q_n^I;\ H(f_1), q_1^N; ...; H(f_m), q_m^N) \in \mathfrak{L}_0$, and similarly for X' and $\hat{X}'$.

[27] See Segal (1990) for a similar result on time-neutrality in an REU setting.

Doubt-neutrality result. *Suppose axioms **A.1** through **A.3** hold. In addition, suppose that $\succeq_N$ obeys stochastic dominance. Then the following three conditions are equivalent:*

(i) $f \sim \delta_f$ *for all* $f \in \mathfrak{L}_o$

(ii) $f \succ f' \;\Rightarrow\; \delta_f \succ \delta_{f'}$ *for all* $f, f' \in \mathfrak{L}_o$

(iii) $\delta_f \succ \delta_{f'} \;\Rightarrow\; f \succ f'$ *for all* $f, f' \in \mathfrak{L}_o$

In words, suppose that an agent has a choice between observing and not observing the outcome of a lottery. Then he is always indifferent, for this type of choice, if and only if the order between any lotteries $f, f' \in \mathfrak{L}_o$ is always strictly preserved. That is, if he strictly prefers f to f' when he expects to observe the outcome, then he also strictly prefers f to f' if he does not expect to see the outcome. Arguably, condition (i) is often violated, even in models that depart significantly from the standard vNM model. Consider, for instance, the following variant of Machina's (1989) mother example. Suppose that a donor to a charity has no strict preference over which worthwhile cause receives the benefit from his donation, but he prefers that it be decided randomly, for reasons of fairness. He may still prefer not to observe which cause receives it, and to remain in doubt (and perhaps this encourages him to donate to an umbrella organization rather a more targeted one). It must therefore be the case that there are some lotteries f, f' over the recipients which he ranks differently based on whether he observes the outcome.

1.6 Closing remarks

This paper provides a representation theorem for preferences over lotteries whose outcomes may never be observed. The agent's perception of the unobserved outcome, relative to his risk-aversion, induces his attitude towards doubt. This relation is captured by his resolved utility function u, his unresolved utility function v and his unresolved decision weighting function w. The model presented here is an extension of the vNM framework, and it does not entail a significant axiomatic departure. However, it can accommodate behavioral patterns that are inconsistent with expected utility, and that have motivated a wide array of different frameworks. For instance, doubt-prone individuals have an incentive to self-handicap, and this incentive is higher if they are less certain about their competence.[28] Doubt-prone individuals are also more likely to choose the status quo bundle, if making a decision is more informative than inaction. In addition, an agent who is risk-neutral may still favor less risky investments, and would pay a firm to invest for him, even if it does not have superior expertise. The agent's attempt to preserve his self-image implies that his utility depends not only on the outcome that results, but also on the action taken. In a political economy context, doubt-proneness encourages political ignorance. When individuals derive more utility from the policies that they are not required to observe, they have *less* incentive to acquire information. Moreover, agents have a greater disutility from acquiring information if they are more ignorant ex-ante.

Finally, note that experiments that address the impact of anticipated regret frequently allow for foregone outcomes that individuals do not observe (see Zeelenberg (1999)). Similarly, in experiments by Dana, Weber and Kuang (2007), subjects deliberately choose to ignore free information concerning the full consequences of their actions. These empirical findings would be useful in determining plausible degrees

[28]Recall that this model does not allow agents to be delusional, since they are unable to mislead themselves into having false beliefs.

of doubt-proneness, although this is outside the scope of this paper.

Chapter 2

Information avoidance and the preservation of self-image

Abstract

There are a number of patterns of observed behavior that appear incompatible with the standard von Neumann-Morgenstern (vNM) Expected Utility model. For instance, behavior associated with anticipated regret, the Allais paradox and preferences for smaller menus are difficult to reconcile with the vNM framework. Evidence that individuals display a status quo bias has also motivated models that differ significantly from Expected Utility. In financial settings, ambiguity aversion has been used to accommodate a safe allocation bias, in which agents choose neither to buy nor short sell an asset for an *interval* of prices. Empirical findings that individuals choose to 'self-handicap' have been explained with notions of self-deception, cognitive dissonance and belief manipulation. Recently, experiments have been conducted in which dictators in dictator games who seem to exhibit preferences for fairness often switch to the 'selfish' choice if they can avoid observing the recipient's allocation.

I show that these seemingly unrelated findings can be accommodated by

a single, natural extension of the vNM model. On an intuitive level, the model is based on the assumption that decision makers wish to preserve their self-image and may be averse to obtaining signals of their self-worth. Agents obey standard axioms, and they are not allowed to manipulate their beliefs in any way nor to display any other form of self-deception. Instead, when choosing a course of action, agents take into consideration what the consequences of their actions will reveal to them about themselves. They cannot ignore bad signals and overweigh good signals, but their actions affect the *amount* of information they expect to receive concerning their self-worth. The agents' preference for controlling the flow of information suffices to accommodate diverse behavioral patterns.

2.1 Introduction

Individuals often let their choices depend on what they expect to learn, even when this information does not have a direct impact on them. Consider, for instance, Zeelenberg's (1999) Dutch lottery example. Subjects in two groups have the choice between buying a lottery ticket and keeping their money. A person in the 'no-feedback' group does not find out whether he would have won if he does not buy the ticket. In contrast, a person in the 'feedback' group observes the outcome of the lottery, independently of his choice. In this experimental setup, subjects in the feedback group buy more tickets, on average, than subjects in the no-feedback group. In another experiment, Zeelenberg and Beattie (1997) find that proposers in an ultimatum game offer significantly less money if they expect to be told the minimal acceptable offer afterwards. The results of these studies are difficult to reconcile with Expected Utility theory.[1]

In both examples, people arguably prefer *not* to observe the outcome of the

[1]More generally, models in which primitive preferences are taken over lotteries over received outcomes are inconsistent with these results.

unchosen lottery.[2] In the Dutch lottery case, individuals may not want to know that they would have won the risky lottery, had they played (Zeelenberg and Pieters (2004)). Similarly, proposers in the ultimatum game may prefer not to learn what the minimal acceptable offer would have been. Consider the dictator game experiment conducted by Dana, Weber and Kuang (2007). Dictators are unsure of the amount that they are actually giving the recipient, as it depends on a hidden lottery whose outcome they do not observe. They could, however, observe this outcome at no cost before making their decision, which allows them to choose the exact quantity given to the recipient. A significant percentage (44%) of dictators chooses to avoid acquiring information. On average, they keep a higher allocation for themselves and leave a smaller allocation to the recipient compared to when there is no uncertainty. Dana, Weber and Kuang conclude that "many subjects behave consistently with a desire to remain ignorant to the consequences of a self-interested choice." These results suggest that there is a connection between a preference to remain ignorant and the empirical findings described here. In this paper, I explore this relation and demonstrate that preferences to remain ignorant are consistent with these findings and other seemingly unrelated patterns of behavior.

It might appear that unobserved outcomes should be irrelevant to the agent's choice. He does not consume the reward that he does not receive, irrespective of whether it is high or low. But it is plausible that the agent is concerned with how his self-worth is affected, not with the foregone consumption per se. In the Dutch lottery example, individuals who realize that they have made the wrong decision presumably feel worse about themselves. In the dictator game, a dictator who believes that he ought to give a fair share to the recipient may experience diminished self-worth if he observes ex-post that the recipient has actually received a small amount. This argument does not require altruism or other-regarding preferences. The agent may be unconcerned with the recipient's utility; his only concern could be the link between

[2] See Northcraft and Ashford (1990), Zeelenberg, van den Bos, van Dijk and Pieters (2002). For a discussion on curiosity, see van Dijk and Zeelenberg (2007)

his self-worth and his code of conduct.[3] In Zeelenberg's ultimatum game, an agent who learns that he has overpaid feels foolish for having been 'ripped-off', and may lower his assessment of his decision making ability.

In principle, the standard von Neumann-Morgenstern Expected Utility theory (henceforth vNM) model is general enough that utility over self-worth can be included in the agent's value function. The vNM model places no restriction on what prizes the agent is allowed to value. There is, however, one fundamental difference between preferences over a consumption good and preferences over self-worth, which is that self-worth, unlike a consumption good, is never observed. Instead, individuals' self-image responds to the inferences they draw from the consequences of their actions. But the vNM model cannot accommodate an individual's preference to avoid a signal over self-worth, as it does not distinguish between the choice to observe and not to observe the outcome. Since there are numerous settings in which self-worth is a major factor, it is important to have a choice theoretic foundation that allows for these preferences. My main objective in this paper is to present such a framework, and to show that it does accommodate, and in fact predict, empirical findings that are inconsistent with the vNM framework, including the Dutch lottery example, the dictator game and the ultimatum game described here.

To characterize the agent's preferences over self-worth, I make use of the VUI (Value of Useless Information) model introduced in the first chapter of the dissertation. The VUI model allows decision makers to exhibit their preferences over observing and not observing the outcome. Agents may strictly prefer to remain ignorant and not observe an outcome (denoted doubt-proneness), or they may instead prefer to acquire information (doubt-aversion). I assume throughout this paper that agents are doubt-prone. They are not employing self-deception, in the sense that they cannot lie to themselves or manipulate their own beliefs in any way. The

[3]See Benabou and Tirole (2007) for a discussion on the connection between self-image and morality. As I later discuss, this model does allow for the interpretation that the agent has other-regarding preferences, provided he remains doubt-prone.

notions of regret and disappointment are not explicitly used in this setting either. Rather, these choices stem from the difference between the agent's expected utility over outcomes that he *eventually* observes and his assessment of lotteries that *never* resolve, from his frame of reference. Since the decision maker does not obtain ex-post utility from the true outcome if he does not know what that outcome is, he does not necessarily evaluate this lottery according to the expectation of a 'non-received' utility.[4]

I demonstrate that taking into account individuals' doubt-proneness over self-worth plays two roles. First, I show that a doubt-prone agent prefers having less choices in his menu and exhibits menu dependence. These results do no use notions of temptation or self-control. While models in which the agent prefers smaller menus typically allow these choices by directly taking preferences over menus as primitives, this framework follows the standard vNM model and only allows for preferences over lotteries. Second, this model can explain a large array of empirical puzzles. These well-known phenomena have motivated the development of models that differ significantly from the standard vNM model and from each other. In particular, the literature on anticipated regret might not appear related to the literature on the status quo bias, which itself seems disconnected from the literature on self-deception. But observed behavior that has motivated these separate modeling branches can be connected by a *single* root cause, namely, the agent's doubt-proneness:

(i) *Anticipated regret.* The Dutch lottery example is typically associated with anticipated regret. Intuitively, an agent who knows he would have won the lottery, if only he had played, regrets his decision. He is assumed to regret this decision more than the decision of buying the lottery and losing.[5] In this paper,

[4]For a full argument that does not appeal to this interpretation of expected utility, see chapter 1 of the dissertation.

[5]See Bell (1982), Loomes and Sugden (1982) and Sugden (1993). See Sarver (2008) for a model of anticipated regret with primitive preferences over menus.

this behavior is a consequence of the agent's doubt-proneness. The safe lottery is more informative of the agent's self-worth in the feedback case than in the no-feedback case, and is therefore less attractive. Moreover, this framework implies that a agent doubt-prone strictly prefers smaller menus over lotteries than larger menus, under some conditions. Each additional lottery in a menu provides the agent with more signals about himself, which he would rather avoid.

Note that this paper does not explicitly model regret. Rather, choices typically associated with regret are entirely due to the individual's preference for avoiding information about his decision making ability. The relationship between regret and self-image has been extensively studied in the psychology literature (see Larrick (1993) for a survey). While the economic models of anticipated regret do not explicitly model self-image, it is usually seen as an implicit factor. In this paper, preservation of self-image is the driving factor that leads to behavior associated with anticipated regret.

Consider the well-known variant of the Allais paradox known as the common ratio effect, which has also motivated models of anticipated regret (Loomes and Sugden (1982)). Given the choice between \$300 with probability 0.8 (and \$0 otherwise) and receiving \$200 for sure, many people prefer \$200. But when given the choice between \$300 with probability 0.4 and receiving \$200 with probability 0.5, a significant percentage of the same individuals choose \$300 with probability 0.4.[6] This preference reversal is a clear violation of the standard vNM model. In this framework, these choices are once again due to the agent's preference to avoid information concerning his self-worth. The agent expects the choice that leads to \$300 to be more informative of his self-worth when compared to the certainty of receiving \$200 than when compared to a 0.5 probability of receiving \$200.

[6]These numbers are drawn directly from Gul (1991), who uses a disappointment aversion explanation of the common ratio effect.

(ii) *Self-deception, cognitive dissonance and self-confidence.* Economic models of self-image generally assume a technology for belief manipulation or temporal inconsistency. For instance, Akerlof and Dickens (1982) use manipulable beliefs and the notion of cognitive dissonance to accommodate agents' preference to remain ignorant. Benabou and Tirole (2002, 2007), Bodner and Prelec (2003) and Benabou (2008) explore different settings in which agents self-signal and have access to belief manipulation.[7]

I follow the same view that self-image is relevant to agents, and that they have a mechanism to draw inferences about themselves. But I demonstrate that experimental results associated with self-deception, such as self-handicapping, hold in this setting as well, without having recourse to belief manipulation or temporal inconsistency. Models which take into account preferences for self-image can therefore use this framework as an axiomatic foundation.

(iii) *Status quo bias and the zero position bias.* The status quo bias refers to individuals' tendency to prefer their current endowment to other alternatives. It is often seen as being irreconcilable with the vNM model, as it does not allow a notion of 'frame of reference'.[8] In contrast, the status quo bias is consistent with the VUI framework, even though it does not explicitly model a frame of reference either. Rather, a doubt-prone agent has a tendency to maintain the status quo if doing so is less informative of his self-worth than other alternatives.

In settings in which keeping the status quo bundle is *not* the least informative, this framework makes predictions that are not necessarily in line with the endowment effect. This paper considers an application which distinguishes the

[7]See also Caplin and Leahy (2001), Compte and Postlewaite (2004) and Brunnermeier and Parker (2005).

[8]See Kahneman and Tversky (1979), Samuelson and Zeckhauser (1988) and Kahneman, Knetch and Thaler (1991).

least informative choice from the status quo in a financial market. In a well-known paper, Dow and Werlang (1992) use ambiguity (uncertainty) aversion to demonstrate that there is a price range at which agents neither buy nor sell an asset. This tendency does not imply that the agent is averse to trading, as in models with a status quo bias or transaction costs. The same result holds using this framework as well.

The economic environment in this paper is as follows. The agent's utility is linearly separable over his financial reward (money) and over his unobserved self-worth. He is not allowed to manipulate his beliefs in any way. That is, he cannot deceive himself into believing that he is better or worse than he is. Nevertheless, his decisions can affect how much information he receives concerning his self-worth. The agent first chooses a menu of lotteries over his financial reward, and subsequently chooses a lottery from within that menu. He expects to observe the outcome of every lottery within the menu he chooses. He does not, however, observe which outcome occurs for any other lottery. In other words, he receives feedback for every element within his chosen menu, and he does not receive feedback for any other element. He then uses the information he acquires as a signal concerning his self-worth, which he uses to update his beliefs in a Bayesian way.

As an illustration of the reasoning used throughout this paper, consider again the Dutch lottery example. In this framework, the agent believes the consequences of the lotteries, given his actions, are informative of his decision making ability. In the feedback case, the agent expects to receive a signal that affects his belief, since he observes the final outcome of both lotteries, regardless of his choice. In the no-feedback case, he obtains the same signal only if he buys the ticket; he does not acquire any information if he chooses the safe lottery. The safe lottery is therefore more informative in the feedback case than in the no-feedback case, while the risky lottery is exactly as informative in both cases. A doubt-prone agent essentially avoids the risky lottery in the no-feedback case because of the information it provides concerning his decision making ability.

An issue of concern is whether the agent's choices allow for his utility over his financial reward to be disentangled from his utility over self-worth. Self-worth is not only unobserved by the agent, it is also unobserved by the modeler. Ideally we would like the agent's choices to reveal his preferences over money. I show that this is indeed the case under some conditions. The utility function over money can be fully recovered and characterized from his choices. This framework also has clear testable implications that are consistent with the empirical findings provided in this discussion. For instance, consider the set $S = \{\{f, g\}, \{f, g, h\}\}$. Suppose that the agent chooses lottery f from menu $\{f, g\}$, and that he also chooses f from menu $\{f, g, h\}$. If given the choice between the two menus, the agent would be indifferent between in a vNM setting. In this framework, however, the doubt-prone agent *strictly* prefers the smaller menu $\{f, g\}$.

The paper is organized as follows. Section 2 presents the framework, and section 3 analyzes the main results. Section 4 focuses on the safe allocation bias, section 5 analyzes the revealed preference implications, section 6 applies this framework to the dictator game and the ultimatum game, and section 7 concludes. All proofs are in the appendix.

2.2 Model

I first introduce the setting. The agent makes a sequence of choices: he first chooses a menu over lotteries over his final reward, in this case money. Given his choice of menu, he knows that he will not only observe the outcome he receives, but also the outcome of every other lottery within his menu. I then characterize the agent's preferences over money and over his own fixed self-worth. The agent's utility over money is of the standard vNM form. Since the agent's self-worth is fixed, the vNM model predicts that it is of no relevance to his decision-making. But in this paper, I take into account that self-worth is unobserved. Furthermore, the agent is *doubt-prone*, meaning that he prefers not to acquire information concerning his self-worth. In the final part of this section, I describe the precise nature of the signal that the agent expects to receive concerning his self-worth.

2.2.1 Setting

Let the agent's final prize be a monetary reward r, drawn from the set $\mathcal{R} = [\underline{r}, \overline{r}] \subset \mathbb{R}$. Let $\mathfrak{L}_{\mathfrak{r}}$ be the set of simple lotteries over $\mathcal{R}$, with typical element $f \in \mathfrak{L}_{\mathfrak{r}}$, and let $2^{\mathfrak{L}_{\mathfrak{r}}}$, with typical element $M \in 2^{\mathfrak{L}_{\mathfrak{r}}}$, be the set of all menus over lotteries over $\mathcal{R}$. The timing of the agent's decision is as follows. He must first choose a menu, which is drawn from a set of menus $S = \{M_1, M_2, ..., M_n\} \subseteq 2^{\mathfrak{L}_{\mathfrak{r}}}$. Suppose that he chooses menu $M_i = \{f_1^i, f_2^i, ...f_{n_i}^i\}$. He then chooses a lottery f_j^i from within menu M_i.[9] The agent knows that he will observe the resulting outcomes $\{r_1^i, ..., r_{n_i}^i\}$ for every lottery within the menu M_i he receives, regardless of his choice. These are the only outcomes he observes; in particular, he does not observe the outcomes from the menus that he has not chosen. He has perfect recall, and so he remembers ex-post the set S from which he chose his menu M_i, and he remembers that he has chosen lottery f_j^i. This information is captured by his ex-post history $H = (r_1^i, ..., r_{n_i}^i; f_j^i; M_i; S)$.

As an example (table 2.1), let the agent first choose between menus from set $S = \{M_1, M_2\}$, where $M_1 = \{f_1^1, f_2^1, f_3^1\}$ and $M_2 = \{f_1^2, f_2^2\}$. Suppose that he chooses M_1, and subsequently chooses lottery f_1^1. He observes the outcomes of lottery f_1^1 (in this case \$100), but also the outcomes of lottery f_2^1 and f_3^1 (in this case \$250 and \$0, respectively). The agent remembers, ex-post, his choice of menu M_1 from S, his choice of lottery f_1^i, and every outcome that he has observed. This history is denoted $H = (\$100, \$250, \$0; f_1^1; M_1; S)$.

[9] Formally, let $C_S(S)$ be the correspondence that denotes the agent's choice of menu, and let $C_M(M_i|S)$ be the correspondence that denotes the agent's choice of lottery from menu M_i. The agent receives menu and lottery $\{M_i, f_j^i\}$, where $M_i \in C_S(S)$ and $f_j^i \in C_M(M_i|S)$. I assume for convenience that the menu and lottery he receives are drawn at random from his choices.

Initial Set	Chosen menu	Chosen lottery	Observed outcomes
$S = \{\{f_1^1, f_2^1, f_3^1\}, \{f_1^2, f_2^2\}\}$	$\{f_1^1, f_2^1, f_3^1\}$	f_1^1	$\{\$100, \$250, \$0\}$

Table 2.1: Sample sequence.

2.2.2 Preferences

In characterizing the agent's preferences, I take into account both his preferences over the consumption good (money) and his preferences over unobserved self-worth. When deciding what menu and lottery to choose, the agent considers not only his expected utility over money, but also how much information he expects to receive concerning his self-worth. I now describe the agent's preferences, and the next subsection describes the inferences he draws, given the consequences of his choices.

The agent has separable utility over the monetary reward r he receives and over his unobserved self-worth (or talent). The agent's utility over money is of the standard vNM form, and is characterized by utility function u_r. Given lottery over money f, the agent's expected utility is $Eu_r(f)$, where the expectation operator is defined in the usual way.[10] The agent's talent is in the set $[\bar{t}, \underline{t}]$.Assume, for simplicity, that the agent only places ex-ante probability $p_t \in (0,1)$ of having high talent $\bar{t}$ and probability $1 - p_t$ of having low talent $\underline{t}$. Under this assumption, the VUI representation for the agent's talent is captured by a function $u_{p_t} : [0,1] \to \mathbb{R}$. Given ex-post belief p of having high talent, his utility is $u_{p_t}(p)$, provided he never observes the resolution of uncertainty.

The function u_{p_t} can be interpreted as the agent's utility function over the *likelihood* that he is of high talent $\bar{t}$. The agent is never certain of being of high talent $\bar{t}$ or low talent $\underline{t}$, and so his final prize is not the actual outcome $\bar{t}$ or $\underline{t}$. Instead, his prize will be the ex-post probability p of being of type $\bar{t}$ (and $1-p$ of being $\underline{t}$). An agent who strictly prefers to avoid information (a doubt-prone agent) is effectively

[10]Specifically, $Eu_r(f) = \sum_{r \in f} p(r|f)u(r)$, where $p(r|f)$ is the probability of r occurring, given lottery f.

risk-averse over this likelihood of being of high talent. In other words, his utility function u_{p_t} is concave, which I assume throughout this discussion.

The standard vNM representation is a special case of this representation; it corresponds to the case in which function u_{p_t} is linear. If u_{p_t} were to be linear, then the agent would always be indifferent between receiving a more informative signal and a less informative signal. His self-worth would play no role in his decision making. I demonstrate in section 5 that under some conditions, the agent has preferences for smaller menus if u_{p_t} is concave, i.e. if he is doubt prone. Combining the agent's utility over money and his utility over self-worth, we obtain the following value function:

Value function

$$W(f^i_j, M_i|S) = Eu_r(f^i_j) \quad + \sum_{\{r^i_1,\ldots r^i_{n_i}\}\in M_i} u_{p_t}\left(p(t|H)\right) \prod_{h=1}^{n_i} p(r^i_h|f^i_h). \tag{2.1}$$

Hence, the agent chooses menu M_i and lottery f^i_j if:

$$\{M_i, f^i_j\} \in \underset{\{M_i\in S, f^i_j\in M_i\}}{\arg\max}\ W(f^i_j, M_i|S). \tag{2.2}$$

The first term of the value function, $Eu_r(f^i_j)$, is the standard expected utility over money. The second term is his expectation, over each possible history, of his utility over talent.[11] That is, suppose that the agent has chosen menu and lottery M_i and f^i_j, and that the the outcomes $\{r^i_1, ..., r^i_{n_i}\}$ occur. He therefore has history $H = (r^i_1, ..., r^i_{n_i}; f^i_j; M_i; S)$, from which he draws an inference concerning his talent, given his prior probability p_t of having high talent. His ex-post probability of having high talent is $p(t|H)$, and so his utility over his talent is $u_{p_t}\left(p(t|H)\right)$.

Recall that the agent having utility over talent is *not*, in itself, an addition to the standard vNM model. The only difference with the vNM model is that the agent has

[11]There is an abuse of notation in the value function, in that the r^i_j's are not elements of M_i, rather it is understood that r^i_j is an attainable outcome from lottery f^i_j in menu M_i.

preferences over whether to *observe* his talent ex-post. The vNM model corresponds to the case where the agent is indifferent, while the VUI model corresponds to the case where the agent is doubt-prone. Yet this difference is sufficient to be the driving factor for all the results that follow. Notice that the agent's primitive preferences do not directly depend on the menu he receives. Any type of menu dependence he exhibits is therefore a consequence of his preferences over money and self-worth.

The next step is to characterize the structure of the agent's signal over his talent. That is, structure must be imposed on $p(t|H)$, the agent's conditional probability of his talent, given the history he receives. For ease of exposition, I restrict attention to a specific signal structure that meets plausible assumptions. Allowing a broader class of signals would essentially yield the same results with little added intuition, while complicating the analysis.

2.2.3 Signals of self-worth

I focus here on the link between decision making ability (talent) and the agent's reward. This connection could be due to a higher ability at accomplishing a task, or it could be due to a better intuition over what constitutes a good decision. Both explanations are consistent with this model. But note that the agent only obtain signals, he has no mechanism for investing in his talent (self-improve) or manipulate his beliefs in any way. This assumption is formally written as follows:

Non-manipulability assumption For any set $S \in 2^{\mathfrak{L}_r}$, any menu $M_i \in S$ and any lottery $f_j^i \in M_i$, the following holds:

$$p_t = \sum_{\{r_1^i, \ldots r_{n_i}^i\} \in M_i} \left(p(t|H) \prod_{h=1}^{n_i} p(r_h^i | f_h^i) \right). \tag{2.3}$$

The non-manipulability assumption allows the agent's posterior beliefs to depend on his ex-post history, but it does not allow his expectation of his talent to be different from his prior. In other words, he expects to obtain a Bayesian signal of talent. Note that the agent makes no inferences from his choice itself. For instance, this agent cannot follow the reasoning "since I am not even willing to walk into a bar and talk to a stranger, I could not possibly be a great seducer."

For the rest of this section, I use as an illustration the set $S = \{\{f, g, h\}\}$, where $f = (\$100, 0.45; \$50, 0.1; \$0, 0.45)$, $g = (\$250, 0.2; \$0, 0.8)$ and $h =$ $(\$50.5, 0.99; \$0, 0.01)$. The agent must choose menu $\{f, g, h\}$, and then chooses between lotteries f, g and h. Suppose throughout the example that the agent chooses lottery f. In addition to non-manipulability, the signal over the agent's talent must satisfy the following properties:

Signal Properties

S.1 The agent views a better outcome from his chosen lottery as a higher signal of his ability. For instance, the agent's posterior probability of having high talent t_h is higher if he receives \$100 than if he receives \$50 (recall that he chooses lottery f), for given outcomes in g and in h. Furthermore, the agent obtains a positive (negative) signal if the outcome of his chosen lottery is higher (lower) than the outcome of an unchosen lottery. For example, suppose that the outcomes of f, g and h are \$50, \$0 and \$50.5. Then he obtains a positive signal comparing the \$50 he receives to \$0, but he receives a negative signal comparing \$50 to \$50.5.

S.2 If an unchosen lottery has a higher outcome, then it is a *lower* signal of the agent's ability. For example, if the outcome of g is \$250, it is a lower signal of his ability than if the outcome is \$0, for a given outcome in f and in h.

S.3 A likely history is less informative than an unlikely one, all else being equal.

Suppose the agent observes that the outcome of lottery h is \$50.5. This outcome is highly probable (probability 0.99) and therefore should not be very informative, compared (for instance) with obtaining \$50.5 from a lottery $h' = (\$50.5, 0.01; \$0, 0.99)$.[12]

S.4 Unchosen menus are less informative than they would be if they are chosen, since the agent never observes which outcomes in the unchosen menus occur.

These properties are sufficient for most of the results and intuition in the rest of this paper. I now focus on a specific signal structure, for which I define additional notation. Consider two lotteries, $f = (r_1, p_1; ...; r_n, p_n)$ and $f' = (r'_1, p'_1; ...; r'_{n'}, p'_{n'})$. Suppose that f does not dominate f', i.e. the best outcome in f' is better than the worst outcome in f. I assume that if a lottery strictly dominates another, then there is no signal to be received from comparing the outcomes. Therefore, there is no need to consider dominated lotteries, as they are never chosen. If the agent chooses lottery f, then he receives a positive signal the outcome r in f is higher than the outcome r' in f'. The *best* signal he receives is when he receives the highest r in f, and when r' in f' is the lowest. Define

$$\overline{p}^*(r_i, r_j | f; f') \equiv \sum_{\{r \in f, r' \in f'\}} p(r; f) p(r'; f') \mathbf{I}_{\{u_r(r_i) - u_r(r_j) \ > \ u_r(r) - u_r(r')\}}.$$

$$\underline{p}^*(r_i, r_j | f; f') \equiv \sum_{\{r \in f, r' \in f'\}} p(r; f) p(r'; f') \mathbf{I}_{\{u_r(r_i) - u_r(r_j) \ < \ u_r(r) - u_r(r')\}}.$$

The notation $\mathbf{I}_{\{\}}$ denotes the indicator function, which has value 1 if the statement in brackets is true and 0 if it is false. Given f and f' and outcomes r_i and r_j, respectively, $\overline{p}^*(r_i, r_j | f; f')$ denotes the probability of receiving a worse relative outcome. That is, $\overline{p}^*$ denotes the probability of receiving outcome r in f and r' in f' such that the difference in utility $u_r(r) - u_r(r')$ is smaller than the difference in utility

[12]Note, however, that a higher outcome is still a higher signal of ability, for any probability. For instance, the agent has a higher belief of his signal if he receives \$100 than if he receives \$50, even though \$50 occurs with much smaller probability.

$u_r(r_i) - u_r(r_j)$. Similarly, $\underline{p}^*$ denotes the probability of receiving a better relative outcome. I use $\overline{p}^*$ as a weight for positive signals; that is, if there is a high likelihood of getting a good outcome, then it is less informative of the agent's talent, in accordance with property S.3. Similarly, I use $\underline{p}^*$ as a weight for negative signals.

The signal drawn for unchosen menus is discussed in the appendix, as it is notationally cumbersome. The main assumption is that if the agent does not observe which outcome of a lottery occurs, then he does not take into account which menu that lottery is from. For instance, whether lottery g is drawn from menu M_j or M_k provides the same information if he receives neither menu M_j or M_k. For now, assume that the set S contains a singleton menu, i.e. $S = \{\{M\}\}$. The agent's ex-ante signal structure is as follows, given history $H = (r_1, ..., r_n; f_j; M; \{M\})$:

Signal structure

$$p(t|H) = p_t \Bigg(1 + \sum_{k \in \{1..n\}} b_{jk} \left(\frac{u_r(r_j) - u_r(r_k)}{u_r(\overline{r}) - u_r(\underline{r})} \right) \mathbf{I}_{\{u_r(r_j) \geq u_r(r_k)\}} (\overline{p}^*(r_j, r_k | f_j; f_k)) \\ - c_S \sum_{k' \in \{1..n\}} \left(\frac{u_r(r_{k'}) - u_r(r_j)}{u_r(\overline{r}) - u_r(\underline{r})} \right) \mathbf{I}_{\{u_r(r_j) < u_r(r_{k'})\}} (\underline{p}^*(r_j, r_{k'} | f_j; f_{k'})) \Bigg) \tag{2.4}$$

The values for parameters b_{jk} and c_S are chosen to ensure that this signal lies in the correct domain and satisfies the non-manipulability assumption (see appendix for details). The joint distribution $p(t, H) = p(t|H)p(H)$ is trivially characterized from the signal structure. Note also that the agent's actions do not provide him with any information in advance, as mentioned previously.

The signal structure is separated in two parts. The first part concerns the positive signal received by the observed outcomes in the chosen menu. It is relevant when the utility of the received outcome $u_r(r_j)$ is higher than the utility of the observed outcome $u_r(r_k)$. This signal is weighted by $\overline{p}^*$. Note that if the agent is highly likely to receive this relative outcome or better, then he acquires less information concerning his decision making ability. This corresponds to the signal property S.3. The second part concerns the negative signal received by the observed outcomes In this case, the signal is weighted by $\underline{p}^*$. Here as well, if the agent is highly likely to receive a bad relative outcome, then he acquires less negative information concerning his decision making ability. In addition, it is immediate that properties S.1 and S.2 are satisfied as well. The signal is positive if the outcome received r_j is higher than the observed outcome r_k, and it increases as $u_r(r_j) - u_r(r_k)$ increases. This completes the discussion of the signal structure, and we are now in position to analyze the results.

2.3 Implications of doubt-proneness

This section analyzes three behavioral patterns that are accommodated with this model, namely the Dutch lottery example, the common ration effect, and the the status quo bias. In most of the examples that follow, there exists a degenerate safe lottery $g = (r, 1)$. While this allows for easier exposition, none of these results depend on a g being degenerate. For instance, $g = (r + \epsilon, 0.5; 4 - \epsilon, 0.5)$ would yield the same results, for small enough ϵ.

2.3.1 Dutch lottery Example

Returning first to Zeelenberg's (1999) Dutch lottery example, the setup of the lotteries falls within this setting. That is, denote the risky lottery $f = (r_h, p; r_l, 1-p)$ and the safe lottery $g = (r, 1)$. The risky lottery f could lead to a high outcome r_h with probability p, or to the lower one r_l (the negative cost of the ticket) with probability $1 - p$. The feedback case, where the agent is forced to observe the resolution of the lottery, corresponds to the set $S = \{\{f, g\}\}$. The agent trivially chooses menu $\{f, g\}$, and then expects to observe the resolution of the risky lottery f, whether he chooses

Initial Set	Chosen menu	Chosen lottery
Feedback: $S = \{\{f, g\}\}$	$\{f, g\}$	f
No feedback: $S' = \{\{f\}, \{g\}\}$	$\{g\}$	g

Table 2.2: Dutch lottery example.

it or not. The no-feedback case, where the agent does not observe the resolution of f if he does not purchase the ticket, is written $S' = \{\{f\}, \{g\}\}$. Suppose that the agent chooses menu $\{f\}$ (and subsequently lottery f). If he chooses menu $\{g\}$ from set S', then he does not observe the outcome of the risky lottery.[13]

In the standard EU case, the agent who chooses lottery f (after choosing menu $\{f\}$) from the feedback set S should also choose lottery f from the no-feedback set S'. But this need not be the case in the VUI model. This framework allows the agent to choose g in the feedback case and f in the no-feedback case, as in table 2.2. Note that lottery g is more informative if it is drawn from set S than it is when drawn from S'. In comparison, lottery f is *exactly* as informative whether it is drawn from the set S or S', since in either case the agent knows that his prize would have been r if he had chosen lottery g. It is therefore possible for a doubt-prone agent to choose lottery g in the no-feedback case to switch to preferring risky lottery f in the feedback case. This allows him to avoid acquiring information about his decision making ability. Notice that it is clear here that a doubt-prone agent would *not* switch preferences in the other direction, i.e. from f in the no-feedback case to g in the feedback case. It may not be immediately clear, using only a general notion of regret, which of the decisions the agent anticipates regretting more; in this setting there is no such confusion, given his doubt attitude.

Consider the special case where $p = 0.5$ and $r = r_{CE}$, the certainty equivalent of lottery f. That is, r_{CE} is the value for which $u_r(r_{CE}) = 0.5u_r(r_h) + 0.5u_r(r_l)$. In this case, the doubt-prone agent is indifferent between receiving f and $g = (r_{CE}, 1)$ in the feedback case, as they are equally informative and have equal expected utility

[13] If the agent has the choice between observing and not observing the resolution of the risky lottery, then the set $\widehat{S} = \{\{f\}, \{g\}, \{f, g\}\}$.

valuations over money. However, he *strictly* prefers receiving lottery g (i.e. his certainty equivalent) in the no-feedback case than lottery f. This result is formally proven in the appendix.

2.3.2 Common ratio effect

Consider again the no-feedback case $S' = \{\{f\}, \{g\}\}$. Now let $r = \$200$, $p = 0.8$, $r_h = \$300$ and $r_l = 0$, so that $f = (\$300, 0.8; \$0, 0.2)$ and $g = (\$200, 1)$. Compare this to a second no-feedback case, $\tilde{S} = \{\{\tilde{f}\}, \{\tilde{g}\}\}$, where $\tilde{f} = (\$300, 0.4; \$0, 0.6)$ and $\tilde{g} = (\$200, 0.5; \$0, 0.5)$. In case S', lottery f is more informative than g. If the agent chooses f (after choosing menu $\{f\}$), he still knows the outcome of g would have been r. If instead he chooses g, he obtains exactly the outcome that he expects. Since he does not observe the resolution of f, he acquires no new information. A doubt-prone agent may therefore take this reasoning into account when choosing g. In the case $\tilde{S}$, it is less clear whether $\tilde{f}$ is more informative than $\tilde{g}$, since he does not observe the resolution of the other lottery in either case. In any case, the difference in informativeness between $\tilde{f}$ and $\tilde{g}$ in set $\tilde{S}$ is *smaller* than the difference in informativeness between f and g in S'. Therefore, the same doubt-prone agent who chooses g in S' to avoid information may now switch to $\tilde{f}$ in $\tilde{S}$ (table 2.3).

These numbers and choices correspond exactly to Gul's (1991) version of the Allais Paradox (also known as the common consequence or as the common ratio effect, see Kahneman and Tversky (1979)), which motivates his model of disappointment aversion. Note that these choices cannot be explained using the standard vNM model, in which the agent who chooses lottery g in S' also chooses lottery $\tilde{g}$ in $\tilde{S}$. Yet they are entirely consistent in this model, and may in fact appear very plausible. Notice that the set $S = \{\{f, g\}\}$ corresponds to a Dutch lottery, and that an individual may choose lottery f in S, even if he chooses g in S'. Whether this pattern holds for this specific example could be tested in an experimental setting.

Consider now the well-known criticism of the rationality of Allais preferences.

Initial Set	Chosen menu	Chosen lottery
No feedback: $S' = \{\{f\}, \{g\}\}$	$\{g\}$	g
No feedback: $\tilde{S} = \{\{\tilde{f}\}, \{\tilde{g}\})\}$	$\{\tilde{f}\}$	$\{\tilde{f}\}$

Table 2.3: Common Ratio Effect.

Suppose that a fair coin is tossed. Then:

(i) If the coin lands heads, the individual receives \$0. Otherwise, he receives lottery $f = (\$300, 0.5; \$0, 0.5)$, which then resolves.

(ii) If the same coin lands heads, then the individual still receives \$0. Otherwise, he receives $g = (\$200, 1)$.

Notice that the probabilities of each prize for choice (i) are the same as in lottery $\tilde{f}$, and similarly for choice (ii) and lottery $\tilde{g}$. It may seem that the Allais preferences indicate that the agent prefers choice (i) to (ii) ex-ante, but then switches to preferring choice (ii) once the coin lands heads.[14] But this is not the case in this model. The agent expects to observe whether the coin lands heads or tails regardless of his choice. This corresponds to set S', provided that he does not observe the final resolution of g in choice (ii). Therefore, the agent who prefers g in S' still prefers choice (ii) here, both ex-ante and ex-post, even if he chooses $\tilde{f}$ in set $\tilde{S}$.[15]

2.3.3 Status quo bias

This model accommodates the status quo bias using the same reasoning as for the explanation of the common ratio effect and the Dutch lottery. As an illustration, consider again lottery $f = (r_h, p; r_l, 1 - p)$, and let lottery $f' = (r'_h, p'; r'_l, 1 - p')$. Suppose that the expected utility over money is identical for f and f':

[14]See Segal (1990) for the counterargument that these preferences are not implied unless the reduction axiom holds. I present a different counterargument here, which is closely related to the regret explanation by Loomes and Sugden (1982) of the isolation effect in two-stage gambles. The intuition of their argument is essentially the same.

[15]The objects used in this section do not formally allow a correlation between lotteries' resolutions. The extension is straightforward, although the notation is cumbersome.

$$pu_r(r_h) + (1-p)u_r(r_l) = p'u_r(r'_h) + (1-p')u_r(r'_l).$$

Consider two cases:

(i) The agent receives lottery f, and before observing its outcome, is asked whether he would exchange it for lottery f'.

(ii) The agent receives lottery f', and is asked, before it resolves, whether he would exchange it for lottery f.

In the vNM framework, the agent is indifferent, in either case, between keeping his current endowment and switching. But there is strong empirical evidence that individuals often strictly prefer to keep their current allocation. This result is known as the status quo bias. It is consistent with the VUI model if the following assumption is made: the agent believes that he must observe the outcome of the lottery he receives. The weaker assumption that the agent is more likely to observe the outcome also suffices, but it is outside the scope of the present analysis. Given this assumption, the two cases can be written as follows:

(i) Set $S = \{\{f, f'\}, \{f\}\}$. Keeping the current allocation corresponds to choosing menu $\{f\}$ and lottery f. Switching lotteries corresponds to choosing menu $\{f, f'\}$ and receiving lottery f'.

(ii) Set $\tilde{S} = \{\{f, f'\}, \{f'\}\}$. Keeping the current allocation corresponds to choosing menu $\{f'\}$ and lottery f. Switching lotteries corresponds to choosing menu $\{f, f'\}$ and receiving lottery f.

It is clear that in both cases, a doubt-prone agent *strictly* prefers to keep his current endowment. The two lotteries yield the same expected utility over money, while switching lotteries provides more information on his decision making ability,

Initial Set	Chosen menu	Chosen lottery
Feedback: $S = \{\{f, f'\}, \{f\}\}$	$\{f\}$	f
No feedback: $\tilde{S} = \{\{f, f';\}, \{f'\}\}$	$\{f'\}$	f'

Table 2.4: Status quo bias.

which he would rather avoid (table 2.4). Note that once again, this result is not only consistent with the VUI model, it is in fact expected.

While the reasoning used here is consistent with the status quo bias for this specific case, there are other settings in which the predictions of this model depart from the implications of the status quo bias. I present an example of this divergence in the next section.

2.4 Safe allocation bias

This section places the the agent in a financial setting in which he chooses whether to buy an asset, to sell it short or to hold a zero position. I show that there is a bias towards the (safe) zero position. I specifically use the example analyzed by Dow and Werlang (1992), as it fits naturally to this setting. Suppose that there is one unit of an asset, at price P. I assume that the agent places probability q that the value of the asset will be r_h and probability $1-q$ that the value of the asset will be r_l.

If the agent chooses to buy the asset, then he receives $r_h - P$ (the good outcome) with probability q and $r_l - P$ with probability $1-q$. This corresponds to lottery $f_b = (r_h - P, q; r_l - P, 1-q)$. If instead he chooses to sell the asset short, then he receives $P - r_h$ (in this case, the bad outcome) with probability q and $P - r_l$ with probability $1-q$. This corresponds to lottery $f_s = (P - r_h, q; P - r_l, 1-q)$. Holding no position corresponds to lottery $g = (0, 1)$. Suppose that there exists some price P for which $r_h - P > 0 > r_l - P$.

As noted by Dow and Werlang, the standard vNM model with (local) risk-neutrality predicts that there is a unique P^* such that the agent buys the asset (chooses f_b) if $P < P^*$, sells it short f_s if $P > P^*$, and takes no position (g) (or is indifferent) otherwise. With a notion of ambiguity aversion, there is instead a price *interval* $[P_B, P_G]$ at which the agent chooses the zero position. He buys the asset for price $P < P_B$, and sells the asset for price $P > P_G$. The key behind Dow and Werlang's result is that under ambiguity aversion, the probabilities add up to less than 1.[16] That is, instead of q and $1 - q$, the agent effectively places probability q_h and q_l, with $q_h + q_l < 1$. Note that his result is different from what would be obtained with the status quo bias or a transaction cost: in those cases, the agent would be biased towards his current portfolio, which need not be the zero position.

This framework yields the same result as Dow and Werlang's, but here the probabilities add up to 1. That is, the agent places probability q that the value of the asset will be r_h and probability $1 - q$ that the value of the asset will be r_l. The reasoning is the same as in the previous sections; the agent does not observe the outcome of the asset unless he chooses to buy it or sell it short. The agent's choice corresponds to set $S = \{\{f_b, f_s\}, g\}$. Note that if he buys (sells) the the asset (f_b), he is certain what he would have received had he sold (bought) the asset (f_g), and he is also certain of what he would have received if he had kept the safe allocation (g).[17] In contrast, if he chooses g, he does not observe the resolution of f_b or f_s, and so obtains less information on his decision making ability. Therefore, there is a price interval $[P_B, P_G]$ at which the agent prefers g to both f_b and f_s. For a low enough price $P < P_B$, he prefers f_b, and for a high enough price $P > P_G$, he prefers f_s. This result is formally proven in the appendix.

Note that this result holds even though the agent is risk-neutral. We can take the zero position g to mean a safe portfolio of bonds, and f_b and f_s to be risky stocks. It is clear, by the same reasoning, that although the agent is risk-neutral, there is still an equity premium.

[16] See, for instance, Schmeidler (1989) and Gilboa and Schmeidler (1989).

[17] Technically, choice S is not identical to his choice, since the two lotteries are correlated; there is no history for which outcome $r_h - P$ from f_b and outcome $r_l - P$ from $f - s$ can occur. This framework can be adapted in a straightforward fashion to allow for this correlation.

2.4.1 Relation to ambiguity

It is perhaps surprising that there is overlap between the predictions of this framework and those obtained by taking into account ambiguity aversion. My aim here is to elucidate this connection by focusing on a variation of Ellsberg's paradox. The argument is not formal, and only relies on similarities.

In an experiment by Halevy (2007), there are four different urns, $U1$, $U2$, $U3$ and $U4$, each with 10 balls:

(i) Urn $U1$ has 5 red balls and 5 black balls.

(ii) The distribution of red and black balls in urn $U2$ is unknown.

(iii) The number of red balls in $U3$ is uniformly distributed between 0 and 10.

(iv) There are 10 red balls with probability 0.5 and 10 black balls with probability 0.5.

Suppose now that the agent is presented with an urn from which one ball is drawn, and must choose a color. If a ball of that color is drawn, he receives $2. He is asked which urn he prefers. The only urn in which there is 'ambiguity' is urn $U2$, since the probabilities are objective in lotteries $U1$, $U3$ and $U4$. Ignoring urn $U2$, the agent, according to the standard vNM model, should be indifferent between urns $U1$, $U3$ and $U4$. Instead, Halevy (2007) finds that two patterns emerge, with even frequency: the first group is indifferent between urns $U1$ and $U4$, but ranks urn $U3$ as worse. The second group prefers urn $U1$ to $U3$ to $U4$. Both groups act in a way that is inconsistent with the vNM model; I focus here only on the second group.

There is perhaps a 'flavor' of ambiguity in urn $U4$, which motivates why the agent prefers $U1$ to $U4$, even without ambiguity. But the interpretation is difficult from a rigorous ambiguity viewpoint, since probabilities are objective both for urn $U1$ and urn $U4$ (it is consistent, however, with a violation of the reduction axiom, see Segal (1990)). This framework can account for the second group's preferences, but not the first. It depends, however, on the exact setup of the experiment. Consider the following two alternatives:

1. All the lotteries have taken place, but the agent does not know the results yet.

2. The number of balls in urns $U3$ and $U4$ has already been determined, but has not been shown to the agent. For urn 3, if the agent chooses the color black, the second draw (the draw of the ball from the urn) occurs after 3 seconds. If he chooses the color red, then the draw occurs after 4 seconds. In other words, a different randomizing device is used, and these draws are uncorrelated.

3. None the lotteries have taken place, and a different randomizing device is used for each draw and for each stage.

For case 1, this framework makes the same prediction as the vNM model that the agent is indifferent between urns $U1$, $U3$ and $U4$. Since this framework does not distinguish between one-stage lotteries and compound lotteries, all three urns corresponds to set $S_1 = \{\{f, f'\}\}$, where $f = f' = \{\$2, 0.5; \$0, 0.5\}$. For any of these lotteries, if the agent observes black and loses (wins), he knows that he would have won (lost) had he chosen red.[18] The same holds for case 3, in which, the agent has no information on whether he *would have* won; the urns now correspond to set $S_3 = \{\{f\}, \{f'\}\}$. Observing that he drew a black ball does not imply that he should have chosen red, since he has no information on whether red *would* have occurred had he chosen it. This situation is akin to having a choice between two roulette wheels; observing the black ball as the outcome of one roulette wheel does not imply that red would have occurred in the other.

This framework, however, does *not* predict that the agent is indifferent between urns $U1$, $U3$ and $U4$ in case 2. Instead, the agent strictly prefers $U1$ to $U4$ to $U3$,

[18] Some histories cannot occur, such as history $H = (\$20, \$20; f; \{f, f'\}; S)$, because of the correlation between lotteries. As mentioned previously, while this framework does not account for a correlation between lotteries, it can be adapted in a straightforward manner. Similarly, lottery f and f', strictly speaking, should not be identical; this can also be incorporated in the model. Alternatively, we could allow f' to vary slightly from f.

which matches exactly the second group in Halevy's experimental findings. The reason is the following: urn $U1$ corresponds to set $S_3 = \{\{f\}, \{f'\}\}$, as in case 3. But urn $U4$ corresponds to set $S_1 = \{\{f, f'\}\}$, as in case 1. The agent prefers the least possible information on his decision making ability, and therefore prefers urn $U1$. Formally, given set $S = \{\{f, f'\}, \{\{f\}, \{f'\}\}$, the agent prefers menu $\{f\}$ (or $\{f'\}$).[19]

2.5 Revealed Preferences

This section analyzes the strong connection between the agent's preferences for smaller menus and his doubt-proneness are strongly connected. I then show that the agent's utility over money can be recovered from the agents choices. My aim in presenting these results is twofold. First, these results provide testable implications for this model. Second, the ability to extract the agent's utility over money from utility over self-worth demonstrates that these two dimensions of preference are distinct and empirically separable from each other.

Consider the sets $S = \{\{f, g\}\}$, $S' = \{\{f, g, h\}\}$ and $S'' = \{\{f, g\}, \{f, g, h\}\}$. Given set S, suppose that the agent chooses $\{\{f, g\}, f\}$. That is, he trivially chooses menu $\{f, g\}$, and subsequently chooses lottery f. Given set S', he chooses $\{\{f, g, h\}, \{f\}\}$. In the vNM case (and in fact, in most models where the agent's primitive preferences are over lotteries over outcomes), the agent also chooses lottery f in set S''. He is indifferent between receiving menus $\{f, g\}$ and $\{f, g, h\}$, since his choice of lottery is the same in both cases.

In this model, he is no longer indifferent between the two menus. Given set S'', a doubt-prone agent chooses $\{\{f, g\}, f\}$. He strictly prefers the smaller menu $\{f, g\}$, even though he makes the same choice of lottery f.[20] Intuitively, he obtains

[19] See Loomes and Sugden (1982) for a discussion of the correlation between states of the world in a model of regret.

[20] This reasoning assumes h is not a degenerate lottery. If h is a degenerate lottery, then he is indifferent between receiving menus $\{f, g\}$ and $\{f, g, h\}$, since he is certain of the outcome of h in

more information about his self-image if there are more lotteries in his menu. He still obtains information from the unchosen menu, but less than he does from the outcomes he observes, as stated in signal property S.4. This result is formalized in theorem 1. Let $C_S(S)$ denote the choice of menu from S, and let $C_M(M|S)$ denote the choice of lottery from menu M, given initial set S.

Theorem 2.1: Preference for smaller menus. *Suppose that the agent is strictly doubt-prone. Take any* $S = \{M_1, M_2, ...M_n\}$, *where* $M_1 \subset M_2 ... \subset M_n$, *and where no lottery is degenerate. In addition,* $C_M(M_1|\{M_i\}) = C_M(M_2|\{M_j\}) = ... = C_M(M_n|\{M_n\})$. *Then the agent strictly prefers the smallest menu* M_1 *in* S, *i.e.* $C_S(S) = M_1$. *Furthermore, he chooses lottery* $C_M(M_1|S) = C_M(M_1|\{M_i\})$.

The reason for the condition that no lottery is degenerate is that degenerate lotteries are equally informative whether they are chosen or unchosen. The next theorem answers the concern over whether doubt-proneness over self-image can be separated from preferences over the financial reward. Theorem 2 demonstrates that the agent's choices allow us to precisely characterize his utility function u_r.[21]

Theorem 2.2: Charactizerization of utility over money. *Function* u_r *can be uniquely characterized from the agent's choices, up to positive affine transformation.*

The precise mechanism for eliciting function u_r is described in the appendix. An implication of theorem 2 is that we can isolate the effect of the agent's doubt-proneness, even though the object over which he is doubt-prone, talent, is never observed. That is, we can study the exact tradeoff between the agent's expected utility over money and his doubt-proneness, if we have collect enough data on his decisions.

---

either case.

[21]See the appendix for a more rigorous version of theorem 2.

2.6 Games with uncertainty

In this section, I revisit the two games mentioned in the introduction, namely Zeelenberg and Beattie's (1997) ultimatum game and Dana, Weber and Kuang's (2007) dictator game. Neither of these settings is actually a game: the amount that the recipient in Zeelenberg and Beattie's game is predetermined, and the recipient does not play a role in the dictator game. The analysis of the two examples is straightforward.

There are two separate cases in the ultimatum game: in case (i), the player making an offer will be told ex-post the exact amount that the recipient would have accepted. In case (ii), the player is not told the amount. Zeelenberg and Beattie (1997) find that players in case (i) make a less generous offer, on average, than players in case (ii). The reasoning used throughout previously in this paper applies here as well: there is an asymmetry between the informativeness of each choice in case (i) compared to case (ii). Suppose, for instance, that the agent can split \$10 dollars in two ways: he can either (a) keep \$8 dollars for himself and \$2 for the recipient or (b) split it equally, and they each receive \$5 or he can . He places probability 0.5 that the recipient accepts \$2 dollars, but he is certain that the recipient always accepts \$5\$. Suppose that the dictator is not altruistic, that is, his utility does not depend on the recipient's monetary reward. Then this corresponds exactly to the Dutch lottery example: in case (i), the agent's set is $S = \{\{f, g\}\}$, where $f = (\$8, 0.5; 0, 0.2)$ and $g = (\$5, 1)$, and in case (ii), the set is $S' = \{\{f\}, \{g\}\}$ (table 2.2). Once again, the reason that an agent may choose f in case S but switch to g in case S' is not that he anticipates 'regretting' his decision in set S if he picks g and realizes f would have been acceptable. Rather, he avoids f in set S' because of how informative it is, since he knows what he would have received, had he chosen g. In set S, however, he observes the resolutions of both f and g, and so the difference in informativeness between the two lotteries is smaller.

Dana, Weber and Kuang's (2007) dictator game, however, requires an adaption of this framework. In a dictator game, the recipient does not have the possibility to

Box A	Dictator: 6 Recipient: 1
Box B	Dictator: 5 Recipient: 5

Table 2.5: Baseline Case.

	Heads	Tails
Box A	Dictator: 6 Recipient: 1	Dictator: 6 Recipient: 5
Box B	Dictator: 5 Recipient: 5	Dictator: 5 Recipient 1

Table 2.6: Baseline Case.

refuse the offer. The dictator must therefore have a reason to share the endowment, otherwise he would not take into account what the recipient receives. Consider the setup used by Dana, Weber and Kuang. The dictator has a choice between two options, A and B. In case (i), the 'baseline' case, (table 2.5), there is no uncertainty; if he chooses option A then he receives \$6 and the recipient receives \$1. If instead he chooses option B, then both he and the recipient receive \$5. In case (ii), the 'hidden information' case (table 2.6), he still receives \$6 he chooses A, and he still receives \$5 if he chooses B. But now, the recipient's allocation is determined by a coin toss, prior to the dictator's choice. If it lands heads, then the recipient's allocation is \$1 for option A and \$5 in option B, and if it lands tails then it is the other way around. The dictator has the choice, before making his decision, to observe the coin toss at no cost. Otherwise, he never observes what the recipient receives. As for the recipient, he does not find out whether or not the dictator has seen the outcome of the coin toss.

A significant number of dictators choose not to observe the outcome of the coin toss in the hidden information case, and are more likely to choose option A ; on average more agents choose option A in this setup then in the baseline case (thereby giving the recipient a smaller share, on average). Notice that this behavior is inconsistent with both altruism and self-interested behavior: if the agents are indeed self-interested, then they should always choose option A, even in the case without uncertainty. If they are altruistic, then they should strictly prefer to observe the coin toss before making the decision.

If the agent's self-image is fixed, as I have assumed throughout this paper, then in fact it will not be affected by whether he makes the generous offer for certain ($5 for the recipient) or the selfish one ($1 for himself). I instead assume that the agent believes that there is a moral action and an immoral action, and that the action *itself* has a benefit or a cost. That is, his self-image is higher if he acts morally (giving the recipient $5), and lower if he acts immorally (giving the recipient $1. This is indistinguishable from assuming that he prefers the recipient to receive $5 or $1. Suppose that giving the recipient $5 raises his self-image to $\bar{t}$ for certain, and giving him $1 lowers it to $\underline{t}$ for certain. Recall that utility over self-image is denoted $u_{p_t}(p)$, where p is the probability of having self-image $\bar{t}$.

Consider first a dictator weakly prefers option A. Then $u_r(6) + u_{p_t}(1) \geq u_r(5) + u_{p_t}(0)$. Equivalently, the difference in his utility over money between options A and B is greater than the difference between his utility over self-image: $u_r(6) - u_r(5) \geq w(1) - u_{p_t}(0)$. In case (ii), in which there is uncertainty, if this dictator chooses to observe the outcome of the coin toss before making his decision, he also still chooses option A. If instead, he chooses *not* to observe the outcome of the coin toss, then he still chooses option A, since he obtains more for himself, and the recipient receives the same on average as in option B. He therefore prefers to observe the outcome of the coin toss if:

$$u_r(6) + 0.5\,(u_{p_t}(1) + u_{p_t}(0)) > u_r(6) + u_{p_t}(0.5). \tag{2.5}$$

But since w is strictly concave, this conditions never holds. He therefore strictly prefers not to observe the coin toss, and to choose A for certain (denote this agent 'type AA').

Now consider a dictator who strictly prefers option B in case (i). Then $u_r(6) -$

$u_r(5) < u_{p_t}(1) - u_{p_t}(0)$. In case (ii), if he chooses to observe the outcome of the coin toss, then his expected utility is: $0.5\,(u_r(6) + u_r(5)) + u_{p_t}(1)$. If he chooses not to observe the outcome, then he prefers option A, and his utility is $u_r(6) + u_{p_t}(0.5)$. This dictator chooses not to observe the outcome of the coin toss if:

$$
\begin{aligned}
&u_r(6) + u_{p_t}(0.5) > 0.5\,(u_r(6) + u_r(5)) + u_{p_t}(1) \\
&\Longrightarrow 0.5(u_r(6) - u_r(0.5)) > u_{p_t}(1) - u_{p_t}(0.5).
\end{aligned}
\tag{2.6}
$$

Together, these conditions imply that

$$
0.5\,(u_{p_t}(1) + u_{p_t}(0)) < u_{p_t}(0.5). \tag{2.7}
$$

Since the agent is doubt-prone (u_{p_t} is concave), this condition can hold (I denote those for whom it holds type BA, and those for whom it does not type BB). Note that it could *not* hold if the agent were doubt-neutral or doubt-averse. Summarizing, there are three possible types of dictators:

- Type AA: Dictators who choose A in both case (i) and case (ii). They must also prefer not to observe the outcome of the coin toss.

- Type BA: Dictators who choose B in case (i), avoid observing the outcome of the coin toss, and choose A in case (ii).

- Type B: Dictators who choose B in case (i), observe the outcome of the coin toss, and choose whichever option provides the recipient with $5.

This framework does *not* allow dictators who choose to observe the outcome of the coin toss and still choose the option for which the recipient receives $1 . It also does not allow dictators not to observe the outcome and then choose option B.

The predictions of this model seem to fit the data well (table 2 in Dana, Weber and Kuang (2007)): of the 32 dictators, 14 (44%) chose not to observe the outcome of the coin toss. Of those 14 individuals, 12 (86%) chose option A (consistent with types I and II).[22] Of the 18 individuals who chose to observe the coin toss, 15 (83%) chose the option for which the recipient receives \$5 (consistent with type III).[23] Note that types I and II are conflated in case (ii), and that we expect more individuals (those of type II) to choose the fair allocation in case (i). This pattern also emerges: in case (i), 14 out of 19 (74%) choose B, compared to 6 out of 16 (38%) who choose the fair option in case (ii).[24]

2.7 Closing remarks

I have shown in this paper that taking into account preferences for avoiding information relevant to self-worth can accommodate a number of empirical findings. In addition, this framework serves as a link between different branches of the literature, namely models of anticipated regret, models of self-image and models in which the agent has a bias towards a reference point. The agent's objective is to preserve his self-image; that is, he wishes to acquire the least amount of information concerning his decision making ability. In this sense, the agent's bias is towards remaining as close to his prior as possible. The same reasoning can be used in Dana, Weber and Kuang's (2007) dictator game with uncertainty. Their empirical findings seem at odds with both self-interested preferences and altruistic preferences, but the predictions of this framework appear to fit the data well.

I conclude with the observation that individuals sometimes *seek* information over self-image, rather than avoid it. For example, part of the appeal of gossip, arguably, is to compare oneself with others. Sports, games and other forms of competition serve as forums for obtaining a more precise signal of one's ability in different fields. I have

[22] Note that the 2 individuals who chose option B are receiving less for themselves, and are leaving less for the recipient, on average, than if they observed the outcome.

[23] One agent chose a smaller allocation both for the recipient and for himself.

[24] It is also the case that 13 out of 16 (81%) choose A when the dictator receives \$6 and the recipient receives \$5, but here the self-interested choice cannot be disentangled from the moral choice.

ignored the notion of doubt-aversion (or curiosity) in this paper, but an extension of this model could allow agents' attitude towards doubt to vary. In particular, we could consider an economy in which agents have the same prior self-image but heterogeneous degrees of curiosity. Intuitively, a higher degree in curiosity may confer an advantage to otherwise identical agents.

Appendix A

Appendix A

The appendix is structured as follows. Part 1 explains why the standard EU model is inappropriate when the agent does not expect to observe the resolution of uncertainty. Part 2 provides an example of the 'preservation of self-image' application. All the proofs are in part 3.

A.1 Limitations of the standard EU model

This example illustrates the problem with using the standard vNM EU model when there are outcomes that the agent never expects to observe. Consider the simple case of an agent who has performed a task and does not know how well he has done. There are no future decisions that depend on his performance. For example, as a simple adaptation of Savage's omelet, suppose that the agent does not know whether he has fed his guests a good omelet or a bad one. With probability p_t, he has done well ($\bar{t}$), and with probability $(1 - p_t)$ he has done badly ($\underline{t}$). He prefers having done well to having done badly, although this will have no future repercussions. Given the choice between remaining forever in doubt (D) and perfectly resolving the uncertainty, (ND), it might appear that he compares:

$$U_D = p_t u(\bar{t}) + (1 - p_t) u(\underline{t})$$

to

$$U_{ND} = p_t u(\bar{t}) + (1 - p_t)u(\underline{t})$$

and that since $U_D = U_{ND}$, he is indifferent. But U_D is *not* necessarily the right function to use if he chooses to remain in doubt, because from his frame of reference the final outcome will not be $\bar{t}$ or $\underline{t}$. That is, he does not expect to 'obtain' ex-post utility $u(\bar{t})$ or $u(\underline{t})$ because he does not expect to observe either $\bar{t}$ or $\underline{t}$. As it is not clear what his perception of the consequence is if he does not expect the uncertainty to be resolved (from his viewpoint), his expected utility is undetermined. In its current form, the standard EU model does not offer a method for evaluating this choice. Using U_D effectively ignores that the relevant frame of reference is the agent's, not the modeler's.[1]

Redefining the outcome space to include the observation itself does not eliminate the problem. Suppose that the outcome space is taken to be $Z = \{\bar{t}_D, \underline{t}_D, \bar{t}_{ND}, \underline{t}_{ND}\}$ where $\bar{t}_D$ represents the outcome that he did well but doubts it, $\bar{t}_{ND}$ that he did well and does not doubt it, and so forth. He therefore compares the following:

$$U_D = p_t u(\bar{t}_D) + (1 - p_t)u(\underline{t}_D)$$

to

$$U_{ND} = p_t u(\bar{t}_{ND}) + (1 - p_t)u(\underline{t}_{ND})$$

It is difficult to interpret the meaning of the consequence 'did well, but doubts it' from his frame of reference, since it is not clear what it means to be in doubt if he knows that he has done well. In addition, his preferences over $\bar{t}_D$ and $\underline{t}_D$ are completely pinned down. Consider the two extremes, $p_t = 1$ and $p_t = 0$. When $p_t = 1$, there is no intrinsic difference between U_D and U_{ND}, since he knows that he has done well. Hence, $u(\bar{t}_D) = u(\bar{t}_{ND})$. Similarly, when $p_t = 0$, he knows he has done badly, and so $u(\underline{t}_D) = u(\underline{t}_{ND})$. It then follows that $U_D = U_{ND}$ for any $p_t \in [0, 1]$. This definition of the outcome space is essentially the same as simply $Z = \{\bar{t}, \underline{t}\}$. His indifference between remaining in doubt and not remaining

[1] This issue is not resolved by starting with preferences over lotteries as primitives. In the standard framework, the agent has primitive preferences over lotteries over outcomes, and he is not allowed to choose between lotteries whose resolution he observes and lotteries whose resolution he does not observe. He is therefore not given the option to express those preferences.

in doubt is a consequence of following this approach, it is not implicit from the standard EU model.

Redefining the outcome space so that his utility is constant if he remains in doubt is even more problematic. Suppose that $Z = \{\bar{t}_{ND}, \underline{t}_{ND}, D\}$, letting $\bar{t}_{ND}$ be the outcome 'talented and he does not remain in doubt (he observes the outcome)', $\underline{T}_{ND}$ be the outcome 'untalented and he observes it', and letting D mean that he does not observe the outcome, hence remaining in doubt. He now compares:

$$U_D = u(D)$$

to

$$U_{ND} = p_t u(\bar{t}_{ND}) + (1 - p_t) u(\underline{t}_{ND})$$

However, in the limit $p_t \to 1$, U_D should approach U_{ND}, which only occurs if $u(D) = u(\bar{t}_{ND})$. But in that case, as $p_t \to 0$, U_D does not approach U_{ND}, and so there is an unavoidable discontinuity.

A.2 Applications

Numerical Example (Preservation of Self-image)

The following is a more general version of the numerical example provided in the main body of the paper. Suppose he puts in effort $e \in [0, 1]$, and obtains reward $m \in [0, 100]$. He also has an unobserved talent $t \in [0, 1]$. The agent is doubt-prone and risk-averse for both resolved and unresolved lotteries on talent. Specifically, $u = at^{1/2}$ for some $a > 0$, and v = t. His expected utility of money is linearly separable from his utility of talent, and is equal to his expected reward Em. He therefore maximizes:

$$\tilde{W}(e) = Em(e) - C(e)$$

where $C(e) \equiv u\left(v^{-1}(Ev(t))\right) - \sum_m p(m|e) u \circ v^{-1}(Ev(t|m,e))$

The agent's prior is q that talent $t = 0$, and $1 - q$ that talent $t = 1$. He can put in level $e \in [\underline{e}, \overline{e}]$. Given that he has talent $t = 1$ or $t = 0$ and puts in effort e, his respective probabilities of obtaining monetary reward $m = 100$ are $p(100|t = 1, e) = e$ and $p(100|t = 0, e) = be$, for $b \in [0, 1)$.

Note that the ostrich effort e_0 in this example is $e = 0$, since he is certain to obtain $m = 0$, independently of his talent. It follows from the probabilities given above that:

$$p(\$0|1, e) = 1 - e$$

$$p(\$0|0, e) = 1 - be$$

$$p(100|e) = e(q + b(1 - q)$$

$$p(\$0|e) = 1 - e(q + b(1 - q))$$

$$p(1|100, e) = \frac{q}{q + b(1 - q)}$$

Solving:

$W(e) = 100 * p(100|e) + a\left(p(0|e)p(\bar{t})p(0|\bar{t}, e)\right)^{1/2} + a\left(p(100|e)p(\bar{t})p(100|\bar{t}, e)\right)^{1/2}$

$= e(100\beta + a(\beta q)^{1/2}) + aq^{1/2}\left(1 - e(1 + \beta) + \beta e^2\right)^{1/2}$

where $\beta = q + b(1 - q)$. Let $\gamma = 100\beta + a(\beta q)^{1/2}$, and $D = \frac{4\gamma^2}{a^2 q}$. Then, from the first order conditions, we obtain:

$e^2(\beta C - 4\beta^2) + e(4\beta - C)(1 + \beta) + C - (1 + \beta)^2 = 0$

The example in the text corresponds to the case $b = 0$, $q = 1/2$, and so $\beta = 1/2$, $\gamma = 50 + \frac{a}{2}$, and $d = 2D = \left(\frac{200}{a} + 2\right)^2$.

A.3 Proofs

Lemma 1 (Informed certainty equivalent). *Proof.* Define $\succeq_N$ in the same way as in the text, i.e. $\delta_f \succeq \delta_{f'} \Leftrightarrow f \succeq_N f'$ (and similarly for $\sim_N$, $\succ_N$). Note that $\succeq_N$ inherits continuity, and so there exists a function $H : \mathfrak{L}_o \to \mathbf{Z}$ such that $\delta_{H(f)} \sim_N f$ for all $f \in \mathfrak{L}_o$.

By the certainty axiom **A.3**, it follows that $\delta_{H(f)} \sim \delta_{\delta_{H(f)}}$. Hence $\delta_{H(f)} \sim \delta_f$.

Main Representation Theorem. *Proof.* Let $X = (z_1, q_1^I; z_2, q_2^I; ...; z_n, q_n^I;\ f_1, q_1^N; f_2, q_2^N; ...; f_m, q_m^N)$. By lemma 1, $\delta_f \sim \delta_{H(f)}$ for any $f \in \mathfrak{L}_o$. Hence, by a well-known implication of the independence axiom **A.4**, $X \sim \tilde{X}$, where $\tilde{X} = (z_1, q_1^I; z_2, q_2^I; ...; z_n, q_n^I;\ H(f_1), q_1^N; H(f_2), q_2^N; ...; H(f_m), q_m^N)$, and so $X \sim \tilde{X}$. Defining $\tilde{Y}$ similarly, $Y \sim \tilde{Y}$. By transitivity, $X \succ Y \Rightarrow \tilde{X} \succ \tilde{Y}$. Note that all lotteries $\tilde{X}$ and $\tilde{Y}$ are one-stage lotteries, with final outcomes as prizes. Define the preference relation $\succ_I$ in the following way: $X \succ Y \Rightarrow \tilde{X} \succ_I \tilde{Y}$. All the EU axioms hold on $\succ_I$, and so $\tilde{X} \succ \tilde{Y}$ if and only if $W(\tilde{X}) > W(\tilde{Y})$, where

$$W(\tilde{X}) = \sum_{i=1}^{n} q_i^I u(z_i) + \sum_{i=1}^{m} q_i^N u\left(H(f_{z_i})\right)$$

and W is unique up to positive affine transformation. But since $X \succ Y \Rightarrow \tilde{X} \succ \tilde{Y}$, it follows that $X \succ Y$ if and only if $W(\tilde{X}) > W(\tilde{Y})$.

To obtain the representation of H: axioms **A.1-A.4** imply that $\succeq_N$ is a weak order and that Jensen-continuity holds. The proof for the RDU representation of $\succeq_N$ then follows from Wakker (1994). Then, for any $f \in \mathfrak{L}_o$, $\delta_H(f) \sim_N f$. Since $w(1) = 1$, it follows that $v(H(f)) = v^{-1}\left(V_{RDU}(f)\right)$, and hence $H(f) = v^{-1}\left(V_{RDU}(f)\right)$, which completes the proof.

Theorem 2. *Proof.* Case 1 is shown below, and case 2 can be proven in a similar way (by changing all the signs). Suppose RDU holds for $\succeq_N$.
There are two cases two consider:

(a) f, f' have more than 2 elements:
Let $f = (z_1, p_1; ...z_i; p_i; z_{i+1}, p_{i+1}; ...; z_n, p_n)$, $f' = (z_1, p_1; ...z_i'; p_i; z_{i+1}', p_{i+1}; ...; z_n, p_n) \in \mathfrak{L}_o$ such that $f \sim_N f'$, and $(z_i', z_{i+1}') \subset (z_i, z_{i+1})$. Suppose that, for some $a \in (0,1)$ and some $z \in (z_i', z_{i+1}')$,

$$af + (1-a)\delta_z \succeq_N af' + (1-a)\delta_z$$

Since RDU holds:

$$f \sim_N f' \Rightarrow V_{RDU}(f) = V_{RDU}(f')$$

$$\Rightarrow v(z_1) + \sum_{j=2}^{i-1} w(p_j^*)[v(z_j) - v(z_{j-1})] + w(p_i^*)[v(z_i) - v(z_{i-1})] + w(p_{i+1}^*)[v(z_{i+1}) - v(z_i)]$$
$$+w(p_{i+2}^*)[v(z_{i+2}) - v(z_{i+1})] + \sum_{j=i+3}^{n} w(p_j^*)[v(z_j) - v(z_{j-1})] \quad =$$
$$v(z_1) + \sum_{j=2}^{i-1} w(p_j^*)[v(z_j) - v(z_{j-1})] + w(p_i^*)[v(z_i') - v(z_{i-1})] + w(p_{i+1}^*)[v(z_{i+1}') - v(z_i')]$$
$$+w(p_{i+2}^*)[v(z_{i+2}) - v(z_{i+1}')] + \sum_{j=i+3}^{n} w(p_j^*)[v(z_j) - v(z_{j-1})]$$

$$\Rightarrow \frac{w(p_{i+1}^*) - w(p_{i+2}^*)}{w(p_i^*) - w(p_{i+1}^*)} = \frac{v(z_i') - v(z_i)}{v(z_{i+1}) - v(z_{i+1}')} \tag{A.1}$$

Note that $af + (1-a)\delta_z = (z_1, ap_1; ...z_i; ap_i; z, 1-a; z_{i+1}, ap_{i+1}; ...; z_n, ap_n)$, where the ranking of z is due to $z \in (z_i', z_{i+1}') \subset (z_i, z_{i+1})$. Similarly, $af' + (1-a)\delta_z = (z_1, ap_1; ...z_i'; ap_i; z, 1-a; z_{i+1}', ap_{i+1}; ...; z_n, ap_n)$. Using the condition

$$af + (1-a)\delta_z \succeq_N af' + (1-a)\delta_z$$

it follows that

$$\Rightarrow v(z_1) + \sum_{j=2}^{i-1} w(ap_j^* + 1 - a)[v(z_j) - v(z_{j-1})] + w(ap_i^* + 1 - a)[v(z_i) - v(z_{i-1})]$$
$$+w(ap_{i+1}^* + 1 - a)[v(z) - v(z_i)] + w(ap_{i+1}^*)[v(z_{i+1}) - v(z)]$$
$$+w(ap_{i+2}^*)[v(z_{i+2}) - v(z_{i+1})] + \sum_{j=i+3}^{n} w(ap_j^*)[v(z_j) - v(z_{j-1})] \quad \geq$$
$$v(z_1) + \sum_{j=2}^{i-1} w(ap_j^* + 1 - a)[v(z_j) - v(z_{j-1})] + w(ap_i^* + 1 - a)[v(z_i') - v(z_{i-1})]$$
$$+w(ap_{i+1}^* + 1 - a)[v(z) - v(z_i')] + w(ap_{i+1}^*)[v(z_{i+1}') - v(z)]$$
$$+w(ap_{i+2}^*)[v(z_{i+2}) - v(z_{i+1}')] + \sum_{j=i+3}^{n} w(ap_j^*)[v(z_j) - v(z_{j-1})]$$

$$\Rightarrow \frac{w(ap^*_{i+1}) - w(ap^*_{i+2})}{w(ap^*_i + 1 - a) - w(ap^*_{i+1} + 1 - a)} \geq \frac{v(z'_i) - v(z_i)}{v(z_{i+1}) - v(z'_{i+1})} \tag{A.2}$$

Combining (1) and (2), we obtain:

$$\frac{w(ap^*_{i+1}) - w(ap^*_{i+2})}{w(ap^*_i + 1 - a) - w(ap^*_{i+1} + 1 - a)} \geq \frac{w(p^*_{i+1}) - w(p^*_{i+2})}{w(p^*_i) - w(p^*_{i+1})} \tag{A.3}$$

Note that this does not depend on the utility function v, but only on the weighting function w. Take any $\tilde{f} = (\tilde{z}_1, p_1; ...\tilde{z}_i; p_i; \tilde{z}_{i+1}, p_{i+1}; ...; \tilde{z}_n, p_n)$,
$\tilde{f}' = (\tilde{z}_1, p_1; ...\tilde{z}'_i; p_i; \tilde{z}'_{i+1}, p_{i+1}; ...; \tilde{z}_n, p_n)$ and $\tilde{z}$ such that $\tilde{z} \in (\tilde{z}'_i, \tilde{z}'_{i+1}) \subset (\tilde{z}_i, \tilde{z}_{i+1})$. It must be that $a\tilde{f} + (1-a)\delta_{\tilde{z}} \succeq_N a\tilde{f}' + (1-a)\delta_{\tilde{z}}$. Suppose not, i.e. suppose that $a\tilde{f}' + (1-a)\delta_{\tilde{z}} \succ_N a\tilde{f} + (1-a)\delta_{\tilde{z}}$. Then, redoing a similar calculation to the one above, we obtain:

$$\frac{w(ap^*_{i+1}) - w(ap^*_{i+2})}{w(ap^*_i + 1 - a) - w(ap^*_{i+1} + 1 - a)} < \frac{w(p^*_{i+1}) - w(p^*_{i+2})}{w(p^*_i) - w(p^*_{i+1})} \tag{A.4}$$

which contradicts (3). Hence ISC holds for this case.

(b) f, f' have exactly 2 elements:

Let $f = (z_1, 1-p; z_2, p), f' = (z'_1, 1-p; z'_2, p) \in \mathfrak{L}_o$ such that $f \sim_N f'$, and $(z'_1, z'_2) \subset (z_1, z_2)$. Suppose that, for some $a \in (0,1)$ and some $z \in (z'_1, z'_2)$. If $\succeq_N$ satisfies RDU, then:

$$f \sim_N f' \Rightarrow v(z_1) + w(p)[v(z_2) - v(z_1)] = v(z'_1) + w(p)[v(z'_2) - v(z'_1)]$$

$$\Rightarrow w(p) = \frac{v(z'_1) - v(z_1)}{[v(z'_1) - v(z_1)] + [v(z_2) - v(z'_2)]}$$

$$\Rightarrow \frac{w(p)}{1 - w(p)} = \frac{v(z'_1) - v(z_1)}{v(z_2) - v(z'_2)} \tag{A.5}$$

Since $af + (1-a)\delta_z = ((z_1, a(1-p); z, 1-a; z_2, ap)$ and $af' + (1-a)\delta_z = ((z'_1, a(1-p); z, 1-a; z'_2, ap)$, the condition $af + (1-a)\delta_z \succ_N af' + (1-a)\delta_z$ implies (using a similar calculation to the one used for obtaining (3)) that

$$\Rightarrow \frac{w(ap)}{1 - w(ap + 1 - a)} \geq \frac{v(z'_1) - v(z_1)}{v(z_2) - v(z'_2)} \tag{A.6}$$

and combining (4) and (5), it follows that

$$\Rightarrow \frac{w(ap)}{1-w(ap+1-a)} \geq \frac{w(p)}{1-w(p)} \tag{A.7}$$

As before, this does not depend on the $v's$, but only on the weighting function w. Take any $\tilde{f} = (\tilde{z}_1, 1-p; \tilde{z}_2, p)$, $\tilde{f}' = (\tilde{z}_1', p_1; \tilde{z}_2', p_2)$ and $\tilde{z}$ such that $\tilde{z} \in (\tilde{z}_1', \tilde{z}_2') \subset (\tilde{z}_1, \tilde{z}_2)$. It must be that $a\tilde{f} + (1-a)\delta_{\tilde{z}} \succeq_N a\tilde{f}' + (1-a)\delta_{\tilde{z}}$. Suppose not, i.e. suppose that $a\tilde{f}' + (1-a)\delta_{\tilde{z}} \succ_N a\tilde{f} + (1-a)\delta_{\tilde{z}}$. Then, redoing a similar calculation to the one above, we obtain:

$$\Rightarrow \frac{w(ap)}{1-w(ap+1-a)} < \frac{w(p)}{1-w(p)} \tag{A.8}$$

which contradicts (7). Hence ISC holds for this case as well, which completes the proof.

The following lemma is used in the proof of theorem 4:

Lemma 2t. *Let $w : [0,1] \to [0,1]$. Take any $p, q, p', p' \in [\underline{p}, \overline{p}] \subseteq [0,1]$ such that $p > p' > q'$, $q > q'$. Then if w is concave on $[\underline{p}, \overline{p}]$:*

$$\frac{w(p)-w(q)}{p-q} \leq \frac{w(p')-w(q')}{p'-q'}$$

if w is convex on $[\underline{p}, \overline{p}]$:

$$\frac{w(p)-w(q)}{p-q} \geq \frac{w(p')-w(q')}{p'-q'}$$

Proof. The proof is only shown for a concave function w. We make use of the following well-known result that a function f is concave if and only if for any $\tilde{p} > \tilde{q} > \tilde{r}$,

$$\frac{f(\tilde{p})-f(\tilde{q})}{\tilde{p}-\tilde{q}} \leq \frac{f(\tilde{p})-f(\tilde{r})}{\tilde{p}-\tilde{r}} \leq \frac{f(\tilde{q})-f(\tilde{r})}{\tilde{q}-\tilde{r}} \tag{A.9}$$

We now directly prove the claim for each of the three possible cases:

(i) $p > q > p' > q'$

Using (9) twice,

$$\frac{w(p) - w(q)}{p - q} \leq \frac{w(q) - w(p')}{q - p'} \leq \frac{w(p') - w(q')}{p' - q'}$$

(ii) $p > p' > q > q'$

Using (9) twice,

$$\frac{w(p) - w(q)}{p - q} \leq \frac{w(p') - w(q)}{p' - q} \leq \frac{w(p') - w(q')}{p' - q'}$$

(iii) $p > p' = q > q'$

In this case, the result follows immediately from (9):

$$\frac{w(p) - w(q)}{p - q} \leq \frac{w(q) - w(q')}{q - q'} = \frac{w(p') - w(q')}{p' - q'}$$

which completes the proof.

Theorem 3. *Proof.* Suppose that $\succeq_N$ satisfies RDU. We first show (A) that the weighting function w is concave implies that for any $f = (z_1, p_1; ...z_i; p_i; z_{i+1}, p_{i+1}; ...; z_n, p_n)$, $f' = (z_1, p_1; ...z_i'; p_i; z_{i+1}', p_{i+1}; ...; z_n, p_n) \in \mathfrak{L}_o$ such that $f \sim_N f'$, and $(z_i', z_{i+1}') \subset (z_i, z_{i+1})$, and for all $a \in (0, 1)$ and $z \in (z_i, z_{i+1})$,

$$af + (1 - a)\delta_z \succeq_N af' + (1 - a)\delta_z$$

We then prove the converse (B).

Proof of (A) Suppose that the weighting function w is concave. We proceed by contradiction. There are two cases to consider:

(a) f, f' have more than two elements: Let $f = (z_1, p_1; ...z_i; p_i; z_{i+1}, p_{i+1}; ...; z_n, p_n)$, $f' = (z_1, p_1; ...z_i'; p_i; z_{i+1}', p_{i+1}; ...; z_n, p_n) \in \mathfrak{L}_o$ such that $f \sim_N f'$, and $(z_i', z_{i+1}') \subset (z_i, z_{i+1})$.

Suppose there exists some $a \in (0,1)$ and some $z \in (z_i, z_{i+1})$ such that $af' + (1-a)\delta_z \succ_N af + (1-a)\delta_z$. Using the derivation of theorem 3, it follows that

$$\frac{w(ap^*_{i+1}) - w(ap^*_{i+2})}{w(ap^*_i + 1 - a) - w(ap^*_{i+1} + 1 - a)} < \frac{w(p^*_{i+1}) - w(p^*_{i+2})}{w(p^*_i) - w(p^*_{i+1})} \tag{A.10}$$

We now show:

(I) $w(ap^*_{i+1}) - w(ap^*_{i+2}) \geq a\left(w(p^*_{i+1}) - w(p^*_{i+2})\right)$

Note that $p^*_{i+1} > p^*_{i+2} > ap^*_{i+2}$, since $a \in (0,1)$, and using the definition of p^*. It is immediate that $ap^*_{i+1} > ap^*_{i+2}$. It follows, therefore, from lemma 2t, that:

$$\frac{w(p^*_{i+1}) - w(p^*_{i+2})}{p^*_{i+1} - p^*_{i+2}} \leq \frac{w(ap^*_{i+1}) - w(ap^*_{i+2})}{ap^*_{i+1} - ap^*_{i+2}}$$

Rearranging, we obtain $w(ap^*_{i+1}) - w(ap^*_{i+2}) \geq a\left(w(p^*_{i+1}) - w(p^*_{i+2})\right)$.

(II) $w(ap^*_i + 1 - a) - w(ap^*_{i+1} + 1 - a) \leq a\left(w(p^*_i) - w(p^*_{i+1})\right)$

Note that $ap^*_i + 1 - a > p^*_i$, since $a, p^*_i \in (0,1)$ implies that $1 - a > p^*_i(1-a)$. Similarly, $ap^*_{i+1} + 1 - a > p^*_{i+1}$, and we know that $p^*_i > p^*_{i+1}$. Using lemma 2t, it follows that:

$$\frac{w(ap^*_i + 1 - a) - w(ap^*_{i+1} + 1 - a)}{(ap^*_i + 1 - a) - \left(ap^*_{i+1} + 1 - a\right)} \leq \frac{w(p^*_i) - w(p^*_{i+1})}{p^*_i - p^*_{i+1}}$$

Rearranging, we obtain $w(ap^*_i + 1 - a) - w(ap^*_{i+1} + 1 - a) \leq a\left(w(p^*_i) - w(p^*_{i+1})\right)$

Combining (I) and (II) (noting that both sides of (II) are greater than zero), it follows that

$$\frac{w(ap^*_{i+1}) - w(ap^*_{i+2})}{w(ap^*_i + 1 - a) - w(ap^*_{i+1} + 1 - a)} \geq \frac{w(p^*_{i+1}) - w(p^*_{i+2})}{w(p^*_i) - w(p^*_{i+1})} \tag{A.11}$$

which is a contradiction of (10).

(b) f, f' have exactly 2 elements:

Let $f = (z_1, 1-p; z_2, p)$, $f' = (z'_1, 1-p; z'_2, p) \in \mathfrak{L}_o$ such that $f \sim_N f'$, and $(z'_1, z'_2) \subset (z_1, z_2)$. Suppose there exists some $a \in (0,1)$ and some $z \in (z_1, z_2)$ such that $af' +$

$(1-a)\delta_z \succ_N af + (1-a)\delta_z$. Using the derivation of theorem 3, it follows that

$$\frac{w(ap)}{1-w(ap+1-a)} < \frac{w(p)}{1-w(p)} \tag{A.12}$$

We now show:

(I) $w(ap) \geq aw(p)$

$a \in (0,1)$ and so $p > ap > 0$. It follows from the well-known result (9) used in proving lemma 2t that:

$$\frac{w(p)-w(0)}{p} \leq \frac{w(ap)-w(0)}{ap-0}$$

Using $w(0) = 0$ and rearranging, we obtain $w(ap) \geq aw(p)$

(II) $1 - w(ap+1-a) \leq a\,(1-w(p))$

Note that $1 > ap+1-a > p$, since it is immediate from $a, p \in (0,1)$ that $a > ap$ and $1-a > p(1-a)$.

Using (9) again,

$$\frac{w(1)-w(ap+1-a)}{1-(ap+1-a)} \leq \frac{w(1)-w(p)}{1-p}$$

Using $w(1) = 1$ and rearranging, we obtain that $1-w(ap+1-a) \leq a\,(1-w(p))$.

Combining (I) and (II), we obtain

$$\frac{w(ap)}{1-w(ap+1-a)} \geq \frac{w(p)}{1-w(p)} \tag{A.13}$$

which contradicts (12).

Proof of (B) Suppose that for any $f = (z_1, p_1; ...z_i; p_i; z_{i+1}, p_{i+1}; ...; z_n, p_n)$, $f' = (z_1, p_1; ...z'_i; p_i; z'_{i+1}, p_{i+1}; ...; z_n, p_n) \in \mathfrak{L}_o$ such that $f \sim_N f'$, and $(z'_i, z'_{i+1}) \subset (z_i, z_{i+1})$, and for all $a \in (0,1)$ and $z \in (z_i, z_{i+1})$,

$$af + (1-a)\delta_z \succeq_N af' + (1-a)\delta_z$$

We proceed as follows: (a) we first show that there is no interval $[\underline{p}, \overline{p}] \subseteq [0,1]$ on which w

is strictly convex; (b) we then show that there is no interval $[\underline{p}, \overline{p}] \subseteq [0,1]$ such that for all $p \in [\underline{p}, \overline{p}]$, $w(p)$ is 'under the diagonal', i.e. $\frac{w(\overline{p})-w(p)}{\overline{p}-p} > \frac{w(\overline{p})-w(\underline{p})}{\overline{p}-\underline{p}} > \frac{w(p)-w(\underline{p})}{p-\underline{p}}$ (note that with stronger smoothness assumptions this would be sufficient for concavity); (c) we use results (a) and (b) to prove that w must be concave. We first note that it follows from the claim and from the derivation of theorem 3 that:

$$\frac{w(ap_1) - w(ap_2)}{w(ap_0 + 1 - a) - w(ap_1 + 1 - a)} \geq \frac{w(p_1) - w(p_2)}{w(p_0) - w(p_1)} \tag{A.14}$$

for all $0 \leq p_2 < p_1 < p_0 \leq 1$ and $a \in (0,1)$.

(a) We proceed by contradiction: suppose there does exist an interval $[\underline{p}, \overline{p}] \subseteq [0,1]$ on which w is strictly convex. Let $\underline{p} < p_2 < p_1 < p_0 < \overline{p}$, and let $\{\frac{\underline{p}}{p_2}, \frac{1-\overline{p}}{1-p_0}\} < a < 1$. It follows that $\underline{p} < ap_2 < ap_1 < ap_1 + 1 - a < ap_0 + 1 - a)\overline{p}$. Using lemma 2t, it follows that:

$$\frac{w(p_1) - w(p_2)}{p_1 - p_2} > \frac{w(ap_1) - w(ap_2)}{ap_1 - ap_2} \tag{A.15}$$

$$\frac{w(ap_0 + 1 - a) - w(ap_1 + 1 - a)}{(ap_0 + 1 - a) - (ap_1 + 1 - a)} > \frac{w(p_0) - w(p_1)}{p_0 - p_1} \tag{A.16}$$

Rearranging and combining (15) and (16), it follows that

$$\frac{w(ap_1) - w(ap_2)}{w(ap_0 + 1 - a) - w(ap_1 + 1 - a)} < \frac{w(p_1) - w(p_2)}{w(p_0) - w(p_1)}$$

which contradicts (14).

(b) We proceed again by contradiction: suppose that there does exist an interval $[\underline{p}, \overline{p}] \subseteq [0,1]$ such that $\frac{w(\overline{p})-w(p)}{\overline{p}-p} > \frac{w(\overline{p})-w(\underline{p})}{\overline{p}-\underline{p}} > \frac{w(p)-w(\underline{p})}{p-\underline{p}}$ for all $p \in [\underline{p}, \overline{p}]$.
Let $a = 1 - (\overline{p} - \underline{p}) + \epsilon$, for an arbitrarily small ϵ. Let $\tilde{p} = \underline{p}/a$. Using result (a), $[\tilde{p}, \tilde{p} + \delta]$ cannot be strictly convex, for any $\delta \in (0, 1 - \tilde{p}]$. We can therefore find $\{p_0, p_1, p_2\} \in [\tilde{p}, \tilde{p} + \delta]$ such that $p_2 < p_1 < p_0$ and

$$\frac{w(p_1) - w(p_2)}{p_1 - p_2} \geq \frac{w(p_0) - w(p_1)}{p_0 - p_1} \tag{A.17}$$

As δ, ϵ become arbitrarily small (and $a\delta \leq \epsilon$), $ap_2 \to \underline{p}$, $ap_0 + 1 - a \to \overline{p}$ and

$\{ap_2, ap_1, ap_1 + 1 - a, ap_0 + 1 - a\} \in [\underline{p}, \overline{p}]$. We therefore have that for small enough δ, ϵ,

$$\frac{w(ap_0 + 1 - a) - w(ap_1 + 1 - a)}{(ap_0 + 1 - a) - (ap_1 + 1 - a)} > \frac{w(\overline{p}) - w(\underline{p})}{\overline{p} - \underline{p}} \tag{A.18}$$

and

$$\frac{w(\overline{p}) - w(\underline{p})}{\overline{p} - \underline{p}} > \frac{w(ap_1) - w(ap_2)}{a(p_1 - p_2)} \tag{A.19}$$

Combining (18) and (19):

$$\frac{w(ap_1) - w(ap_2)}{w(ap_0 + 1 - a) - w(ap_1 + 1 - a)} < \frac{p_1 - p_2}{p_0 - p_1} \tag{A.20}$$

Combining (17) and (20), we obtain:

$$\frac{w(ap_1) - w(ap_2)}{w(ap_0 + 1 - a) - w(ap_1 + 1 - a)} < \frac{w(p_1) - w(p_2)}{w(p_0) - w(p_1)}$$

which contradicts (14).

(c) We now prove that w is concave. Suppose not, i.e. suppose there exist $0 \leq p < q < r < 1$ such that

$$\frac{w(r) - w(q)}{r - q} > \frac{w(q) - w(p)}{q - p} \tag{A.21}$$

Let $a = 1 - (r - q) + \epsilon$, for an arbitrarily small ϵ. Let $\tilde{p} = q/a$. Using result (a), $[\tilde{p} - \delta, \tilde{p}]$ cannot be strictly convex, for any $\delta \in (0, \tilde{p}]$. We can therefore find $\{p_0, p_1, p_2\} \in [\tilde{p} - \delta, \tilde{p}]$ such that $p_2 < p_1 < p_0$ and

$$\frac{w(p_1) - w(p_2)}{p_1 - p_2} \geq \frac{w(p_0) - w(p_1)}{p_0 - p_1} \tag{A.22}$$

As δ, ϵ become arbitrarily small (and $a\delta \leq \epsilon$), $ap_1 \to q$, $ap_0 + 1 - a \to r$, $\{ap_2, ap_1\} \in (p, q]$ and $\{ap_1 + 1 - a, ap_0 + 1 - a\} \in [q, r]$.

Using result (b), we have can find some (small enough) δ, ϵ such that

$$\frac{w(ap_1) - w(ap_2)}{a(p_1 - p_2)} \leq \frac{w(q) - w(p)}{q - p} \tag{A.23}$$

$$\frac{w(ap_0 + 1 - a) - w(ap_1 + 1 - a)}{(ap_0 + 1 - a) - (ap_1 + 1 - a)} \geq \frac{w(r) - w(q)}{r - q} \tag{A.24}$$

Combining (21, (23) and (24) we have

$$\frac{w(ap_1) - w(ap_2)}{w(ap_0 + 1 - a) - w(ap_1 + 1 - a)} < \frac{p_1 - p_2}{p_0 - p_1} \tag{A.25}$$

Combining (22) and (25), we have

$$\frac{w(ap_1) - w(ap_2)}{w(ap_0 + 1 - a) - w(ap_1 + 1 - a)} < \frac{w(p_1) - w(p_2)}{w(p_0) - w(p_1)}$$

which contradicts (14), and completes the proof.

Theorem 4. *Suppose that axioms **A.1** through **A.4** and the RDU axioms hold, and let u and v be the utility functions associated with the resolved and unresolved lotteries, respectively, and w be the decision weight associated with the unresolved lotteries. In addition, suppose that u, v are both differentiable. Then:*
(i) If there exists $p \in (0,1)$ such that $p < w(p)$, then there exists an $f \in \mathfrak{L}_\mathsf{o}$ such that $\delta_f \succ f$. Similarly, if there exists $p' \in (0,1)$ such that $p' > w(p')$, then there exists an $f' \in \mathfrak{L}_\mathsf{o}$ such that $f' \succ \delta'_f$.
(ii) If $\succeq$ exhibits doubt-aversion, then $p \geq w(p)$ for all $p \in (0,1)$. Moreover, if u exhibits stronger diminishing marginal utility than v (i.e. $u = \lambda \circ v$ for some continuous, weakly concave, and increasing λ on $v([\underline{z}, \bar{z}])$), then $\succeq_N$ violates quasi-concavity. (that is, there exists some $f', f'' \in \mathfrak{L}_\mathsf{o}$, and $\alpha \in (0,1)$ such that $f' \succ f''$ and $f'' \succ_N \alpha f' + (1-\alpha) f''$).
Similarly, if $\succeq$ exhibits doubt-proneness, then $p \leq w(p)$ for all $p \in (0,1)$. Moreover, if v exhibits stronger diminishing marginal utility than u , then $\succeq_N$ violates quasi-convexity. (that is, there exists some $f', f'' \in \mathfrak{L}_\mathsf{o}$, and $\alpha \in (0,1)$ such that $f' \succ f''$ and $\alpha f' + (1-\alpha) f'' \succ_N f'$).

Proof. (i) Suppose not, i.e. suppose that there exists $p \in (0,1)$ such that $p < w(p)$, and that $f \succeq \delta_f$ for all $f \in \mathfrak{L}_o$. Let $f_\epsilon = (z; 1-p; z+\epsilon, p)$ for some $z \in \mathbf{Z}$, $p \in \mathfrak{L}_o$, $0 < \epsilon < \bar{z} - z$. Since $f \succeq \delta_f$, by continuity (and using the certainty axiom), there exists a $\tilde{z}_\epsilon \in (z, z+\epsilon)$ such that $f \succeq \left[\delta_{\tilde{z}_\epsilon} \sim \delta_{\delta_{\tilde{z}_\epsilon}}\right] \succeq \delta_f$. Hence:

$$(1-p)u(z) + pu(z+\epsilon) \geq u(\tilde{z}_\epsilon)$$

$$w(p)\left(v(z+\epsilon) - v(z)\right) + v(z) \leq v(\tilde{z}_\epsilon)$$

Rearranging:

$$p \geq \frac{u(\tilde{z}_\epsilon) - u(z)}{u(z+\epsilon) - u(z)}$$

$$w(p) \leq \frac{v(\tilde{z}_\epsilon) - v(z)}{v(z+\epsilon) - v(z)}$$

Hence:

$$\frac{u(\tilde{z}_\epsilon) - u(z)}{u(z+\epsilon) - u(z)} - \frac{v(\tilde{z}_\epsilon) - v(z)}{v(z+\epsilon) - v(z)} \leq p - w(p)$$

But as $\epsilon \to 0$, $\frac{u(\tilde{z}_\epsilon)-u(z)}{u(z+\epsilon)-u(z)} \to \frac{u'(z)}{u'(z)}$, and $\frac{v(\tilde{z}_\epsilon)-v(z)}{v(z+\epsilon)-v(z)} \to \frac{v'(z)}{v'(z)}$, by differentiability. Since the left-hand-side goes to $1-1=0$ in the limit, while the right-hand-side does not change, it must be that $0 \leq p - w(p)$. But this is a contradiction, since $p < w(p)$.
The second part of the result can be proved in a similar manner, for the case $p' > w(p')$.

(ii) The result is only shown for doubt-aversion; a similar reasoning holds for doubt-proneness. By the contrapositive of (i), it is immediate that if $f \succeq \delta_f$ for all $f \in \mathfrak{L}_o$, then $w(p) \leq p$ for all $p \in (0,1)$. Now suppose that $f \succ \delta_f$ for some f, and that u is a (weakly) concave transformation of v. If w is not concave, then $\succeq_N$ cannot be quasi-concave, by Wakker (1994) theorem 25. Since $w(0) = 0$, $w(1) = 1$, $w(p) \geq p$ for a concave function. We have that $w(p) \leq p$, and so it suffices to show that $w(p) < p$ for some p. Suppose not. That is, $w(p) = p$ for all p. Since u is more concave than v, it must be that $u^{-1}(EU(f)) \leq v^{-1}(EV(f))$(that is, the certainty equivalent of f for the informed lotteries is not bigger than the certainty equivalent of f for the unresolved lotteries, by a well known result). However, since $f \succ \delta_f$, it must also be that $u^{-1}(EU(f)) > v^{-1}(EV(f))$, which is a contradiction.

Note that if $f \sim \delta_f$ for all $f \in \mathfrak{L}_o$, than trivially, u is a linear transformation of v, and $w(p) = p$.

Corollary 4.1. *Proof.* To prove (i) $\Rightarrow$ (ii):If $\succeq_N$ displays mean-preserving risk-aversion, then $w(p)$ is convex, by Chew, Epstein and Safra (1986) or Grant, Kajii and Polak (2000). Since $w(0) = 0$, $w(1) = 1$, it must be that $p \geq w(p)$. Since $\delta_f \succeq f$, it follows from result (ii) that $p \leq w(p)$. Hence $w(p) = p$, implying that $\succeq_N$ satisfies expected utility.

Since $\delta_f \succeq f$ for all $f \in \mathfrak{L}_o$, and both u and v are of EU form, u must be a concave transformation of v. This is well-known, see for instance Kreps-Porteus (1978).

The other direction, (ii) $\Rightarrow$ (i), is trivial: if u and v are concave then they both display mean-preserving risk aversion by well known results, and if u is a concave transformation of v then $\delta_f \succeq f$ for all $f \in \mathfrak{L}_o$.

Theorem 5.

Proof. If $u(z) = v(z)$ for all $z \in \mathbf{Z}$, then $\delta_f \succeq f$ if and only if

$$u(z_1) + \sum_{i=2}^{m} [u(z_i) - u(z_{i-1})] w(p_i^*) \geq \sum_{i=1}^{m} u(z_i) p(z_i) \tag{A.26}$$

$$\Leftrightarrow u(z_1) + \sum_{i=2}^{m} [u(z_i) - u(z_{i-1})] w(p_i^*) \geq u(z_1) + \sum_{i=2}^{m} [u(z_i) - u(z_{i-1})] p_i^* \tag{A.27}$$

$$\Leftrightarrow \sum_{i=2}^{m} [u(z_i) - u(z_{i-1})](w(p_i^*) - p_i^*) \geq 0. \tag{A.28}$$

This expression is always true if and only if $w(p) \geq p$ for all $p \in [0, 1]$. For the agent to be doubt-prone, the inequality in (A.28) must be strict somewhere, hence $w(p) > p$ for some $p \in (0, 1)$. Now suppose $u = \lambda \circ v$ for some continuous, weakly concave and increasing λ. By theorem 4, the agent is doubt-prone everywhere only if $p \leq w(p)$. Now suppose that $w(p) > p$. Then using the same argument as above, we have:

$$v(z_1) + \sum_{i=2}^{m} [v(z_i) - v(z_{i-1})] w(p_i^*) \geq \sum_{i=1}^{m} v(z_i) p(z_i). \tag{A.29}$$

Hence:

$$u\left(v^{-1}\left(v(z_1)+\sum_{i=2}^{m}[v(z_i)-v(z_{i-1})]w(p_i^*)\right)\right)\geq u\left(v^{-1}\left(\sum_{i=1}^{m}v(z_i)p(z_i)\right)\right). \tag{A.30}$$

But by concavity of $u\left(v^{-1}(\cdot)\right)$, we know that

$$u\left(v^{-1}\left(\sum_{i=1}^{m}v(z_i)p(z_i)\right)\right)\geq \sum_{i=1}^{m}u(z_i)p(z_i), \tag{A.31}$$

with strict inequality somewhere, hence the agent is doubt-prone everywhere. This completes the proof.

Preservation of self-image. For an agent who is doubt-prone and risk-averse for both resolved and unresolved lotteries, the following holds:

$$C(e)\equiv u\circ v^{-1}(Ev(t))-\sum_m p(m|e)u\circ v^{-1}(Ev(t|m,e))\geq 0$$

Proof. Note that $u\circ v^{-1}(\cdot)$ is concave. Hence

$$\begin{aligned}\sum_m p(m|e)u\circ v^{-1}(Ev(t|m,e)) &\leq u\circ v^{-1}\left(\sum_m p(m|e)(Ev(t|m,e))\right)\\ &\leq u\circ v^{-1}\left(\sum_m p(m|e)\sum_t \frac{p(m|t,e)p(t)}{p(m|e)}v(t)\right)\\ &\leq u\circ v^{-1}\left(\sum_m\sum_t p(m|t,e)p(t)v(t)\right)\\ &\leq u\circ v^{-1}\left(\sum_t\sum_m p(m|t,e)p(t)v(t)\right)\\ &\leq u\circ v^{-1}\left(\sum_t p(t)v(t)\right)=u\circ v^{-1}(Ev(t))\end{aligned}$$

Doubt-neutrality result. *Proof.* If (i) holds, then it is trivial that (ii) and (iii) hold as well.

To show that (ii) $\Rightarrow$ (i):

Suppose not. Then there exists an $f \in \mathfrak{L}_o$ such that either $f \succ \delta_f$ or $\delta_f \succ f$. Suppose $f \succ \delta_f$. Then by lemma 1, there exists an $H(f) \in \mathbf{Z}$ such that $\delta_f \sim \delta_{H(f)}$. By transitivity, $f \succ \delta_f \Leftrightarrow f \succ \delta_{H(f)}$, and so by (ii), $\delta_f \succ \delta_{\delta_{H(f)}}$. By transitivity again, $\delta_{H(f)} \succ \delta_{\delta_{H(f)}}$, but this violates the certainty axiom **A.1**. Now suppose that $\delta_f \succ f$. Then $\delta_{H(f)} \succ f$, and by (ii), $\delta_{\delta_{H(f)}} \succ \delta_f \Leftrightarrow \delta_{\delta_{H(f)}} \succ \delta_{H(f)}$, which violates **A.1**.

To show that (iii) $\Rightarrow$ (i):

Suppose not. Then there exists an $f \in \mathfrak{L}_o$ such that either $f \succ \delta_f$ or $\delta_f \succ f$. Suppose that $f \succ \delta_f$. Note that by continuity, it is also the case that there exists an $\tilde{H} \in \mathbf{Z}$ such that $f \sim \delta_{\tilde{H}(f)}$. By the certainty axiom **A.1**, $\delta_{\tilde{H}(f)} \sim \delta_{\delta_{\tilde{H}(f)}}$. By transitivity, $\delta_{\delta_{\tilde{H}(f)}} \succ \delta_f$, and by (iii), $\delta_{\tilde{H}(f)} \succ f$. But this is a contradiction. Now suppose that $\delta_f \succ f$. Then $\delta_f \succ \delta_{\delta_{\tilde{H}(f)}} \Leftrightarrow f \succ \delta_{\tilde{H}(f)}$ which is a contradiction.

Appendix B

Appendix B

The appendix is structured as follows. I first present the details of the signal structure, and then analyze the Dutch Lottery and the safe allocation bias. I then prove theorem 1 and theorem 2. In the discussion that follows, all lotteries are ordered from best to worse; that is, if $f = (r_1, p_1; ...r_m; p_m)$, then $r_1 > ... > r_m$.

Signal structure

Recall that:

$$\overline{p}^*(r_i, r_j|f; f') \equiv \sum_{\{r\in f, r'\in f'\}} p(r; f)p(r'; f')\mathbf{I}_{\{u_r(r_i)-u_r(r_j)\ >\ u_r(r)-u_r(r')\}}.$$

$$\underline{p}^*(r_i, r_j|f; f') \equiv \sum_{\{r\in f, r'\in f'\}} p(r; f)p(r'; f')\mathbf{I}_{\{u_r(r_i)-u_r(r_j)\ <\ u_r(r)-u_r(r')\}}.$$

For set S, let

$$K_{-i}(S) = \{g \in \mathfrak{L}_{\mathfrak{r}}|g \in M,\ where\ M \in S\backslash M_i\}.$$

This set contains all the lotteries that are not in menu M_i, but it does not specify which menu they are from. Now suppose that g' ia a lottery in an unchosen menu, while f is *still the received lottery. Let $\alpha > 1$. Define $\overline{p}^*_K$ and $\underline{p}^*_K$ in an analogous way:*

$$\overline{p}^*_K(r_i, r_j|f; g') \equiv (1 - p(r_i|f)) \sum_{\{r\in f, r'\in g'\}} p(r; f)p(r'; g')^{\alpha}\mathbf{I}_{\{u_r(r_i)-u_r(r_j)\ >u_r(r)-u_r(r')\}}.$$

$$\underline{p}^*_K(r_i, r_j|f; g') \equiv (1 - p(r_i|f)) \sum_{\{r\in f, r'\in g'\}} p(r; f)p(r'; g')^{\alpha}\mathbf{I}_{\{u_r(r_i)-u_r(r_j)\ <u_r(r)-u_r(r')\}}.$$

Weights $\overline{p}^*_K(r_i, r_j|f; f')$ and $\underline{p}^*_K(r_i, r_j|f; f')$ will be used to assess the signal over unchosen menus. In accordance with signal property S.4, a lottery in an unchosen menu gives the agent a less informative signal than if it were chosen, compared to if it were in a chosen menu, unless it is degenerate. Note that if the agent receives an outcome $r_i|f$ for sure, then since $1 - p(r_i|f)) = 0$, the agent obtains no information from the unchosen menu. Let $X(g)$ be the set of outcomes that can be reached with positive probability in g. That is, if $g = (r_1, p_1; ...r_{n_g}, p_{n_g})$, then $X(g) = \{r_1, ...r_{n_g}\}$. The signal structure, given history $\mathrm{H} = (r_1, ..., r_n; f_j; M; S)$, *is:*

$$\begin{aligned} p(t|\mathrm{H}) = p_t\Bigg(1 + &\sum_{k\in\{1..n\}} b_{jk}\left(\frac{u_r(r_j) - u_r(r_k)}{u_r(\overline{r}) - u_r(\underline{r})}\right)\mathbf{I}_{\{u_r(r_j)\geq u_r(r_k)\}}(\overline{p}^*(r_j, r_k|f_j; f_k)) \\ &-c_S \sum_{k'\in\{1..n\}}\left(\frac{u_r(r_{k'}) - u_r(r_j)}{u_r(\overline{r}) - u_r(\underline{r})}\right)\ \mathbf{I}_{\{u_r(r_j)<u_r(r_{k'})\}}(\underline{p}^*(r_j, r_{k'}|f_j; f_{k'})) \qquad \text{(B.1)} \\ + &\sum_{g_h\in K_{-i}(S)}\sum_{\tilde{r}\in X(g_h)}\left(\frac{u_r(r_j) - u_r(\tilde{r})}{u_r(\overline{r}) - u_r(\underline{r})}\right)\big[\mathbf{I}_{\{u_r(r_j)\geq u_r(\tilde{r})\}}b_{K,jh}(\overline{p}^*(r_j, \tilde{r}|f_j; g_h)) \\ &+\ \ \mathbf{I}_{\{u_r(r_j)<u_r(\tilde{r})\}}c_S(\underline{p}^*(r_j, \tilde{r}|f_j; g_h))\big]\Bigg) \qquad \text{(B.2)} \end{aligned}$$

The first part of the expression, (B.1), *corresponds to the signal received from the chosen menu. The second part (between* (B.1) *and* (B.2)*), corresponds to the signal received from unchosen menus. The differences between the two are as follows. The weights $\overline{p}^*_K$ and $\underline{p}^*_K$ associated with the unchosen menus are smaller than for the chosen menus, since unchosen menus are less informative. In addition, since the agent does not know which outcome has occurred, the signal uses an ex-*ante *viewpoint for the lotteries from unchosen menus, and an ex-*post *viewpoint for the lottery whose outcome he actually receives. Finally, note*

that we are using set $K_{-i}(S)$, so that the agent does not take into consideration which unchosen menu a lottery is from, or whether it is in more than one menu. I now turn to characterizing the constants. Let $N(S)$ be the total number of lotteries in set S. This number is useful for ensuring that the agent's ex-post belief is in the correct range, $[0,1]$ (this number is higher than necessary, but it is convenient). Finally, let $c \in (0,1)$ be a parameter. Again, the upper bound (1) is chosen to make sure that the ex-post belief is within $[0,1]$. The lower bound (0) is chosen to ensure that the signal goes in the desired direction. Then:
$c_S = \frac{c}{N(S)-1} \min\{\frac{1-p_t}{p_t}, \frac{p_t}{1-p_t}\}$.

Index b_{jk} is defined to ensure that the non-manipulability assumption holds. Specifically:

$$b_{jk} = \frac{c_S \sum_{r_l \in X(f_j)} \sum_{r_m \in X(f_k)} p(r_l|f_j)p(r_m|f_k)\left(u_r(r_m) - u_r(r_l)\right) \quad \mathbf{I}_{\{u_r(r_l)<u_r(r_m)\}}(\underline{p}^*(r_l, r_m|f_j; f_k))}{\sum_{r_{l'} \in X(f_j)} \sum_{r_{m'} \in X(f_k)} p(r_{l'}|f_j)p(r'_m|f_k)\left(u_r(r_{l'}) - u_r(r_{m'})\right) \mathbf{I}_{\{u_r(r_{l'})\geq u_r(r_{m'})\}}(\overline{p}^*(r_{l'}, r_{m'}|f_j; f_k))} \tag{B.3}$$

It is clear that comparing any two lotteries f and f' at the ex-ante stage, the expectation of t is still $p(t)$. The only difference with $b_{K,jk}$ is that the weighting functions $\overline{p}^*$ and $\underline{p}^*$ are replaced with $\overline{p}^*_K$ and $\underline{p}^*_K$:

$$b_{K,jk} = \frac{c_S \sum_{r_l \in X(f_j)} \sum_{r_m \in X(f_k)} p(r_l|f_j)p(r_m|f_k)\left(u_r(r_m) - u_r(r_l)\right) \quad \mathbf{I}_{\{u_r(r_l)<u_r(r_m)\}}(\underline{p}^*_K(r_l, r_m|f_j; f_k))}{\sum_{r_{l'} \in X(f_j)} \sum_{r_{m'} \in X(f_k)} p(r_{l'}|f_j)p(r'_m|f_k)\left(u_r(r_{l'}) - u_r(r_{m'})\right) \mathbf{I}_{\{u_r(r_{l'})\geq u_r(r_{m'})\}}(\overline{p}^*_K(r_{l'}, r_{m'}|f_j; f_k))} \tag{B.4}$$

It is straightforward to show that the desired signal properties S.1 through S.4 hold. We now proceed to the Dutch lottery example.

Dutch lottery example

Consider the special case of the Dutch lottery example in which $f = (r_h, 0.5; r_l, 0.5)$ and the safe lottery $g = (r_{CE}, 1)$, where r_{CE} is the monetary certainty equivalent of lottery f.

Initial Set	Chosen menu	Chosen lottery
Feedback: $S = \{\{f, g\}\}$	$\{f, g\}$	f
No feedback: $S' = \{\{f\}, \{g\}\}$	$\{g\}$	g

Table B.1: Dutch lottery example.

That is, $u_r(r_{CE}) = 0.5u_r(r_h) + 0.5u_r(r_l)$. In this case, the doubt-prone agent is indifferent between receiving f and $g = (r_{CE}, 1)$ in the feedback case:

First, note that $u(r_h) - u_r(r_{CE}) = u_r(r_{CE}) - u_r(r_l)$. Denote the parameter b_jk from the signal structure b_{12} if f is chosen, and b_{21} if g is chosen. Then:

$$b_{12} = c_S \frac{u_r(r_{CE}) - u_r(r_l)}{u_r(r_h) - u(r_{CE})} = c_S \tag{B.5}$$

and similarly,

$$b_{21} = c_S \frac{u_r(r_h) - u_r(r_{CE})}{u_r(r_{CE}) - u_r(r_l)} = c_S. \tag{B.6}$$

Using again

$$u(r_h) - u_r(r_{CE}) = u_r(r_{CE}) - u_r(r_l) \tag{B.7}$$

$$\Rightarrow \frac{u_r(r_h) - u_r(r_{CE})}{u_r(\overline{r}) - u_r(\underline{r})} = \frac{u_r(r_{CE}) - u_r(r_l)}{u_r(\overline{r}) - u_r(\underline{r})}, \tag{B.8}$$

we obtain:

$$\begin{aligned} W(f, \{f, g\} | \{\{\{f, g\}\}\}) = 0.5\big(u_r(r_h) + u_r(r_l) + u_{p_t}\left(p_t\left(1 + 0.5c_S \frac{u_r(r_h) - u_r(r_{CE})}{u_r(\overline{r}) - u_r(\underline{r})}\right)\right) \\ + u_{p_t}\left(p_t\left(1 - 0.5c_S \frac{u_r(r_h) - u_r(r_{CE})}{u_r(\overline{r}) - u_r(\underline{r})}\right)\right)\big) = W(g, \{f, g\} | \{\{f, g\}\}). \end{aligned} \tag{B.9}$$

The agent is therefore indifferent between the two. But now, consider the no-feedback case.

It is immediate that lottery f is exactly as informative as in the feedback case. Hence,

$$W(f,\{f\},\{g\}|\{\{f\},\{g\}\}) = W(f,\{f,g\}|\{\{f,g\}\}) = W(g,\{f,g\}|\{\{f,g\}\}). \qquad \text{(B.10)}$$

but now, lottery g must be less informative than before, by signal property S.4. Since the agent is doubt-prone (u_{p_t} is concave), it is immediate that

$$W(g,\{f\},\{g\}|\{\{f\},\{g\}\}) > W(g,\{f,g\}|\{\{f,g\}\}). \qquad \text{(B.11)}$$

which implies that

$$W(g,\{f\},\{g\}|\{\{f\},\{g\}\}) > W(f,\{f,g\}|\{\{f,g\}\}). \qquad \text{(B.12)}$$

Hence, the doubt-prone agent who is indifferent between f and g in the feedback case strictly prefers lottery g in the no-feedback case.

Safe allocation bias

Let $f_b = (r_h - P, q; r_l - P, 1-q)$, $f_s = (P - r_h, q; P - r_l, 1-q)$, and $g = (0,1)$.I assume that f_b and f_s are correlated, so that if the agent chooses f_b and $r_h - P$ occurs, then he knows that lottery f_s would have yielded $P - r_h$. If he receives $r_l - P$ from lottery f_b, then he knows that lottery f_s would have yielded lottery $P - r_l$ (and similarly he chooses lottery f_s. While this correlation is technically outside the scope of this framework, the extension is straightforward; I assume exactly the same signal structure as before, but I do not allow for histories in which f_b yields and $r_h - P$ while f_s yields $P - r_l$, and similarly for $r_l - P$ and $P - r_l$. I assume that if the agent chooses g, then he does not observe the resolution of either f_b or f_s. The extension to allow the correlation is immediate, but notationally cumbersome. I denote set $S = \{\{f_b, f_s\}, g\}$, though formally this does not take into account the correlation between f_b and f_s. Suppose that the possible price range is an interval $[\underline{P}, \overline{P}]$. The agent is risk-neutral, so that that there exists exactly one price $\hat{P}$ for which the agent would be indifferent, in the standard EU setting, between f_b, f_s and g.

I now show that in this framework, there is a price range $[P_B, P_G]$ at which the agent *chooses lottery g. In addition, he prefers f_s for $P > P_G$ and f_b for $P < P_B$. It is clear that for any price such that $r_l - P > 0$, the agent always chooses f_b, and that for any price such that $r_h - P < 0$, the agent always chooses f_S. We now focus on the range of prices for which $r_h - P > 0 > r_l - P$. For simplicity, set $u(\overline{r}) - u(\underline{r}) = 1$, this has no impact on the result. If the agent chooses lottery f_b, then his value function is:*

$$\begin{aligned} W(f_b|\{f_b,f_s\},S) =& q u_r(r_h - P) + (1-q)u_r(r_l - P) \\ &+ q u_{p_t}\left(p_t\left(1 + b_{b0}(u_r(r_h - P) - u(0)) + b_{bs}u_r((r_h - P) - u_r(P - r_l))\right)\right) \\ &+ (1-q)u_{p_t}\left(p_t\left(1 + c(u_r(0) - u_r(P - r_l)) + c(u_r(P - r_l) - u_r(r_l - P))\right)\right). \end{aligned} \tag{B.13}$$

where b_{b0}, b_{bs} are the parameters associated with the comparison of $r_h - P$ from f_b and 0 from g, and the comparison $r_h - P$ from f_b and $P - r_h$ from f_s, respectively. If he chooses lottery f_s, then his value function is:

$$\begin{aligned} W(f_s|\{f_b,f_s\},S) =& q u_r(P - r_h) + (1-q)u_r(P - r_l) \\ &+ q u_{p_t}\left(p_t\left(1 + b_{s0}(u(P - r_l) - u(0)) + b_{sb}u_r((P - r_l) - u_r(r_l - P))\right)\right) \\ &+ (1-q)u_{p_t}\left(p_t\left(1 + c(u_r(0) - u_r(P - r_h)) + c(u_r(r_h - P) - u_r(P - r_h))\right)\right). \end{aligned} \tag{B.14}$$

where b_{s0} and b_{sb} are the parameters associated with the comparison of $P - r_l$ from f_s and 0 from g, and the comparison $P - r_h$ from f_s and $r_h - P$ from f_b, respectively.

Finally, the value function for choosing g is simply $W(g|\{g\},S) = u_r(0)$, since the agent acquires no new information.

Now consider again price $P = \hat{P}$, for which the agent would be indifferent, in the standard EU model (with risk-neutrality) between f_b,f_s and g. That is:

$$qDu_r(r_h - \hat{P}) + (1-q)u_r(r_l - \hat{P}) = u(0) = qu_r(\hat{P} - r_h) + (1-q)u_r(\hat{P} - r_l). \tag{B.15}$$

By property S.4, it is clear that choosing g is less informative than either f_b or g, and since the expected utility over money is the same for each, it must be that $W(g|\{g\}, S) > W(f_b|\{f_b, f_s\}, S)$ and $W(g|\{g\}, S) > W(f_b|\{f_b, f_s\}, S)$, at price P^*. It is clear from expression (B.13)*, that $W(f_b|\{f_b, f_s\}, S)$ is decreasing in P: the expected utility over money decreases with P, and the self-image term decreases as well, since the informativeness increases as P increases (given our signal structure, since the difference $(u_r(P - r_l) - u_r(r_l - P))$ increases, informativeness increases). As for the term $W(f_s|\{f_b, f_s\}, S)$ from* (B.14)*, it unambiguously increases: the expected utility over money increases with P, and the informativeness decreases as $(u_r(r_h - P) - u_r(P - r_h))$ decreases. Finally, the term $W(g|\{g\}, S) = u_r(0)$, is not affected by P. Hence, there is some P_G (possibly $\overline{P}$) such that $W(f_s|\{f_b, f_g\}, S) > W(g|\{g\}, S)$ and $W(f_s|\{f_b, f_g\}, S) > W(f_b|\{f_b, f_g\}, S)$ for $P \geq P_G$. Using a similar reasoning for P decreasing, it follows that there is some P_B (possibly ($\underline{P}$) such that $W(f_b|\{f_b, f_g\}, S) > W(g|\{g\}, S)$ and $W(f_b|\{f_b, f_g\}, S) > W(f_s|\{f_b, f_g\}, S)$ for $P < P_B$.*

Theorem 2.1: Preference for smaller menus. *Suppose that the agent is strictly doubt-prone. Take any $S = \{M_1, M_2, ...M_n\}$, where $M_1 \subset M_2... \subset M_n$, and where no lottery is degenerate. In addition, $C_M(M_1|\{M_i\}) = C_M(M_2|\{M_j\}) =$ $... = C_M(M_n|\{M_n\})$. Then the agent strictly prefers the smallest menu M_1 in S, i.e. $C_S(S) = M_1$. Furthermore, he chooses lottery $C_M(M_1|S) = C_M(M_1|\{M_i\})$.*

Proof. Assume the agent is strictly doubt-prone, so that u_{p_t} is strictly concave. Since $C_M(M_1|S) = ...C_M = (M_N|S)$, it suffices to show that $C_S(f_l|M_i, S) > C_S(f_l|M_j, S)$, where $M_i \subset M_j$, for any $f_l \in M_i$ (the superscript of f_l is omitted, as it is understood that $f_l = f_l^i = f_l^j$, where $f_l^i \in M_i$ and $f_l^j \in M_j$). That is, it suffices to show that:

$$W(f_l, M_i|S) > W(f_l, M_j|S). \tag{B.16}$$

That is,

$$\begin{aligned} Eu_r(f_l) \quad &+ \sum_{\{r_1^i,...r_{n_i}^i\}\in M_i} u_{pt}\left(p(t|H)\right)\prod_{h=1}^{n_i} p(r_h^i|f_h^i) > \\ Eu_r(f_l) \quad &+ \sum_{\{r_1^j,...r_{n_j}^j\}\in M_j} u_{pt}\left(p(t|H)\right)\prod_{h=1}^{n_j} p(r_h^j|f_h^j) \end{aligned} \tag{B.17}$$

Since the expected utility of f_l is $Eu_r(f_l)$ for both, we focus on u_{p_t}. The rest of the proof will show that the expected utility (using u_{p_t}) of the possible histories generated by the larger menu M_j is less preferred than a mean preserving spread of the possible histories of M_i. By the concavity of u_{p_t}, well-known results in the literature that the EU (using u_{p_t}) of the possible histories of M_i is preferred to the EU of the possible histories of M_j.

We first consider the case for where $M_j \subset M_N$ (i.e. $j < N$, before considering the case $M_j = M_N$ (i.e. $j = N$).

(a) Suppose that $M_j \subset M_N$. Notice that $K_{-i}(S) = K_{-j}(S) = M_N$. In other words, since it is not relevant which unchosen menu the unchosen lotteries come from, the expectation of the signal that the agent receives from lotteries that are not in menu M_i is the same as that of the signal that the agent receives from lotteries that are not in menu M_j. It is straightforward to show that $Eu_{p_t}(f_l|M_i, S) = Eu_{p_t}(f_l|M_i, \{M_i, M_N\})$, and similarly, $Eu_{p_t}(f_l|M_j, S) = Eu_{p_t}(f_l|M_j, \{M_j, M_N\})$. Now consider the menu $M_i' = \{f_1, ...f_{n_i}, f_{n_i+1}\}$, where $f_{n_i+1} \in M_j$ but not in M_i. Let S' be the set $\{M_1, ...M_{i-1}, M_i', M_{i+1}, ..M_N\}$. Since $Eu_{p_t}(f_l|M_i', S') = Eu_{p_t}(f_l|M_i', M_i', M_N)$, it suffices to show that $Eu_{p_t}(f_l|M_i, \{M_i, M_N\}) > Eu_{p_t}(f_l|M_i', \{M_i', M_N\})$, since we can extend the reasoning by appending lotteries to M_i' until we obtain the set M_j, that is, we would prove:

$$[Eu_{p_t}(f_l|M_i, \{M_i, M_N\}) = Eu_{p_t}(f_l|M_i, \{M_i, M_N\})] > Eu_{p_t}(f_l|M_i', \{M_i', M_N\}) > Eu_{p_t}(f_l|M_i'', \{M_i'', M_N\}) > ... > [Eu_{p_t}(f_l|M_j, \{M_j, M_N\}) = Eu_{p_t}(f_l|M_j, \{M_j, S\})] . \tag{B.18}$$

We now show that $Eu_{p_t}(f_l|M_i', \{M_i', M_N\}) > Eu_{p_t}(f_l|M_i', \{M_i', M_N\})$. First, consider *all* histories of form

$$H(r_l|r_1, ...r_{l-1}, r_{l+1}, ...r_n) = (r_1, ...r_l; ...r_{n_i}; f_l, M_i; \{M_i; M_N\},$$

where *only* the $r_l \in X(f_l)$ varies, for a fixed $r_1, ...r_{l-1}, r_{l+1}, ...r_n$. Similarly, consider, all histories of form

$$H'(r_l, r_{n+1}|r_1, ...r_{l-1}, r_{l+1}, ...r_n) = (r_1, ...r_l; ...r_{n_i}, r_{n_i+1}; f_l, M_i'; \{M_i'; M_N\},$$

where this time, *only* the $r_l \in X(f_l)$ and $r_{n+1} \in f_{n+1}$vary, for the same fixed $r_1, ...r_{l-1}, r_{l+1}$, $, ..., r_n$ as for history $H(r_l|r_1, ...r_{l-1}, r_{l+1}, ...r_n)$. Notice that for any history $H'(r_l, r_{n+1}|r_1, ...r_{l-1}, r_{l+1}, ...r_n)$,

$$
\begin{aligned}
&p(t|H'(r_l, r_{n+1}|r_1, ...r_{l-1}, r_{l+1}, ...r_n)) = \\
&p_t\Bigg(1 + \sum_{k\in\{1..n+1\}} b_{lk}\left(\frac{u_r(r_l) - u_r(r_k)}{u_r(\overline{r}) - u_r(\underline{r})}\right)\mathbf{I}_{\{u_r(r_l)\geq u_r(r_k)\}}(\overline{p}^*(r_l, r_k|f_l; f_k)) \\
&-c_S \sum_{k'\in\{1..n\}} \left(\frac{u_r(r_{k'}) - u_r(r_l)}{u_r(\overline{r}) - u_r(\underline{r})}\right) \quad \mathbf{I}_{\{u_r(r_l)<u_r(r_{k'})\}}(\underline{p}^*(r_l, r_{k'}|f_l; f_{k'})) \\
&+ \sum_{g_h\in K_{-i}(S)} \sum_{\tilde{r}\in X(g_h)} \left(\frac{u_r(r_l;) - u_r(\tilde{r})}{u_r(\overline{r}) - u_r(\underline{r})}\right)\big[\mathbf{I}_{\{u_r(r_l)\geq u_r(\tilde{r})\}} b_{K,lh}(\overline{p}^*(r_l, \tilde{r}|f_l; g_h)) \\
&+ \quad \mathbf{I}_{\{u_r(r_l)<u_r(\tilde{r})\}} c_S(\underline{p}^*(r_l, \tilde{r}|f_l; g_h))\big]\Bigg)
\end{aligned} \tag{B.19}
$$

$$
\begin{aligned}
&= p_t\Bigg(1 + \sum_{k\in\{1..n\}} b_{lk}\left(\frac{u_r(r_l) - u_r(r_k)}{u_r(\overline{r}) - u_r(\underline{r})}\right)\mathbf{I}_{\{u_r(r_l)\geq u_r(r_k)\}}(\overline{p}^*(r_l, r_k|f_l; f_k)) \\
&-c_S \sum_{k'\in\{1..n\}} \left(\frac{u_r(r_{k'}) - u_r(r_l)}{u_r(\overline{r}) - u_r(\underline{r})}\right) \quad \mathbf{I}_{\{u_r(r_l)<u_r(r_{k'})\}}(\underline{p}^*(r_l, r_{k'}|f_l; f_{k'})) \\
&+ \sum_{g_h\in K_{-i}(S)} \sum_{\tilde{r}\in X(g_h)} \left(\frac{u_r(r_l;) - u_r(\tilde{r})}{u_r(\overline{r}) - u_r(\underline{r})}\right)\big[\mathbf{I}_{\{u_r(r_l)\geq u_r(\tilde{r})\}} b_{K,lh}(\overline{p}^*(r_l, \tilde{r}|f_l; g_h)) \\
&+ \quad \mathbf{I}_{\{u_r(r_l)<u_r(\tilde{r})\}} c_S(\underline{p}^*(r_l, \tilde{r}|f_l; g_h))\big]\Bigg) + \\
&b_{l(n+1)}\left(\frac{u_r(r_l) - u_r(r_{(n+1)})}{u_r(\overline{r}) - u_r(\underline{r})}\right)\mathbf{I}_{\{u_r(r_l)\geq u_r(r_{n+1})\}}(\overline{p}^*(r_l, r_{n+1}|f_l; f_{n+1})) \\
&-c_S\left(\frac{u_r(r_{n+1}) - u_r(r_l)}{u_r(\overline{r}) - u_r(\underline{r})}\right) \quad \mathbf{I}_{\{u_r(r_l)<u_r(r_{n+1})\}}(\underline{p}^*(r_l, r_{n+1}|f_l; f_{n+1})) \\
&= p(t|H(r_l|r_1, ...r_{l-1}, r_{l+1}, ...r_n)) + \theta(r_l, r_n + 1)
\end{aligned} \tag{B.20}
$$

where

$$
\begin{aligned}
\theta(r_l, r_n + 1) = b_{l(n+1)}&\left(\frac{u_r(r_l) - u_r(r_{(n+1)})}{u_r(\overline{r}) - u_r(\underline{r})}\right)\mathbf{I}_{\{u_r(r_l)\geq u_r(r_{n+1})\}}(\overline{p}^*(r_l, r_{n+1}|f_l; f_{n+1})) \\
-c_S&\left(\frac{u_r(r_{n+1}) - u_r(r_l)}{u_r(\overline{r}) - u_r(\underline{r})}\right) \quad \mathbf{I}_{\{u_r(r_l)<u_r(r_{n+1})\}}(\underline{p}^*(r_l, r_{n+1}|f_l; f_{n+1})).
\end{aligned} \tag{B.21}
$$

This notation is used for simplicity, in that the normalization of of $b_{l(n+1}$ and c_S does depend on the number of lotteries in $\{M_i', M_n$.Now, let

$$t_v(r_l) = \sum_{r_{n+1} \in X(f_{n+1})} p(r_{n+1}|f_{n+1})\theta(r_l, r_{n+1}). \tag{B.22}$$

By (strict) concavity of u_{p_t}, it is clear that :

$$u_{p_t}(p(t|H) + t_v(r_l)) > \sum_{r_{n+1} \in X(f_{n+1})} p(r|f_{n+1})u_{p_t}(p(t|H') = \sum_{r_{n+1} \in X(f_{n+1})} p(r_{n+1}|f_{n+1}) \left(u_{p_t}(p(t|H) + \theta(r_l, r_{n+1})\right). \tag{B.23}$$

for any H and associated H's as described above. It follows that:

$$\sum_{r_l \in X(f_l)} p(r_l|f_l)u_{p_t}(p(t|H) + t_v(r_l)) > \sum_{r_l \in X(f_l), r_{n+1} \in X(f_{n+1})} p(r_l|f_l)p(r_{n+1}|f_{n+1})u_{p_t}(p(t|H'). \tag{B.24}$$

But notice that for any H as defined above (i.e. allowing only the r_l's to vary), the random variable associated with $p(t|H) + t_v(r_l)$ is a mean preserving spread of $p(t|H)$, in the Rothschild and Stiglitz (1970) and Diamond and Stiglitz (1973) 'fat tail' sense. That is, it is clear that the mean is the same, since it means these conditions: first,

$$\sum_{r_l \in X(f_l)} p(r_l|f_l)t_v(r_l) = 0, \tag{B.25}$$

by construction of the signal (the $b_{j(k+1)}$ and c_S are normalized so that this holds, to satisfy the non-manipulability assumption). Second, it is straightforward to show that $t_v(r_l) > t_v(\tilde{r}_l)$ if $r_l > \tilde{r}_l$ (from which it also follows that $t_v(r_l) > 0$ for the maximal $r_l X(f_l)$, and $t_v(r_l) < 0$ for the minimal $r_l \in X(f_l)$. That is, $t_v(r_l)$ is positive for the highest r_l, and decreases monotonically as r_l diminishes (and crosses the 0 for some r_l). This can easily be seen from the observation that $\theta(r_l, r_n) > \theta(\tilde{r}_l, r_n)$ for any r_n

and for any $r_l > \tilde{r}_l$ (the formal proof is trivial), from which it follows that:

$$\sum_{r_{n+1} \in X(f_{n+1})} p(r_{n+1}|f_{n+1})\theta(r_l, r_{n+1}) > \sum_{r_{n+1} \in X(f_{n+1})} p(r_{n+1}|f_{n+1})\theta(\tilde{r}_l, r_{n+1}). \quad \text{(B.26)}$$

It is also clear that for any $r_l > \tilde{r}_l$ and H as defined earlier,

$$u_{p_t}\left(p(t|H(r_l|r_1, ...r_{l-1}, r_{l+1}, ...r_n))\right) > u_{p_t}\left(p(t|H(\tilde{r}_l|r_1, ...r_{l-1}, r_{l+1}, ...r_n))\right). \quad \text{(B.27)}$$

Since $p(t|H) + t_v(r_l)$ is a mean preserving spread of $p(t|H)$, in the Rothschild and Stiglitz(1970) and Diamond and Stiglitz (1973) sense and u_{p_t} is strictly concave, it follows from their well-known results that

$$\sum_{r_l \in X(f_l)} p(r_l|f_l)u_{p_t}(p(t|H)) \geq \sum_{r_l \in X(f_l)} p(r_l|f_l)u_{p_t}(p(t|H) + t_v(r_l)). \quad \text{(B.28)}$$

for any H. Combining (B.24) and (B.28), we have that

$$\sum_{r_l \in X(f_l)} p(r_l|f_l)u_{p_t}(p(t|H)) > \sum_{r_l \in X(f_l), r_{n+1} \in X(f_{n+1})} p(r_l|f_l)p(r_{n+1}|f_{n+1})u_{p_t}(p(t|H'). \quad \text{(B.29)}$$

for any H and associated H'. Finally, summing over all the histories H's, it follows from (B.29) that:

$$\sum_{\{r_1,...r_{n_i}\} \in M_i} u_{p_t}\left(p(t|H)\right) \prod_{h=1}^{n_i} p(r_h^i|f_h^i) > \sum_{\{r_1,...r_{n_{i+1}}\} \in M_i'} u_{p_t}\left(p(t|H')\right) \prod_{h=1}^{n_{i+1}} p(r_h^j|f_h^j)$$

and so $Eu_{p_t}(f_l|M_i, \{M_i, M_N\}) > Eu_{p_t}(f_l|M_i', \{M_i', M_N\})$, and as shown earlier, this suffices to show that (B.17) holds.

(b) We now consider the case where $M_j = M_N$. It is no longer the case $K_{-j}(S) = M_N$, instead $K_{-j}(S) = M_{N-1}$. It suffices to show that $C_S(f_l|M_{j-1}, S) > C_S(f_l|M_j, S)$. It is straightforward to show that $Eu_{p_t}(f_l|M_j, S) = Eu_{p_t}(f_l|M_j, \{M_j, M_{j-1}\})$, and as discussed earlier, $Eu_{p_t}(f_l|M_{j-1}, S) = Eu_{p_t}(f_l|M_{j-1}, \{M_{j-1}, M_j\})$ (since $j = N$). It therefore suffices to show that

$$Eu_{p_t}(f_l|M_j, \{M_j, M_{j-1}\}) < Eu_{p_t}(f_l|M_{j-1}, \{M_{j-1}, M_j\}). \tag{B.30}$$

The signal structure for lotteries in unchosen menus allows the proof of part (a) above can easily be adapted to show that

$$Eu_{p_t}(f_l|M_j, \{M_j, M_{j-1}\}) < Eu_{p_t}(f_l|M_j, \{M_j, M_j\}). \tag{B.31}$$

Using the proof of part (a) directly, it is immediate that

$$Eu_{p_t}(f_l|M_j, \{M_j, M_j\}) < Eu_{p_t}(f_l|M_{j-1}, \{M_{j-1}, M_j\}). \tag{B.32}$$

Combining (B.31) and (B.32), we obtain (B.30), which concludes the proof.

Theorem 2.2: Charactizerization of utility over money. *Function u_r can be uniquely characterized from the agent's choices, up to positive affine transformation.*

Proof. Take any $f = (r_h, 0.5; r_l, 0.5)$ and $f_{r'} = (r', 1)$, for some r'. Consider the set $S' = \{\{f, f_{r'}\}\}$. Let the signal structure parameter associated with choosing lottery f from set $S' = \{\{f, f_{r'}\}\}$ be b'_{12}, and the parameter associated with choosing $f_{r'}$ be b'_{21}. We first show that if the agent is indifferent between lottery f and lottery $f_{r'}$, then it must be that $r' = r_{CE}$, the monetary certainty equivalent of lottery f, in the sense that $u_r(r_{CE}) = 0.5u_r(r_h) + 0.5u_r(r_l)$. First, recall that it was previously shown (see the Dutch lottery example in the appendix) that the agent is indifferent between receiving f and $f_{r'}$, if $r' = r_{CE}$. Now, suppose that $r' > r_{CE}$. Denote the lottery $f_{r_{CE}} = (r_{CE}, 1)$, and let the signal structure parameters associated with choosing lottery $f_{r_{CE}}$ from set $S_{CE} = (f, f_{r_{CE}})$ be b_{21}. Recall that $b_{21} = c_S$, as was previously shown.

Now,

$$b'_{21} = c_S \frac{u_r(r') - u_r(r_l)}{u_r(r_h) - u_r(r')}. \tag{B.33}$$

$$b'_{21} = c_S \frac{u_r(r_h) - u_r(r')}{u_r(r') - u_r(r_l)}. \tag{B.34}$$

Since $r' > r_{CE}$, $u_r(r_h) - u_r(r') < u_r(r_h) - u_r(r_{CE})$, and $u_r(r') - u_r(r_l) > u_r(r_{CE}) - u_r(r_l)$. Hence $b'_{21} < [b_{21} - c_S]$. Now,

$$\begin{aligned} W(f, \{f, f_{r'}\} | \{\{f, f_{r'}\}\}) = 0.5\big(u_r(r_h) + u_r(r_l) + u_{p_t}\left(p_t\left(1 + 0.5c_S \frac{u_r(r') - u_r(r_l)}{u_r(\overline{r}) - u_r(\underline{r})}\right)\right) \\ + u_{p_t}\left(p_t\left(1 - 0.5c_S \frac{u_r(r') - u_r(r_l)}{u_r(\overline{r}) - u_r(\underline{r})}\right)\right)\big) \end{aligned} \tag{B.35}$$

and

$$\begin{aligned} W(f_{r'}, \{f, f_{r'}\} | \{\{f, f_{r'}\}\}) = u_r(r') + 0.5\big(u_{p_t}\left(p_t\left(1 + 0.5c_S \frac{u_r(r_h) - u_r(r')}{u_r(\overline{r}) - u_r(\underline{r})}\right)\right) \\ + u_{p_t}\left(p_t\left(1 - 0.5c_S \frac{u_r(r_h) - u_r(r')}{u_r(\overline{r}) - u_r(\underline{r})}\right)\right)\big). \end{aligned} \tag{B.36}$$

Since $r_{CE} < r'$ by assumption, it must be that

$$[0.5(u_r(r_h) + u_r(r_l)) = u_r(r_{CE})] < u_r(r'). \tag{B.37}$$

Focusing on the u_{p_t} terms, notice that both (B.35) and (B.36) have the same expected value of the term inside the function, p_t. It also the case that

$$u_r(r') - u_r(r_l) > [u_r(r_{CE}) - u_r(r_l) = u_r(r_h) - u_r(r_{CE})] > u_r(r_h) - u_r(r'). \tag{B.38}$$

Hence, using a standard mean preserving spread argument,

$$0.5\big(u_{p_t}\left(p_t\left(1+0.5c_S\frac{u_r(r')-u_r(r_l)}{u_r(\overline{r})-u_r(\underline{r})}\right)\right)+u_{p_t}\left(p_t\left(1-0.5c_S\frac{u_r(r')-u_r(r_l)}{u_r(\overline{r})-u_r(\underline{r})}\right)\right)\big)$$
$$<$$
$$0.5\big(u_{p_t}\left(p_t\left(1+0.5c_S\frac{u_r(r_h)-u_r(r')}{u_r(\overline{r})-u_r(\underline{r})}\right)\right)+u_{p_t}\left(p_t\left(1-0.5c_S\frac{u_r(r_h)-u_r(r')}{u_r(\overline{r})-u_r(\underline{r})}\right)\right)\big). \tag{B.39}$$

Combining (B.37) and (B.39), it must be that
$W(f,\{f,f_{r'}\}|\{\{f,f_{r'}\}\}) < W(f,\{f,f_{r'}\}|\{\{f,f_{r'}\}\})$ if $r' > r_{CE}$. Using a similar argument, it must also be the case that $W(f,\{f,f_{r'}\}|\{\{f,f_{r'}\}\}) > W(f_{r'},\{f,f_{r'}\}|\{\{f,f_{r'}\}\})$. Therefore, since $W(f,\{f,f_{r'}\}|\{\{f,f_{r'}\}\}) = W(f,\{f,f_{r'}\}|\{\{f,f_{r'}\}\})$ when $r = r_{CE}$, the agent is indifferent between f and $f_{r'}$ if and only if $r' = r_{CE}$. Furthermore, we know that if the agent prefers f_r' to f in set S', then $r' > r_{CE}$; similarly, if the agent prefers f to $f_{r'}$, then $r' < r_{CE}$.

We now use an algorithm for finding the value of $u_r(r)$ for any $r \in [\underline{r},\overline{r}]$, noting first that u_r is unique up to positive affine transformation. (Of course, u_{p_t} responds accordingly; that is, $u_r(r)+u_{p_t}(p_t)$ is unique up to positive affine transformation, so that for function $\tilde{u}_r(r)+\tilde{u}_{p_t}(p_t)$, if $\tilde{u}_r = au_r(.)+b$ for some $a > 0$ and b, then it must be that $\tilde{u}_{p_t}(.) = au_{p_t}(.)$, for the same a.) We arbitrarily choose $u_r(\underline{r}) = \underline{u}_r$, and $u_r(\overline{r}) = \overline{u}_r$, where $\underline{u}_r < \overline{u}_r \in \mathbb{R}$. The algorithm is as follows:

Let $\overline{r}_1 = \overline{r}$, $\underline{r}_1 = \underline{r}$, and $f_1 = (0.5,\overline{r}_1;0.5,\underline{r}_1)$, and let $r_{CE,1}$ be the value such that the agent is indifferent between f_1 and $f_{r_{CE,1}}$ in set $S_1 = \{\{f,f_{r_{CE,1}}\}\}$, i.e. $C_M(\{f,f_{r_{CE,1}}\}|S_1) = \{f,f_{r_{CE,1}}\}$. Let $L > 1$ be an integer. We now use the following 'for' loop. Note that $u_r(r_{CE,1} = 0.5\overline{u}_r + 0.5\underline{u}_r$, as shown above.

For $j = 1$ to L (incrementing j by 1 at each loop),

begin

if $r_{CE,j} = r$, then

let $u(r) = u_r(r_{CE,j})$, and exit the loop.

else if $r_{CE,j} < r$ then

let $\overline{r}_{j+1} = \overline{r}_j$,

let $\underline{r}_{j+1} = r_{CE,j}$,

else (i.e., if $r_{CE,j} > r$),

let $\overline{r}_{j+1} = r_{CE,j}$,

let $\underline{r}_{j+1} = \underline{r}_j$,

endif

let $f_{j+1} = (0.5, \overline{r}_{j+1}, 0.5, \underline{r}_{j+1})$,

let $r_{CE,j}$ be the value such that the agent is indifferent between f_j and $f_{r_{CE,j}}$ in set $S_j = \{\{f, f_{r_{CE,j}}\}\}$, implying that $u_r(r_{CE,j}) = 0.5u_r(\overline{r}_j) + 0.5\underline{u}_r(\underline{r}_j)$.

end loop.

If this program does produce a value for $u_r(r)$, then this concludes the proof. Otherwise, since u_r is continuous in $[\underline{r}, \overline{r}]$, the series $u_r(r_{CE,j})$ converges to $u_r(r)$, as L goes to infinity. Since $u_r(r)$ is the limit of this series, this concludes the proof.

Bibliography

[1] M. Abdellaoui. A genuine rank-dependent generalization of the von Neumann-Morgenstern expected utility theorem. *Econometrica*, 70(2):717–736, 2002.

[2] S. Adam, S. Wiggins, P Whyte, M. Bloch, M. Shokeir, H. Soltan, W. Meschino, A. Summers, O. Suchowersky, J. Welch, M. Huggins, J. Theilmann, and M. Hayden. Five year study of prenatal testing for huntington's disease: demand, attitudes, and psychological assessment. *Journal of Medical Genetics*, 30:549–556, 1993.

[3] G. Akerlof and W. Dickens. The economic consequences of cognitive dissonance. *American Economic Review*, 72(3):307–319, 1982.

[4] M. Allais. Le comportement de l'homme rationnel devant le risque: Critique des postulats de l'ecole americaine. *Econometrica*, 21(4):503–546, 1953.

[5] L. Bartels. Uninformed votes: Informational effects in presidential elections. *American Journal of Political Science*, 40(1):194–230, 1996.

[6] D. Bell. Regret in decision making under uncertainty. *Operations Research*, 30:961–981, 1982.

[7] D. Bell and P. Fishburn. Probability weights in rank-dependent utility with binary even-chance independence. *Journal of Mathematical Psychology*, 47:244–258, 2003.

[8] R. Benabou. Ideology. *Journal of the European Economic Association*, 6(2–3):321–352, 2008.

[9] R. Benabou and J. Tirole. Self-confidence and personal motivation. *Quartely Journal of Economics*, 117(3):871–915, 2002.

[10] R. Benabou and J. Tirole. Belief in a just world and redistributive politics. *Quartely Journal of Economics*, 121(2):699–746, 2006.

[11] R. Benabou and J. Tirole. Identity, dignity and taboos: Beliefs as assets. *Working paper*, 2007.

[12] S. Berglas and E. Jones. Drug choice as a self-handicapping strategy in response to noncontingent success. *Journal of Personality and Social Psychology*, 36(4):405–417, 1978.

[13] R. Bodner, , and D. Prelec. Self-signaling and diagnostic utility in everyday decision making. In I. Brocas and J. Carrillo, editors, *The Psychology of Economic Decisions. Vol. 1: Rationality and Well-Being*. Oxford University Press, 2003.

[14] M. Brunnermeier and J. Parker. Optimal expectations. *American Economic Review*, 95(4):1092–1118, 2005.

[15] A. Caplin and J. Leahy. Psychological expected utility theory and anticipatory feelings. *Quartely Journal of Economics*, 116(1):55–79, 2001.

[16] J. Carrillo and T. Mariotti. Strategic ignorance as a self-disciplining device. *Review of Economic Studies*, 67:529–544, 2000.

[17] S. H. Chew, E. Karni, and Z. Safra. Risk aversion in the theory of expected utility with rank dependent probabilities. *Journal of Economic Theory*, 42:370–381, 1986.

[18] O. Compte and A. Postlewaite. Confidence-enhanced performance. *American Economic Review*, 94(5):1536–1557, 2004.

[19] J. Dana, R. Weber, and J. Kuang. Exploiting moral wiggle room: experiments demonstrating an illusory preference for fairness. *Economic Theory*, 33:67–80, 2007.

[20] E. Dekel. An axiomatic characterization of preferences under uncertainty: Weakening the independence axiom. *Journal of Economic Theory*, 40:304–318, 1986.

[21] P. Diamond and J. Stiglitz. Increases in risk and in risk aversion. *Journal of Economic Theory*, 8(3):337–360, 1974.

[22] E. Diecidue and P. Wakker. On the intuition of rank-dependent utility. *Journal of Risk and Uncertainty*, 23(3):281–298, 2001.

[23] D. Dillenberger. Preferences for one-shot resolution of uncertainty and allais-type behavior. *Working paper*, 2008.

[24] J. Dow and S. Werlang. Uncertainty aversion, risk aversion, and the optimal choice of portfolio. *Econometrica*, 60(1):197–204, 1992.

[25] R. Dudley. *Real Analysis and Probability.* Cambridge University Press, New York NY, 2002.

[26] D. Ellsberg. Risk, ambiguity, and the savage axioms. *Quarterly Journal of Economics*, 75(4):643–669, 1961.

[27] D. Ellsberg. *Risk, Ambiguity and Decision*. Garland Publishing, New York NY, 2001.

[28] L. Epstein and S. Zin. Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. *Econometrica*, 57(4):937–969, 1989.

[29] I. Gilboa and D. Schmeidler. Maxmin expected utility with a non-unique prior. *Journal of Mathematical Economics*, 18:141–153, 1989.

[30] I. Good. *Good Thinking: The Foundations of Probability and Its Applications*. University of Minnesota Press, 1983.

[31] S. Grant, A. Kajii, and B. Polak. Intrinsic preference for information. *Journal of Economic Theory*, 83:233–259, 1998.

[32] S. Grant, A. Kajii, and B. Polak. Temporal resolution of uncertainty and recursive non-expected utility models. *Econometrica*, 68(2):425–434, 2000.

[33] F. Gul. A theory of disappointment aversion. *Econometrica*, 59(3):667–686, 1991.

[34] Y. Halevy. Ellsberg revisited: An experimental study. *Econometrica*, 75(2):503–536, 2007.

[35] M. Hemphill. Pretesting for huntington's disease: An overview. *Hastings Center Report*, 3(3):12–13, 1973.

[36] D. Kahneman, J. Knetsch, and R. Thaler. Anomalies: The endowment effect, loss aversion, and status quo bias. *Journal of Economic Perspectives*, 5(1):193–206, 1991.

[37] D. Kahneman and A. Tversky. Prospect theory: An analysis of decisions under risk. *Econometrica*, 47(2):263–291, 1979.

[38] E. Karni. *Decision Making Under Uncertainty*. Harvard University Press, Cambridge MA, 1985.

[39] E. Karni and Z. Safra. Rank-dependent probabilities. *Economic Journal*, 100(401):487–495, 1990.

[40] B. Koszegi. Ego utility, overconfidence, and task choice. *Journal of the European Economic Association*, 4(4):673–707, 2006.

[41] D. Kreps. *Notes on the Theory of Choice*. Westview Press, Boulder CO, 1988.

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

[43] R. Larrick. Motivational factors in decision theories: The role of self-protection. *Psychological Bulletin*, 113(3):440–450, 1993.

[44] G. Loomes and Sugden R. Regret theory: An alternative theory of rational choice under uncertainty. *The Economic Journal*, 92(368):805–824, 1982.

[45] M. Machina. Choice under uncertainty: Problems solved and unsolved. *Economic Perspectives*, 1(1):121–154, 1987.

[46] M. Machina. Dynamic consistency and non-expected utility models of choice under uncertainty. *Journal of Economic Literature*, XXVII:1622–1668, 1989.

[47] G. Northcraft and S. Ashford. The preservation of self in everyday life: The effects of performance expectations and feedback context on feedback inquiry. *Organizational Behavior and Human Decision Processes*, 41:42–64, 1990.

[48] S. Pinker. *The Stuff of Thought*. Penguin Group, New York NY, 2007.

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

[50] J. Quiggin. A theory of anticipated utility. *Journal of Economic Behavior and Organization*, 3:323–343, 1982.

[51] M. Rothschild and J. Stiglitz. Increases risk: I. a definition. *Journal of Economic Theory*, 2:225–243, 1970.

[52] W. Samuelson and R. Zeckhauser. Status quo bias in decision making. *Journal of Risk and Uncertainty*, 1:7–59, 1988.

[53] T. Sarver. Anticipating regret: Why fewer options may be better. *Econometrica*, 76(2):263–305, 2008.

[54] L. Savage. *The Foundations of Statistics*. Dover Publications, New York NY, 1972.

[55] D. Schmeidler. Subjective probability and expected utility without additivity. *Econometrica*, 57(3):571–587, 1989.

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

[57] L. Selden. A new representation of preferences over "certain x uncertain" consumption pairs: The "ordinal certainty equivalent" hypothesis. *Econometrica*, 46(5):1045–1060, 1978.

[58] R. Sugden. An axiomatic foundation for regret theory. *Journal of Economic Theory*, 60:159–180, 1993.

[59] A Tyler, O. Quarrell, L. Lazarou, A. Meredith, and P. Harper. Exclusion testing in pregnancy for huntington's disease. *Journal of Medical Genetics*, 27:488–495, 1990.

[60] E. van Dijk and M. Zeelenberg. When curiosity killed regret: Avoiding or seeking the unknown in decision-making under uncertainty. *Journal of Experimental Social Psychology*, 43:656–662, 2007.

[61] P. Wakker. Separating marginal utility and probabilistic risk aversion. *Theory and Decision*, 36:1–44, 1994.

[62] G. Wu. Anxiety and decision making with delayed resolution of uncertainty. *Theory and Decision*, 46:159–198, 1999.

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

[64] M. Zeelenberg. Anticipated regret, expected feedback and behavioral decision making. *Journal of Behavioral Decision Making*, 12:93–106, 1999.

[65] M. Zeelenberg and J. Beatty. Consequences of regret aversion 2: Additional evidence for effects of feedback on decision making. *Journal of Experimental Social Psychology*, 72:63–78, 1997.

[66] M. Zeelenberg and R. Pieters. Consequences of regret aversion in real life: The case of the dutch postcode lottery. *Organizational Behavior and Human Decision Processes*, 93:155–168, 2004.

[67] M. Zeelenberg, K. van den Bos, E. van Dijk, and R. Pieters. The inaction effect in the psychology of regret. *Journal of Personality and Social Psychology*, 82(3):314–327, 2002.