Probabilistic Certainty Equivalents as an Explanation of Preference Reversals and Other Behavioural Anomalies

Pavlo R. Blavatskyy* and Wolfgang R. Köhler†

University of Zürich

February 2007

(preliminary and incomplete)

Abstract

We develop a model that describes how individuals determine certainty equivalents. Individuals are characterized by a probability measure on the set of preference relations. Each preference relation is monotone, complete, transitive, and convex. Determination of the certainty equivalent is modelled as a sequence of hypothetical binary choices. In each hypothetical choice, the lottery is compared to a monetary outcome for sure, which is adjusted upwards or downwards until the preference for the lottery switches. This model offers a straightforward and intuitive explanation for several well-known phenomena such as the fourfold pattern of risk attitudes, the preference reversal phenomenon and the discrepancy between certainty equivalent and probability equivalent elicitation methods. A re-examination of several experimental studies shows that the model explains why different elicitation methods produce results which differ systematically. Two new experiments test and confirm the predictions of the model.

Key words: loss aversion, expected utility theory, prospect theory, natural experiment

JEL Classification codes: C91, D81

---

* Institute for Empirical Research in Economics, University of Zurich, Winterthurerstrasse 30, CH-8006 Zurich, Switzerland, tel.: +41(0)446343586, fax: +41(0)446344978,
e-mail: pavlo.blavatskyy@iew.unizh.ch

† Corresponding author, Institute for Empirical Research in Economics, University of Zurich, Winterthurerstrasse 30, CH-8006 Zurich, Switzerland, tel.: +41(0)446343588, fax: +41(0)446344978,
e-mail: wolfgang.koehler@iew.unizh.ch

1. Introduction

Most experiments on decision making under risk involve two different types of decision problems: choice between lotteries and the pricing of lotteries (i.e., the elicitation of certainty equivalents). There exists a striking dichotomy between choice and pricing in experiments. Choice tasks rarely involve more than two alternatives. Subjects are presented with the (possibly degenerate) lotteries and are asked to indicate their preference. In contrast, in pricing tasks, subjects have to determine themselves the minimum selling or maximum buying price of a lottery.

Experiments that try to elicit the exact certainty equivalent use payoff schemes to provide incentives under which the price that the subject states implies a compound lottery that the subject chooses.[1] From a theoretical perspective, the pricing task is equivalent to a choice task where subjects choose among a large number of compound lotteries. The difference is that the pricing task is not a binary choice task but involves large number of compound lotteries which are much more complex than those which are used in standard binary choice tasks. But while tasks are similar from a theoretical perspective, experiments phrase the tasks differently. Specifically, we do not know of a single study which elicits certainty equivalents via asking subjects to choose from the large list of compound lotteries which are implied by the incentive scheme that is used. Furthermore, experimental instructions and examples for subjects (usually) do not refer to compound lotteries.

At first glance, the difference between choice and pricing seems to be rather subtle. However, there is plenty of experimental evidence that this difference matters. One of the most prominent examples is the preference reversal phenomenon. Preference reversal refers to a particular discrepancy between choice and pricing. Standard preference reversal occurs if subjects

[1] We discuss in section 2.3 experiments which use a list of binary choices to infer an approximate value for the certainty equivalent.

choose a lottery that gives a moderate payoff with high probability (the P-bet) over a lottery that gives a high payoff with small probability (the \$-bet) but state a higher certainty equivalent for the P-bet than for the \$-bet. Non-standard preference reversal refers to the opposite pattern. The preference reversal anomaly refers to the observation that standard reversals are much more common that non-standard ones.

The observation of preference reversals led to a large literature which attempts to explain the phenomenon (see Seidl 2002 for a review of the literature). Most explanations which have brought forward by economists concentrate on the question whether the elicitation procedure elicits the true certainty equivalent. These explanations retain the assumption that preferences are context-free but relax the independence and/or reduction axioms (Holt, 1986, Karni and Safra, 1987; Segal, 1988). However, Tversky et al. (1990) and Cubitt et al. (2004) use payoff schemes which do not rely on the independence and reduction axioms and still find the same patterns of preference reversals.[2] Loomes and Sugden (1983) show that violations of transitivity as proposed by regret theory can explain preference reversal. Most psychologists view preference reversal as a result of a violation of procedure invariance (e.g., Tversky et al. 1990). The most popular explanations are the scale compatibility hypothesis, the prominence hypothesis, and the anchoring and adjustment effect. These theories model behavior in pricing tasks differently than in choice tasks.

We take a different approach. We go back to the basic notion that a preference relation is a binary relation. Therefore, we consider individuals who can only make binary choices. Hence if an individual faces a complex decision problem, then she splits the decision problem into (hypothetical) binary choice problems.

Consider the statement of a certainty equivalent. To determine the certainty equivalent of

[2] Tversky et al. and Cubitt et al. use an ordinal payoff scheme to provide incentives. Under this scheme, subjects play either the lottery which they have chosen in the binary choice question or the play the lottery for which they stated a higher price.

a lottery, the individual has to compare the lottery to monetary amounts for sure until she determines the amount for sure where her preference switches. Of course, if the subject has deterministic preferences and if she can make an arbitrary large number of comparisons, then the outcome of this process should be exactly the certainty equivalent as it is usually defined by the preference relation.

But although most decision theories are deterministic, there exists overwhelming evidence from experiments that choices involve some degree of randomness. For example, Hey and Orme (1994) and Ballinger and Wilcox (1997) report that 25% and 20.8% of decisions in binary choice tasks are reversed if subjects face the same decision problem for a second time. Additionally, given the short amount of time that subjects take on average to make a decision in experiments, it seems to be a strong assumption that subjects indeed make an arbitrarily large number of comparisons. Therefore, we consider individuals who

- have probabilistic preferences
- make only a finite number of comparisons

We model the determination of certainty equivalents as a sequence of binary decision problems. Consider an individual who is asked to state her certainty equivalent of a lottery. We assume that the individual randomly chooses an amount between the lowest and the highest outcome of the lottery, that she draws a preference relation and then compares the amount to the lottery. Depending on whether the amount or the lottery is preferred, the individual increases or decreases the amount to which the lottery is compared, draws a new preference relation and compares the lottery to the new amount, and so on. The sequence of comparisons stops when the preferred alternative switches (i.e., if the amount is preferred in the last decision problem while the lottery has been preferred in all earlier decision problems or vice versa). We assume that the individual states the average of the last two monetary amounts to which the lottery has been compared as her certainty equivalent.

Given a lottery and a probability measure over preference relations, our model generates a probability distribution over the certainty equivalents that subjects state. In section 2.3 we show that our model explains why different methods to elicit certainty equivalents produce results which differ systematically. In section 3 we show that our model offers a simple and intuitive explanation of the preference reversal phenomenon. Specifically, our model explains why pricing and choice decisions differ and why subjects tend to underprice P-bets and to overprice $-bets. Tversky and Thaler (1990, p.208) conclude that 'the preference reversal phenomenon is an example of a general pattern, rather than a peculiar characteristic of choice between bets'. And while our basic model is about the determination of certainty equivalents, our approach can viewed as a general model of how subjects price alternatives. Therefore, we apply the concept of the stepwise procedure to the determination of probability equivalents and time equivalents. We then show that our model also explains preference reversals in probability and time.

Throughout the paper, we compare the predictions of our model with experimental results of a variety of papers. In section 4, we report the results from an experiment with two new kinds of decision problems which allow to directly test particular assumptions and predictions of our model. Results from the first decision problem show that bounds on the set of possible minimum selling prices have the systematic effect on elicited prices that our model predicts. In the second decision problem, subjects have an incentive to state at every point in time the best 'estimate' of their minimum selling price for a given lottery but are able adjust the price at zero cost. Not only do we observe frequent adjustments as our model predicts. Also the adjustment pattern closely resembles the sequence of hypothetical binary choices from the stepwise procedure in our model. Section 5 concludes.

2. The Model

2.1. Preferences

Let X be a convex and bounded subset of $\mathbb{R}$ with generic element x. A lottery L is defined on a finite subset $X_L \in X$. Lottery $L\left(x_1, q_1; ...x_n, q_n\right)$ delivers outcome x_i with probability $q_i > 0$, $i \in \{1, ..., n\}$ where $\bigcup x_i = X_L$ and $\sum q_i = 1$. A degenerate lottery that delivers outcome x_i with probability one is denoted by x_i. Let $\underline{x}_L = \min\{X_L\}$ and $\overline{x}_L = \max\{X_L\}$. Let $\mathcal{L}$ be the set of lotteries that are defined on X. A preference relation $\succsim_P$is a binary relation on $\mathcal{L}$. Strict preference and indifference are defined in the usual way. We make the following assumptions about $\succsim_P$:

A1: **Completeness**: For all $L_1, L_2 \in \mathcal{L}$, we have $L_1 \succsim_P L_2$ or $L_2 \succsim_P L_1$ or both.

A2: **Transitivity**: For all $L_1, L_2, L_3 \in \mathcal{L}$, if $L_1 \succsim_P L_2$ and $L_2 \succsim_P L_3$, then $L_1 \succsim_P L_3$.

A3: **Monotonicity**: For any $x_i, x_j \in X$ with $x_i > x_j$, we have $x_i \succ_P x_j$.

A4: **Convexity**: For any $L_1, L_2, L_3 \in \mathcal{L}$ and $\alpha \in (0, 1)$, if $L_1 \succ_P L_2$ and $L_3 \succsim_P L_2$, then $\alpha L_1 + (1 - \alpha)L_3 \succsim_P L_2$. If $L_1 \succ_P L_2$ and $L_1 \succsim_P L_3$, then $L_1 \succsim_P \alpha L_2 + (1 - \alpha)L_3$.

Let $\mathcal{P}$ be a set of preference relations $\succsim_P$. Let μ be a probability measure on $\mathcal{P}$. Let $N(L_1 \succsim L_2)$ be the set of preference relations such that $L_1 \succsim_P L_2$. Let EV_L be the expected value of lottery L. Lottery L is risky if X_L contains at least two elements. If there is no risk of confusion, we will refer to a risky lottery as lottery and to a degenerate lottery as amount for sure.

2.2. Choices and the Pricing of Lotteries

Preferences are described by $(\mu, \mathcal{P})$. An individual who faces a (direct) binary choice problem draws a preference relation $\succsim_P \in \mathcal{P}$ according to μ and chooses according to the realized $\succsim_P$. While this definition of binary choice under random preferences is straightforward, it is not obvious how to define certainty equivalents under random preferences and how to interpret the selling or buying prices are elicited in experiments.[3]

We consider individuals who can make only a finite number of binary comparisons. To determine the price of a lottery, the individual has to compare the lottery with different monetary amounts for sure. Obviously, there is no reason to compare the lottery to irrelevant amounts, i.e., to amounts which cannot possibly be the minimum selling/maximum buying price. There are two factors which determine the set of possible prices: preferences and experimental procedures (e.g., see section 4.1). A monetary amount is a possible price of the lottery

- if it possible to state this price given the elicitation procedure and
- if the individual cannot exclude the amount as price by logical deduction.

Definition: The *relevant set* S_L is the convex hull of the set of possible prices of lottery L.[4]

Individuals are endowed with a probability measure over preference relations $(\mu, \mathcal{P})$. We consider individuals who do not know $(\mu, \mathcal{P})$ but know that all preference relations in $\mathcal{P}$ satisfy A1-A4. Since individuals know that all preference relations satisfy convexity, they know that they never prefer some amount less than $\underline{x}_L$ to the lottery and never prefer the lottery to some amount larger than $\overline{x}_L$. In this section, we consider the general case where there are no experimental restrictions on the set of possible prices. Hence, the relevant set of prices is $S_L = [\underline{x}_L, \overline{x}_L]$. Individuals draw some amount $x \in S_L$ at random, draw a preference relation and compare x to the lottery. Once an individual learns whether she prefers x or the lottery, she has additional information about her preferences. Hence, we assume that the individual increases or decreases the amount by $\Delta > 0$ (depending on whether the amount or the lottery

[3] Note that we do not assume that preference relations are continuous. Hence it is possible that for some lottery $L \in \mathcal{L}$ and some $\succsim_P \in \mathcal{P}$ there exists no c_P such that $c_P = \{x \in \mathbb{R} | x \succsim_P L\} \cap \{x \in \mathbb{R} | x \precsim_P L\}$.

[4] The assumption that the relevant set is the convex hull instead of the set of possible prices is made purely for technical reasons.

is preferred in the first comparison) and makes a new comparison, and so on. The price that an individual states is denoted by CE_L.

Formally: Consider an individual who is asked to price lottery L in an incentive compatible way. Individuals use the following procedure P to determine the price:

Procedure *P*:

Step 1: The individual draws an amount $x \in S_L$ at random.

Step 2: The individual draws $\succsim_P \in \mathcal{P}$ according to probability measure μ and compares x and L.

Step 3: If $x \sim_p L$, then $CE_L = x$. If $x \succ_P L$, then Step 2 is repeated where x is replaced by $\max\{x - \Delta, \inf S\}$. If $L \succ_P x$, then Step 2 is repeated where x is replaced by $\min\{x + \Delta, \sup S\}$. The sequence of binary comparisons ends if the preferred alternative switches. In this case, CE_L is the average of the last two amounts to which the lottery has been compared.

If individuals use procedure P, then the price CE_L that an individual states can be described as random variable whose distribution depends on the probability measure over preference relations and on the step-size Δ. Until now, we made only weak assumptions on the preferences. To make predictions about the distributions of elicited prices and to relate or results to the findings in the experimental literature we need to impose two technical assumptions on the distribution of preferences and on the relation between preferences and step-size.

Assumption 5: $\mu(N(EV_L - \delta \succsim L)) = \mu(N(L \succsim EV_L + \delta))$ $\forall \delta$ with $EV_L \pm \delta \in X$.

Assumption 5 implies that individuals are 'on average' risk-neutral, i.e., the probability that an individuals prefers the expected value over the lottery is the same as the probability that the lottery is preferred over its expected value. While this is a fairly strong assumption, it simplifies the exposition of the formal arguments and the description of the intuition behind our results. When we discuss the relation of our model to the observed violations of Expected

Utility Theory we argue that our qualitative results do not change as long as individuals are not too risk-averse relative to the degree of randomness of their preferences.

If $EV_L < \frac{\underline{x}_L+\overline{x}_L}{2}$ let n be the unique integer such that $\underline{x}_L + (n + 0.5)\Delta \leq EV_L$ and $\underline{x}_L+(n+1.5)\Delta > EV_L$. If $EV_L > \frac{\underline{x}_L+\overline{x}_L}{2}$ let n be the unique integer such that $\overline{x}_L-(n+0.5)\Delta \geq EV_L$ and $\overline{x}_L - (n + 1.5)\Delta < EV_L$.

Assumption 6:

$$\prod_{j=1}^{n} \mu(N(L \quad \succsim \quad \underline{x}_L + j\Delta)) < 0.5 \text{ if } EV_L < \frac{\underline{x}_L + \overline{x}_L}{2}$$
$$\prod_{j=1}^{n} \mu(N(L \quad \precsim \quad \overline{x}_L - j\Delta)) < 0.5 \text{ if } EV_L > \frac{\underline{x}_L + \overline{x}_L}{2}$$

Assumption 6 is a joint condition on the distribution of preferences and the step-size. Assumption 6 guarantees that preferences are sufficiently random relative to the step-size.

To see the intuition behind Assumption 6, consider an individual with deterministic preferences who uses the method described above to price a lottery. The larger the step-size Δ, the more random is CE_L. However, this randomness has nothing to do with random preferences but with the assumption that CE_L is equal to the average of the last two amounts to which the lottery has been compared. The randomness of the stated price CE_L is affected by two factors. The first is the randomness of the preferences and the second is the step size. Assumption 6 guarantees that the randomness due to the second factor is small relative to the randomness due to the first factor.

Consider an individual who is asked to price a lottery with expected value close to its lowest possible outcome. As argued above, before the individual draws a preference relation and compares the lottery to some fixed amount, she has no information about which amount makes her indifferent to the lottery. Therefore, she picks some amount between the lowest and the highest outcome of the lottery at random and compares it to the lottery. Since the expected value of the lottery is close to its lowest outcome, this 'starting amount' is likely to be larger than the expected value. It turns out that the individual is not only likely to start at an amount larger than the expected value but she is also likely to end up at an amount that is larger than the expected value.

Theorem 1. *If preferences satisfy A1 - A6 and individuals use the procedure P, then for any lottery $L \in \mathcal{L}$ with $EV_L < \frac{\underline{x}_L + \overline{x}_L}{2}$ the following holds:* $\Pr\{CE_L > EV_L\} > \frac{1}{2}$. *If $EV_L > \frac{\underline{x}_L + \overline{x}_L}{2}$, then* $\Pr\{CE_L < EV_L\} > \frac{1}{2}$.

The proof is presented in the Appendix.

2.3. Elicitation procedures

A variety of methods are used in experiments to elicit certainty equivalents. According to expected utility theory, all elicitation methods are equivalent. However, it is well-known that different elicitation methods produce results which differ systematically. One advantage of our theory is that it allows to predict how the details of the elicitation procedure affect the distribution of elicited certainty equivalents. We briefly describe the most commonly used methods and point out how they relate to the procedure described above.

2.3.1. Becker-DeGroot-Marschak Method

Under the Becker-deGroot-Marschak (1964) method (BDM), certainty equivalents are elicited through a second price auction. Individuals are endowed with a lottery and are asked to state the lowest price for which they would sell the lottery. After this, a random device generates a price. If this price exceeds the stated price, the individual has to give up the lottery in exchange for the randomly generated price. Otherwise she plays the lottery. This is exactly the situation that we outlined above, i.e., the procedure to determine the certainty equivalent (the lowest selling price) is left to the individual. In terms of our model, the experimenter does not specify the sequence of binary comparisons and doesnot specify the step size Δ.

Lemma 1 *Suppose that preferences satisfy A1 - A6 and that individuals use the procedure P. If certainty equivalents are elicited via the Becker-deGroot-Marschak method, then for any lottery* $L \in L$ *with* $EV_L < \frac{\underline{x}_L + \overline{x}_L}{2}$ *the following holds:* $\Pr\{CE_L > EV_L\} > \frac{1}{2}$. *If* $EV_L > \frac{\underline{x}_L + \overline{x}_L}{2}$, *then* $\Pr\{CE_L < EV_L\} > \frac{1}{2}$.

2.3.2. Ascending/Descending Auctions

Another method to elicit certainty equivalents is via a second price auction with an ascending or descending price. Under this format, the price is gradually increased respectively decreased and individuals indicate the first price at which they are willing to buy respectively sell the lottery. In terms of our model, this elicitation procedure fixes a constant step size Δ where Δ is the increment by which price is increased or decreased in each step. Additionally, the elicitation procedure determines the sequence of binary comparisons. Usually, an ascending auction starts at $\underline{x}_L$ and a descending auction starts at $\overline{x}_L$. Of course, if individuals compute their minimum selling or maximum buying price before the auctions starts, we are in same situation as in the BDM-method. Therefore, we consider the case where individuals do not compute their minimum selling or maximum buying price before the auctions starts.

Lemma 2 *Suppose that preferences satisfy A1 - A6.* For any risky lottery $L \in \mathcal{L}$, *if the auction is ascending and starts at* $\underline{x}_L$, *then* $\Pr\{CE_L < EV_L\} > \frac{1}{2}$. *If the auction is descending and starts at* $\overline{x}_L$, *then* $\Pr\{CE_L > EV_L\} > \frac{1}{2}$.

An ascending (descending) auction is equivalent to individuals who follow procedure P except that the first amount to which the lottery is compared is fixed at $\underline{x}$ ($\overline{x}$). Since preferences are random, the probability that the certainty equivalent is relatively small is larger in an ascending auction than in a descending auction.

A variety of elicitation procedures are closely related to the ascending/descending auction. For example, in Butler and Loomes (2007), subjects are endowed with a lottery. Subjects face a sequence of choices between the lottery and a sure amount of money. For half of the subjects the sequence started with the sure amount being equal to the lowest outcome. In each question, the amount was gradually increased. For the other half, the sequence started at the highest outcome and the amount was gradually lowered. Butler and Loomes found a huge difference between the certainty equivalents under the two treatments. As predicted by Lemma 2, elicited certainty equivalents are much smaller for subjects who start at the lowest outcome than for subjects who start at the highest outcome. Schkade and Johnson (1989) asked subjects to indicate their minimum selling price for a lottery by using the mouse to move a pointer along a price scale. They manipulated the starting points by choosing the initial position of the pointer to be either close to the lower or the upper bound of the response scale. They found a large and significant effect of the starting point. For high starting points, average prices were significantly higher. It is a natural assumption that at least some subjects use the amount that is equal to the initial position of the pointer as the first amount that is compared to the lottery. If subjects use procedure P and some subjects use the initial position of the pointer as the first amount that is compared to the lottery, then the expected certainty equivalent that subjects state is larger if the initial position of the pointer is close to the upper end of the possible responses.

2.3.3. List of Binary Choices

The BDM-method and the ascending/descending auction formats are designed to elicit the exact certainty equivalent. Other elicitation methods compute an approximate value of the certainty equivalents from a sequence of observed choices but do not impose the order in which choices are made. Under these methods, individuals face the lottery and a list of monetary amounts. For each amount, individuals have to indicate whether they prefer the amount for sure or the lottery. The certainty equivalent is usually defined as the mid point of the lowest amount that is chosen over the lottery and the highest amount which is not chosen over the lottery. This elicitation method explicitly asks individuals to make a series of binary choices similar to the ones described in our cognitive model above. In terms of our model, this elicitation method fixes the step size Δ but does not specify the order in which binary comparisons are made. Most experiments use computer programs to prevent inconsistent choices. I.e., if the individual indicates that she prefers $\widetilde{x}$ over L then the program automatically fills in that all $x > \widetilde{x}$ are preferred to L and vice versa if the individual indicates that she prefers L over $\widetilde{x}$.

Suppose that the monetary amounts to which the lottery is compared are equally spaced with distance Δ, that EV_L is one of the monetary amounts on the list, that $\mu(N(L \sim EV_L)) = 0$, and that a computer program prevents inconsistent choices.

Lemma 3 *If preferences satisfy A1 - A6, subjects start at random with one of the choice problems and then solve adjacent choice problems, then* $\Pr\{CE_L > EV_L\} > \frac{1}{2}$ if $EV_L < \frac{\underline{x}_L + \overline{x}_L}{2}$ and $\Pr\{CE_L < EV_L\} > \frac{1}{2}$ if $EV_L > \frac{\underline{x}_L + \overline{x}_L}{2}$.

For example, in the seminal paper of Tversky and Kahneman (1992), the computer screen displays a lottery and seven sure amounts that are logarithmically spaced between the lowest and the highest outcome of the lottery. Subjects have to indicate a preference between each of the amounts and the lottery. A computer program prevents inconsistent choices. After subjects state their preferred alternatives, they face a new set of outcomes, linearly spaced between a value 25% higher than the lowest amount accepted and a value 25% lower than the highest amount rejected. The certainty equivalent is defined as the average between the highest amount that the subject is not willing to choose over a risky lottery and the lowest amount that the subject prefers to the lottery. A decision is classified as risk-averse (risk-seeking) when the certainty equivalent of a lottery is below (above) the expected value of the lottery.

Tversky and Kahneman (1992) find that 78% of choices are risk seeking, when the probability of a gain in a two-outcome lottery is less or equal to 0.1, and that 88% of choices are risk averse, when the probability of a gain is greater or equal to 0.5. Additionally, 80% of the choices are risk averse when the probability of a loss in a two-outcome lottery is less or equal to 0.1 and 87% of choices are risk seeking choices when the probability of a loss is greater or equal to 0.5. Except for three subjects, this phenomenon is also observed on the individual level (e.g. Table 4 in Tversky and Kahneman, 1992). Tversky and Kahneman (1992) refer to risk seeking over unlikely gains and probable losses and risk aversion over probable gains and unlikely losses as fourfold pattern of risk attitudes.

In terms of procedure P, the experimental design in Tversky and Kahneman (1992) effectively imposes a fixed step size by presenting a list of sure monetary amounts to the subjects. Subjects are not asked to state their certainty equivalent, instead the certainty equivalent is defined as the average of the two monetary amounts where the preference switches. Consider a subject who randomly starts with one of the decision problems, i.e., who chooses at random one of the amounts on the screen and compares it to the lottery. If the subject then solves adjacent decision problems and if the step size (i.e., the difference between the monetary amounts) and the preferences of the subject satisfy Assumption 5 and 6, we immediately obtain the result from Lemma 3. For the two-outcome lotteries used in Tversky and Kahneman (1992), $EV_L < \frac{\underline{x}_L + \overline{x}_L}{2}$ implies that the lottery delivers either a gain with small probability or a loss with high probability and vice versa for $EV_L > \frac{\underline{x}_L + \overline{x}_L}{2}$. Thus, the fourfold pattern of risk attitudes emerges as a consequence of the experimental design.

3. Preference Reversal

In standard economic theory, preferences are viewed as independent of the task that a subject faces. Hence observed choices should allow direct inference about the underlying preferences. Yet many studies of the preference reversal phenomenon provide welldocumented evidence that subjects attitudes towards lotteries depend on the way how they are elicited. Studies of preference reversal involve pairs of lotteries of similar expected value. In each pair, one lottery (the \$-bet) yields a relatively high outcome with a low probability. The other lottery (the P-bet) yields a modest outcome with a probability close to one. Studies of preference reversal involve two different procedures of eliciting a preference relation. Under one procedure, subjects are asked to choose between lotteries. Under the other procedure, the minimum selling price for each lottery is elicited.

Standard (regular) preference reversal refers to the observation that a subject chooses the P-bet over the \$-bet in a direct binary choice but states a higher minimum selling price for \$-bet than for P-bet. Non-standard (irregular) preference reversal refers to the observation that a subject chooses \$-bet over P-bet in a direct binary choice but states a higher minimum selling price for P-bet than for \$-bet. If choices and elicited certainty equivalents are probabilistic (because there is some error in decision making or because preferences are random), then this should result in preference reversals. The preference reversal anomaly refers to the observation that standard reversals are much more common that non-standard ones.

Bostic et al. (1990) conducted 2 experiments with 41 respectively 21 students who 1) choose between the P-bet and the \$-bet; 2) state minimum selling prices for the P-bet and the \$-bet under the Becker-DeGroot-Marschak (BDM) method; and 3) make a sequence of binary choices between monetary amounts and the P-bet respectively the \$-bet. In the last elicitation procedure, half of subjects started at a monetary amount well above the expected value of the lottery while the other half started at a monetary amount well below the expected value. The amount was subsequently increased or decreased, depending on whether the subject chose the lottery or the amount for sure. In the first experiment, monetary amounts were adjusted by a fixed amount and the procedure was continued until a subject's choice had changed 6 times. In the second experiment, a more elaborate rule was used to adjust the monetary amounts.[5]

Bostic et al. (1990) found that 41% (46% in experiment 2) of choices between the lotteries were inconsistent with the certainty equivalents that were elicited through the BDM-method. In 70% (81%) of all cases in which the P-bet is chosen over the \$-bet, the elicited certainty equivalent for the \$-bet is higher (where certainty equivalents are elicited with the BDM-method). When the \$-bet is chosen, the reversal rate is only 12.4% (3%). The incidence of preference reversals is lower (28% in experiment 1 and 39% in experiment 2) if direct binary choices are compared with the certainty equivalents that are elicited through a sequence of binary comparisons. While the standard reversal is still observed more frequently, the difference between the frequencies is smaller than if choices are compared to certainty equivalents that are elicited under the BDM-method. For the P-bet, the means of the certainty equivalents elicited through the BDM-method and a sequence of binary comparisons are nearly identical. They are typically slightly below the expected value of the lottery. However, for the \$-bet, the mean of the certainty equivalents that are elicited through the BDM-method is significantly above the expected value of the lottery while the mean of the certainty equivalents that are elicited through a sequence of binary comparisons is close to the expected value of the lottery (being above the expected value only in one lottery pair and below the expected value in two lottery pairs).

Note that our assumptions on preferences do not imply any predictions about which lottery

[5] Bostic et al. refer to certainty equivalents that are elicited by the BDM-method as Judged Indifference Points (JIP) and to certainty equivalents that are computed as the result of observed choices as Choice Indifference Points (CIP).

is chosen more frequently. For the certainty equivalents that Bostic et al. elicit with the BDM-method, Lemma 1 implies that $\Pr\{CE_L > EV_L\}_{\$\text{-bet}} > \frac{1}{2}$ and that $\Pr\{CE_L > EV_L\}_{\text{P-bet}} < 0.5$. Bostic et al. use lottery pairs where the expected values of the P-bet and the \$-bet are nearly the same. If expected values are equal, it follows immediately from Lemma 1 that the percentage of reversals given that P-bet $\succ$ \$-bet is higher than the percentage of reversals given that \$-bet $\succ$ P-bet. Hence our model predicts that standard preference reversal is observed more frequently than the non-standard reversal if certainty equivalents are elicited with the BDM-method. Of course, if preferences are sufficiently random, this holds also if the expected values of the P-bet and \$-bet differ slightly.

There exists a large number of papers which elicit certainty equivalents via the BDM-method. These papers consistently replicate the results of Bostic et al. The additional feature in Bostic et al. is that they also elicit certainty equivalents via the sequence of binary choices described above. Hence the remaining puzzle is why the frequency of standard and non-standard preference reversals is lower when Bostic et al. elicit the certainty equivalent through a sequence of binary comparisons. Note that under the procedure in Bostic et al. the sequence of comparisons is determined by the experimenter and is increasing respectively decreasing for half of the subjects. If the starting amounts are spaced equally distant from the expected value, then this procedure is roughly equivalent to the ascending or descending auction that is described in Lemma 2 (except for the elaborate definitions of certainty equivalents as the average of six switching points and the fact that amounts can increase as well as decrease in later comparisons). Then Lemma 2 predicts that we should see overpricing if the starting value is above the expected value and underpricing if the starting value is below the expected value. On average, there should be less systematic over- or underpricing. And indeed, Bostic et al. find that certainty equivalents that are derived from choices are on average very close to the expected value. But if there is less systematic over- or underpricing, our theory predicts that we should observe less frequent preference reversals than when certainty equivalents are elicited with the BDM-method which is exactly what Bostic et al. observe.

A series of papers (Holt 1986, Karni and Safra 1987, and Segal 1988) show that the BDM-method does not elicit the true certainty equivalents if independence and/or the reduction-of-compound-lotteries axiom do not hold. Therefore, Tversky et al. (1990) use an ordinal payoff scheme to control whether violations of the independence and/or the reduction axiom are indeed the causes of preference reversals.[6] Tversky et al. investigate whether preference reversal is the result of intransitive preferences or of procedure invariance[7]. The P-bets and the \$-bets are constructed so that the expected value of the \$-bet is slightly higher. Additionally, the expected value of the \$-bet is in most pairs higher than the gain in the P-bet. Subjects face two tasks: first they state minimum selling prices for the \$-bet and the P-bet and then they choose between the two lotteries, and between some amount X for certain and each of the lotteries. The fixed amount X is close to the smaller of the two expected values of the lotteries. As in nearly all studies, standard preference reversals occur much more frequently (45% of all choice patterns) that non-standard reversals (4%). The P-bet was chosen over the \$-bet in 74% of all direct binary comparisons but the minimum selling price for the P-bet is only in 34% of all cases higher than the minimum selling price for the \$-bet. Since the expected value of the \$-bet is in most pairs higher than the gain in the P-bet it is no surprise that the minimum selling price for the \$-bet is in most cases higher. Tversky et al. conclude that only 10% of the standard preference reversals can be explained by intransitive preferences and that the remaining 90% of reversals are the result of a failure of procedure invariance, i.e., of

[6] Tversky et al. use the following ordinal payoff scheme: 15% of students are randomly selected to play one of the bets for real. To determine the bet, one of the pairs of lotteries is selected. Then it is determined whether the choice or the pricing task determines the payoffs. If the choice task is selected, subjects play the bet which they preferred in the direct binary choice question. If the pricing task is selected, subjects play the bet for which they stated a higher cash equivalent.

[7] Tverky et al. classify preferences as intransitive if the choice patterns are intransitive. I.e., if the P-bet is preferred to the \$-bet and the amount X is preferred to the P-bet and the \$-bet is preferred to X and vice versa if \$-bet $\succ$ P-bet.

the overpricing of the \$-bet, the underpricing of the P-bet, and the combination of the two[8]. Since we assume that preferences are random, our model - like other random utility models - predicts that some fraction of observed choice patterns violates transitivity even though we assume that each preference relation is transitive. The contribution of our model is that it additionally explains the 90% of preference reversals which are caused by the overpricing of the \$-bet and/or the underpricing of the P-bet.

4. Experimental Design

We conducted an experiment to test particular assumptions and predictions of our model. Here we report the results of a pilot study. The main experiment will be conducted in March 2007.

The experiment consisted of 72 decision problems that were presented in random order. There were five types of decision problems: binary choice and four versions of lottery pricing. Lotteries are described in table #. Lotteries 1-3 are the same as in Harbaugh et al. (2003), except that payoffs are in Swiss Francs (CHF) and multiplied by 3.5. Lotteries 4-15 are the same as in lottery set I from Tversky et al. (1990), except that payoffs are in Swiss Francs (CHF) and multiplied by 10. All lotteries have only two outcomes and the lowest outcome is zero. Risky lotteries were described and subsequently played out in terms of the number of red and black cards in a box that contains 100 cards. We used a bar to represent the proportion of red and black cards. If the subject drew a black card, she would receive zero, if she drew a red card, she would receive the high outcome of the corresponding lottery.

In *choice tasks*, subjects were presented with two lotteries (including degenerate lotteries). Subjects could choose one of the lotteries or declare indifference by pressing “I do not care”.

[8]The case for intransitive preferences as explanation of preference reversal becomes even weaker if one takes into account that not all subjects who exhibit intransitive choices also exhibit preference reversals. Only 77% of cases where subjects choices are intransitive and the subject prefers the P-bet to the \$-bet result in preference reversals. And only 34% of cases where subjects choices are intransitive and the subject prefers the \$-bet to the P-bet result in preference reversals (compare Table 1 in the appendix of Tversky et al. 1990).

If a subject declared indifference, one of the lotteries was chosen at random to determine the payoff of the subject. In six choice problems, subjects had to choose from the pairs of P-bets and $-bets used by Tversky et al. (1990). In the remaining 15 choice problems, subjects had to choose between a risky lottery and a monetary outcome for certain. Subjects were asked to choose between the six P-bets used by Tversky et al. (1990) and their expected value and between lottery 1-3 and the expected value of the lottery. In addition, the subject were also asked to choose between the six $-bets from Tversky et al. (1990) and the expected value of the corresponding P-bet.

The experiment included four different lottery pricing tasks. In all pricing tasks, subjects were endowed with a lottery.

In a *standard BDM-task*, the subjects were endowed with a lottery (we used lotteries 1 to 15). Subjects were asked to enter their minimum selling price for the lottery. If one of these tasks was selected to determine the earnings, a random price was drawn from the interval between zero and the highest outcome of the lottery. If the subject stated a price above the price that was drawn, she would play the lottery. Otherwise she would sell the lottery and receive whatever price was drawn.

We used the same lotteries (L1 to L15) in the *time-restricted BDM task*. The set-up and the procedure to determine the earnings was the same as in the standard BDM-task. The only difference was that subjects faced a time limit of 10 seconds to enter a price. Before the lottery appeared on the computer screen, subjects received a short notice which told them that the next question had a time limit. Time started running when the lottery appeared on the screen. A little counter displayed how much time was left. Subjects, who did not enter a price before the time limit expired, received zero if this question was selected to determine the earnings in the experiment.

We used the same lotteries (L1 to L15) in the *restricted BDM task*. In a restricted BDM-task, subjects could only enter prices from some interval. When the lottery appeared on the screen, subjects were told from which interval they could enter prices. All intervals were symmetric around the expected value of the lottery and included either zero (for $-bets) or the highest outcome (for P-bets). If one of these tasks was selected to determine the earnings, a random price was drawn from the interval. If the subject stated a price above the price that was drawn, she would play the lottery. Otherwise she would sell the lottery and receive whatever price was drawn. Subjects received a detailed explanation why it was optimal for them to choose a price equal to the lower bound if their minimum selling price was below the lower bound and similarly for the upper bound. Since lotteries 9 and 15 involve 50%-50% chances, the restricted BDM-task is for these lotteries equivalent to the standard BDM-task.

Lotteries 16 to 21 were used in the *radioline-task*. In the radioline-task, the computer screen displayed the lottery and a list of 30 prices spaced equally apart between zero and the highest outcome of the lottery. Subjects were asked to indicate their minimum selling price by using their mouse to click one of the prices on list. Subjects were told that the question is terminated after some randomly determined time (between 10 and 30 seconds). If subjects did not click a price before the question was terminated, they received a payoff of zero if the question was selected to determine their payoffs. Subjects could change the price as often as they wanted. The computer program recorded the order in which prices are clicked. The last price that a subject clicked was used to determine her payoff. I.e., if one of these tasks was selected to determine the earnings, a random price was drawn from the interval between zero and the highest outcome of the lottery. If the last price that the subject entered was higher than the price that was drawn, she would play the lottery. Otherwise she would sell the lottery and receive whatever price was drawn. Before the lottery appeared on the screen, subjects received a short notice which told them that the next question was radioline-task. The notice explained the type of question, that they would receive a payoff of zero if they did not enter a price before the question was terminated, and that they could change the price at zero cost.

We used a random lottery incentive scheme. At the end of the experiment each subject drew a card from a box with cards numbered from 1 to 72. The number on the card determined the question which was used to compute the payoff of the subject. If the subject had to play a risky lottery, she had to draw a card from the corresponding box. The color of the card that was drawn determined the payoff.

The experiment was conducted on December 15, 2006 in the experimental laboratory of the Institute for Empirical Research in Economics at University of Zürich. Thirty undergraduates (17 male and 13 female) from a variety of majors participated in the experiment. The average age was 21. At the beginning of the experiment, subjects received a copy of the instructions (a translation can be found in the Appendix). Instructions included screenshots for the different tasks. Additionally, the experimenter read aloud the instructions. The experiment lasted about 45 minutes (plus 30 minutes to explain the instructions).[9] Subjects received a 10 CHF show-up fee and whatever they earned in the experiment. Average earnings were 41.4 CHF (approx. \$33.2 or €25.6). The lowest earning was 10 CHF, the highest was 113.4 CHF. At the end of the experiment, the subjects completed a short socio-demographic questionnaire.

4.1. The Effect of Bounds on the BDM Elicitation Method

In the procedure described in section 2.2, subjects start to determine the price of a lottery by drawing a 'starting value' from the relevant set of prices S. The purpose of the restricted BDM-task is to test whether our assumption that subjects draw starting values from S is correct. To test this assumption, we compare the elicited prices from the standard BDM-task with the elicited prices from the restricted BDM-task. In the standard BDM-task, subjects can enter any minimum selling price. Since preference relations are convex, the relevant set from

[9] Subjects did not have to wait until others finished before a new question appeared.

which starting values are drawn in the standard BDM-task is $S_L = [\underline{x}_L, \overline{x}_L]$. In the restricted BDM-task, subjects can state only prices in some interval I_L. Since intervals I are chosen such that I_L is a subset of $[\underline{x}_L, \overline{x}_L]$, from the definition of the relevant set of prices follows that $S_L = I_L$ in the restricted BDM-task.

This generates a clear prediction about the distribution of prices that subjects state in the restricted BDM-task compared to the distribution in the standard BDM-task. Recall that the intervals I from which prices can be chosen are centered around the expected value of the lottery. Consider the standard BDM-task and suppose that $\underline{x}_L < \inf I_L$. If the starting value is smaller than $\inf I_L$, then the probability that the stated certainty equivalent is below or equal $\inf I_L$ is larger than if the subject draws a starting value at random from I_L. Similar, if the starting value is higher than $\sup I_L$, then the probability that the stated certainty equivalent is higher or equal $\sup I_L$ is larger than if the subject draws a starting value at random from I_L. Hence, our theory predicts that the percentage of subjects who state prices equal to the upper (lower) bound in the restricted BDM task is smaller than the percentage of subjects who state a minimum selling price equal or above (equal or below) the bound in the standard BDM task.

If it would instead be true that the distribution of starting values in the restricted BDM task is the same as in the standard BDM task, then the distribution of minimum selling prices would be the same in both tasks. Since subjects can not state prices outside the interval in the restricted BDM task, they state the lower or upper bound of the interval if their minimum selling price lies outside the interval. Hence the percentage of subjects which state the lower bound in the restricted BDM task would be equal to percentage of subjects who state a minimum selling price equal to or below the lower bound in the standard BDM task and similarly for the upper bound.

Similarly, any procedure invariant theory (PI-theory) predicts that percentage of subjects which state the lower bound in the restricted BDM task would be equal to percentage of subjects who state a minimum selling price equal to or below the lower bound in the standard BDM task and similarly for the upper bound. Note that this prediction holds for all PI-theories, regardless of whether they are stochastic or deterministic.

There were 30 subjects who evaluated Lotteries 1 to 15 both in the standard BDM task and the restricted BDM task. In the restricted BDM task, subjects could only state prices in some interval I_L which was centered around the expected value of the lottery. Intervals were chosen such that either the lower bound of I_L was equal to $\underline{x}_L$ or the upper bound equal to $\overline{x}_L$. Since lotteries 9 and 15 involved a 50% chance to win $\overline{x}_L$, bounds did not bind for these lotteries. Table 1 summarizes the results. The fifth row lists the number of answers which were equal to or outside the bounds in the standard BDM-task. The sixth row lists the number of answers which were equal to the bound in the restricted BDM-task.[10]

Table 1

Lottery	1	2	3	4	5	6	7	8	10	11	12	13	14
Prob. win	0.1	0.4	0.8	0.97	0.31	0.81	0.19	0.94	0.89	0.11	0.94	0.39	0.92
Payoff win	70	70	70	40	160	20	90	30	40	400	25	85	20
Interval	0-14	0-56	42-70	37.6-40	0-99.2	12.4-20	0-34.2	26.4-30	31.2-40	0-88	22-25	0-66.3	16.8-20
Number equal or outside bounds	18	0	5	13	5	5	8	13	11	12	11	0	8
Number equal bounds	8	0	4	4	1	2	1	7	3	4	12	0	2

In the standard BDM task, 28% of minimum selling prices are equal to or outside the bounds. However, only 12.3% of stated prices are equal to the bounds in the restricted BDM task. The data clearly show that a restriction on the set of prices has a strong effect on the way how minimum selling prices are determined. Specifically, the data support the assumption,

[10] In a few cases, subjects entered zero as minimum selling price for a $-bet or the highest outcome as minimum selling price for a probability bet. Note that the bounds in the restricted BDM-task did not prohibit these answers. The frequency of these observations was the same in both tasks. Since the prediction of our theory refers to the bound that is not equal to zero or the highest payoff, we did not include these answers in rows 5 and 6.

that lotteries are only compared to relevant amounts. In the restricted BDM-task, the set of relevant amounts is the interval from which prices can be chosen.

To test whether elicited prices differ significantly from the prediction of PI-theories, we compute a sample of prices that are predicted by PI-theories in the restricted BDM-task. To generate the sample of predicted prices, we take the prices that are elicited in the standard BDM-task and replace prices that lie outside the bounds with the respective bound and leave the other prices unchanged.

We first use for each lottery a one-sided two sample Kolmogorov-Smirnov test to check whether the sample of predicted prices comes from the same distribution as the elicited prices under the restricted BDM-task. The null-hypothesis is that both samples are drawn from the same distribution. For P-bets, the alternative hypothesis is that the prices under the restricted BDM-task come from a distribution that first order stochastically dominates the distribution from which the predicted prices are drawn. For $-bets, the alternative hypothesis is that the prices under the restricted BDM-task come from a distribution that is first order stochastically dominated by the distribution from which the predicted prices are drawn. Due to the small number of observations, we can reject the null-hypothesis only for Lottery 1 (with significance level $\alpha = 0.05$) and for lotteries 4, 10, and 11 (with significance level $\alpha = 0.1$).

The Kolmogorov-Smirnov test compares the distribution of elicited and predicted prices but makes no use of within subject comparisons. To use this information, we compare the elicited and predicted prices for each subject and each lottery. Predicted prices are generated by applying the prediction of PI-theories. It follows immediately that PI-theories predict that, for each subject and for each lottery, the probability that the predicted price is smaller than the elicited price in the restricted BDM-task is equal to the probability that the predicted price is larger than elicited price in the restricted BDM-task. Note that this prediction applies to P-bets as well as to $-bets.

If subjects determine the minimum selling price according to procedure P, then the probability that the elicited price in the restricted BDM-task is smaller or larger depends on the location of the interval in the restricted BDM-task. Recall that intervals are centered around the expected value and that the length of the interval was chosen such that either the lower bound is zero or the upper bound equal to he maximum outcome of the lottery. For P-bets, our theory predicts that the probability that the elicited price in the restricted BDM-task is smaller than the minimum selling price in the standard BDM-task is smaller than the probability that the elicited price in the restricted BDM-task is larger. For $-bets, our theory predicts that it is more likely that the elicited price in the restricted BDM-task is smaller.

We use a sign-test to analyze how predicted and elicited prices differ. For each subject, we compute the difference between the elicited price in the restricted BDM-task and predicted price according to PI theories. Since only the sign of the difference matters, we do not need to test each lottery separately but can combine the data from P-bets and from $-bets. The null-hypothesis is that the differences are drawn from a distribution with median zero. Table 2 summarizes the results. The second and third row of table 2 list the percentage of elicited prices under the restricted BDM-task that are higher and lower than the price that is predicted by PI-theories. For two $-bets (lotteries 2 and 13), all subjects stated prices in the interior of I_L - both in restricted and the standard BDM-task. These lotteries are excluded in the third row of table 2.

Table 2

	P-bets (Lottery 3,4,6,8,10,12,14)	$-bets (Lottery 1,2,5,7,11,13)	$-bets where bounds bind (Lottery 1,5,7,11)
% above predicted	44	31	22
% below predicted	19	44	53
p-value	0.000	0.049	0.000

The sign-test shows that we can reject the null-hypothesis that it is equally likely that the elicited price in the restricted BDM-task is smaller respectively larger than the minimum selling price in the standard BDM-task. The comparison between the elicited certainty equivalents in the standard and the restricted BDM-task support our assumption that subjects draw the starting values from the relevant interval.

4.2. The Price Adjustment Pattern

In this section we describe the results of an experiment that is designed to test the assumption that subjects use a stepwise procedure to determine certainty equivalents. If subjects use a stepwise procedure to determine the certainty equivalent of a lottery, then subjects do not immediately know their certainty equivalent. Instead, the best estimate of the certainty equivalent changes over time as the lottery is compared to different amounts until the comparison procedure is finished.

To test this part of our theory, we designed the radioline task where subjects

- have an incentive to state at every point in time their best estimate of the minimum selling price
- can adjust the price if they wish to do so.

Recall that in the radioline task the computer screen displays the lottery and a list of 30 prices spaced equally apart between zero and the highest outcome of the lottery. Subjects are asked to indicate their minimum selling price by using their mouse to click one of the prices on list. Since the question is terminated after some randomly determined time, subjects have an incentive to state a price soon. However, subjects can change the price as often as they want. The computer program records the order in which prices are clicked.

If subjects determine certainty equivalents according to our theory, then we expect to observe the following:

1) If the time restriction binds, we expect to observe frequent price adjustments where subjects adjust prices after they made additional comparisons between the lottery and some monetary amount for sure.

2) For a given subject and lottery, we expect that the sequence of prices is not random but instead follows a predictable pattern.

3) For P-bets, we expect that it is more likely that the first price is below the final price than vice versa. For $-bets, we expect that it is more likely that the first price is higher than the final price than vice versa.

We discuss the results separately for each prediction.

1) We observe a surprisingly large number of price adjustments. Out of 180 observations (30 subjects and 6 lotteries), there are only 34 cases in which subjects entered only a single price. Line 4 in Table 4.1 lists the time until the question was terminated and the average number of prices that are entered. The more time subjects had, the more frequently did they revise their price. A simple linear regression returns an intercept of 0.98 and a slope of 0.125, i.e., if subjects have an additional eight seconds, they revise their price once more. These results show that the determination of a certainty equivalent is a time consuming process.

2) The observation of frequent price adjustments alone does not prove that subjects indeed follow a stepwise procedure where new comparisons depend on the results of old comparisons. Instead, price changes could just be the consequence of subjects who draw a new preference relation and state a new price without there being any connection between the prices. If this explanation is true, the sequence of prices that a subject enters should be random. The computer program recorded the complete sequence of prices that a subject entered. Obviously, sequences are too short to apply standard tests for randomness. To analyze whether prices are random, we look at all subsequences of three consecutively entered prices where the first price is not the same as the third price. If prices are random, then we should expect that for half of

Lottery	16	17	18	19	20	21
Probability win	0.15	0.22	0.35	0.45	0.68	0.85
High outcome	150	120	90	96	60	84
Time limit in sec.	19	27	17	22	21	18
Average # of prices	3.3	4.5	3.67	3.8	3.23	3.83
# first price higher	16	16	12	12	7	6
# first price lower	6	9	4	12	15	16
p-value	0.053	0.23	0.077	1	0.134	0.053
# of observations	30	30	30	30	30	30

Figure 4.1:

these subsequences either the third price is lower than the first although the second prices is larger than the first or that the third price is higher than the first although the second prices is lower than the first. In contrast, our theory predicts that if the second price is higher than the first price, then the third price should not be lower than the first price (and similarly if the second price is lower than the first price). There are 161 subsequences of three prices where the first price is not the same as the third price. Of these 161 subsequences, there are only 14 subsequences (i.e., 8.7%) where either the third price is lower than the first although the second prices is larger than the first or that the third price is higher than the first although the second prices is lower than the first. These results clearly show that the sequences of prices that subjects enter is not random.

3) According to our theory, subjects start to determine the certainty equivalent by comparing the lottery with some amount between the lowest and the highest outcome of the lottery. Hence, the certainty equivalent that a subject states in an experiment depends on the starting value and on the preferences. The design of the radioline task provides strong incentives to enter a price quickly since subjects otherwise risk to earn zero if the question is terminated before they enter a price. This implies that some subjects enter a price before they finished the sequence of comparisons to determine the certainty equivalent. In the early stages of the

sequence of comparisons, the price that subjects enter is influenced more by the starting value and less by the preferences than the final price that subjects enter. If starting values are spread out between the lowest and highest outcome of the lottery, this implies that the first price is likely to be lower than the final price if the lottery has a large probability to win the highest outcome. Similar, the first price is likely to be higher than the final price if the lottery has a small probability to win the highest outcome.

Therefore, our theory predicts that the first price that is entered is more likely to be lower than the final price for P-bets and that the first price is more likely to be higher for \$-bets. Rows 6 and 7 of Table 4.1 list the number of observations where the first price that was entered is lower and higher than the final price, respectively. With the exception of lottery 3, we observe that it becomes more likely that the first price is smaller than the final price when the probability that the high payoff is realized increases. We use a sign-test to test the prediction that the first price that is entered is more likely to be lower than the final price for P-bets and that the first price is more likely to be higher for \$-bets. The null-hypothesis is that it is equally likely that the first price is higher or lower than the final price. Row 8 of Table 4.1 lists the p-values for a two-sided sign-test. As we would expect, the effect is larger and more significant for lotteries with the probabilities close to zero or one (lotteries 1 and 6).

5. Conclusion

The preference reversal phenomenon is the prime example for inconsistent behavior across choice and pricing problems. Since recent studies show that the preference reversal phenomenon prevails if the incentive scheme does not require that the independence and the reduction axiom hold, most researchers attribute the preference reversal phenomenon to preferences that are not context-free. We follow a different approach. We retain the fundamental assumption of standard economics that preferences are context-free.

We go back to the basic notion that a preference relation is a binary relation. Therefore, we consider individuals who can only make binary choices and model the determination of certainty equivalents in terms of binary choice. Our model generates a probability distribution over the certainty equivalents that subjects state. We analyze different experimental procedures that are used to elicit certainty equivalents and show that our model explains why the distributions of elicited certainty equivalents differ systematically. Additionally, our model provides a simple and intuitive explanation for the preference reversal phenomenon. Although we assume that preferences are context-free, our model explains why subjects tend to overprice $-bets and to underprice P-bets which in turn generates the preference reversal phenomenon. To test our theory, we develop a new experimental design. This design allows us to observe how subjects adjust certainty equivalents over time. The pattern of the adjustment process confirms the predictions of our model.

This paper is the first part of a larger research agenda. The fundamental concept behind our approach is to model behavior in complex decision problems in terms of binary choice. This concept can be fruitfully applied to other questions. For example, experiments have shown that subjects frequently violate the choice axiom. Our concept to model complex decision problems in terms of binary choice can explain these violations of the choice axiom (such as an asymmetric dominance effect, similarity effect, etc.).

6. Literature

Ballinger, P. and N. Wilcox (1997) "Decisions, error and heterogeneity" Economic Journal 107, 1090-1105

Becker, G.M., DeGroot, M.H. & Marschak, J. (1964). Measuring utility by a single-response sequential method. Behavioral Science, 9, 226-232

Bostic, R., D. Luce and J. Herrnstein (1990) "The effect on the preference-reversal phenomenon of using choice indifferences" Journal of Economic Behavior and Organization 13, 193-212

Buschena, D. and Zilberman, D. (2000) "Generalized expected utility, heteroscedastic error, and path dependence in risky choice" Journal of Risk and Uncertainty 20, 67-88

Busemeyer, J. and J. Johnson (2005) A dynamic, stochastic, computational model of preference reversal phenomena, Psychological Review, 2005 Oct ;112 (4):841-61

Butler, David and Loomes, Graham: Imprecision as an Account of the Preference Reversal Phenomenon, forthcoming American Economic Review, 2007

Cohen, Michèle, Jean-Yves Jaffray, & Tanios Said (1985), "Individual Behavior under Risk and under Uncertainty: An Experimental Study," Theory and Decision 18, 203–228

Cubitt, R., Munro, A. and C. Starmer, Testing Explanations of Preference Reversal, Economic Journal 114, 2004, p. 709-726

Hershey, John C. & Paul J.H. Schoemaker (1980), "Prospect Theory's Reflection Hypothesis: A Critical Examination," Organizational Behavior and Human Performance 25, 395–418

Hey, J. (2001) "Does repetition improve consistency?" Experimental Economics 4, 5-54

Hey, J.D. and C. Orme (1994) Investigating generalisations of expected utility theory using experimental data, Econometrica 62, 1291-1326

Holt C. (1986) "Preference Reversals and the Independence Axiom", American Economic Review, Vol. 76 (3), p.508-15.

Karni E. and Z. Safra (1987) "Preference Reversals and the Observability of Preferences by Experimental Methods", Econometrica, Vol. 51, p.675-85

Loomes, G. (2005) "Modelling the stochastic component of behaviour in experiments: some issues for the interpretation of data" Experimental Economics 8, p. 301-323

Loomes, G. and R. Sugden, A Rationale for Preference Reversal, American Economic Review 73, 1983, p. 428-432

Machina, M. (1985) "Stochastic Choice Functions Generated from Deterministic Preferences over Lotteries", Economic Journal 95, 575-594

Schkade , D. and E. Johnson (1989), Cognitive Processes in Preference Reversals, Organizational Behavior and Human Decision Processes 44,, p. 203-231

Segal, U., Does the Preference Reversal Phenomenon necessarily contradict the Independence Axiom?, American Economic Review 78, 1988, p. 233-236

Seidl, C., Preference Reversal, Journal of Economic Surveys, 16(5), 2002

Tversky, A. and Kahneman, D. (1992) "Advances in prospect theory: Cumulative representation of uncertainty" Journal of Risk and Uncertainty 5, 297-323

Tversky, A., Slovic, P. and Kahneman, D. (1990) "The Causes of Preference Reversal" American Economic Review 80, 204-217

Tversky, A. and Thaler, R.H., Preference Reversals, Journal of Economic Perspectives 4, 1990, p. 201-211

7. Appendix

Proof of Theorem 1

We prove Theorem 1 for the case where $EV_L < \frac{\underline{x}_L + \overline{x}_L}{2}$. The proof for $EV_L > \frac{\underline{x}_L + \overline{x}_L}{2}$ is similar. Let $\Pr\{CE_L < EV_L\}_s$ be the probability that the stated certainty equivalent is smaller than EV_L if s is the first amount to which the lottery is compared. We consider first $s \in [\underline{x}_L, 2EV_L - \underline{x}_L]$ and then $s \in (2EV_L - \underline{x}_L, \overline{x}_L]$.

Let I be an indicator function such that $I(a)$ takes value one if a is a true statement and zero otherwise. Consider $\underline{s} \in [\underline{x}_L, EV_L)$. Define $\overline{s}$ with $\overline{s}(\underline{s}) = 2EV_L - \underline{s}$. Since preference relations satisfy Assumption 4 (Convexity) we have $\mu(N(\min\{\underline{s} + \triangle, \overline{x}_L\} \succsim L)) = \mu(N(\underline{s} + \triangle \succsim L))$ and $\mu(N(\max\{\overline{s} - \triangle, \underline{x}\} \succsim L)) = \mu(N(\overline{s} - \triangle \succsim L))$, and similar for indifference and strict preference. To simplify notation, we drop the min/max operator whenever possible.

$\Pr\{CE_L < EV_L\}_{\underline{s}} = \mu(N(\underline{s} \succsim L))+$

$\mu(N(L \succ \underline{s})) \cdot [\mu(N(\underline{s} + \triangle \succ L)) \cdot I(\min\{\underline{s} + \frac{\Delta}{2}, \frac{\underline{s}+\overline{x}_L}{2}\} < EV_L)+$

$\mu(N(\underline{s} + \triangle \sim L)) \cdot I(\underline{s} + \Delta < EV_L)+$

$\mu(N(L \succ \underline{s} + \triangle)) \cdot [\mu(N(\underline{s} + 2\triangle \succ L)) \cdot I(\min\{\underline{s} + \frac{3}{2}\Delta, \frac{\underline{s}+\Delta+\overline{x}_L}{2}\} < EV_L)+$

$\mu(N(\underline{s} + 2\triangle \sim L)) \cdot I(\underline{s} + 2\Delta < EV_L) + [...]]]$

And similar for $\overline{s}(\underline{s})$:

$\Pr\{CE_L > EV_L\}_{\overline{s}(\underline{s})} = \mu(N(L \succsim \overline{s}(\underline{s})))+$

$\mu(N(\overline{s}(\underline{s}) \succ L)) \cdot [\mu(N(L \succ \overline{s}(\underline{s}) - \triangle)) \cdot I(\max\{\overline{s}(\underline{s}) - \frac{\Delta}{2}, \frac{\overline{s}(\underline{s})+\underline{x}_L}{2}\} > EV_L)+$

$\mu(N(L \sim \overline{s}(\underline{s}) - \triangle)) \cdot I(\overline{s}(\underline{s}) - \triangle > EV_L)+$

$\mu(N(\overline{s}(\underline{s}) - \triangle \succ L)) \cdot [\mu(N(L \succ \overline{s}(\underline{s}) - 2\triangle)) \cdot I(\max\{\overline{s}(\underline{s}) - \frac{3}{2}\Delta, \frac{\overline{s}(\underline{s})-\Delta+\underline{x}_L}{2}\} > EV_L)+$

$\mu(N(L \sim \overline{s}(\underline{s}) - 2\triangle)) \cdot I(\overline{s}(\underline{s}) - 2\triangle > EV_L) + [...]]]$

To see that $\Pr\{CE_L < EV_L\}_{\underline{s}} = \Pr\{CE_L > EV_L\}_{\overline{s}(\underline{s})}$ note that Assumption 5 implies that $\mu(N(\underline{s} \succsim L)) = \mu(N(L \succsim \overline{s}(\underline{s})))$. Hence $\mu(N(L \succ \underline{s})) = \mu(N(\overline{s}(\underline{s}) \succ L))$. The same argument implies that $\mu(N(\underline{s} + \triangle \succ L)) = \mu(N(L \succ \overline{s}(\underline{s}) - \triangle))$, and so on.

Since $\underline{s} \in [\underline{x}_L, EV_L)$, we have $\overline{s}(\underline{s}) \leq 2EV_L - \underline{x}_L$. Recall that by assumption $EV_L < \frac{\underline{x}_L+\overline{x}_L}{2}$. Hence $\frac{\underline{s}+\overline{x}_L}{2} > EV_L$ and $\frac{\overline{s}(\underline{s})+\underline{x}_L}{2} \leq EV_L$. Note that $\underline{s} + \frac{\Delta}{2} < EV_L$ if and only if $\overline{s}(\underline{s}) - \frac{\Delta}{2} > EV_L$. Hence $I(\min\{\underline{s} + \frac{\Delta}{2}, \frac{\underline{s}+\overline{x}_L}{2}\} < EV_L) = 1$ if and only if $I(\max\{\overline{s}(\underline{s}) - \frac{\Delta}{2}, \frac{\overline{s}(\underline{s})+\underline{x}_L}{2}\} > EV_L) = 1$. A similar argument shows that $I(\underline{s} + \Delta < EV_L) = 1$ if and only if $I(\overline{s}(\underline{s}) - \triangle > EV_L) = 1$, and so on.

Hence for $\underline{s} \in [\underline{x}_L, EV_L)$, we have $\Pr\{CE_L < EV_L\}_{\underline{s}} = \Pr\{CE_L > EV_L\}_{\overline{s}(\underline{s})}$. Using the same logic, one can show that $\Pr\{CE_L > EV_L\}_{\underline{s}} = \Pr\{CE_L < EV_L\}_{\overline{s}(\underline{s})}$ for $\underline{s} \in [\underline{x}_L, EV_L)$. It is straightforward to show that $\Pr\{CE_L > EV_L\}_{EV} = \Pr\{CE_L < EV_L\}_{EV}$. For the first comparison, the individual draws some amount at random from $[\underline{x}_L, \overline{x}_L]$. Recall that CE_L is the average of the last two amounts to which the lottery has been compared (except if the individual was indifferent between lottery and amount). Hence $\Delta > 0$ implies that $\Pr\{CE_L =$

$EV_L\} = 0$. Hence to prove Theorem 1 it is sufficient to show that $\Pr\{CE_L > EV_L\}_s > 0.5$ for all $s \in (2EV_L - \underline{x}_L, \overline{x}_L]$.

Now consider $\widetilde{s} \in (2EV_L - \underline{x}_L, 2EV_L - \underline{x}_L + \Delta]$. Let n be the unique integer such that $\underline{x}_L + (n + 0.5)\Delta \leq EV_L$ and $\underline{x}_L + (n + 1.5)\Delta > EV_L$. Hence $\widetilde{s} - (n + 0.5)\Delta > EV_L$.

Hence

$$\Pr\{CE > EV_L\}_{\widetilde{s}} \geq 1 - \mu(N(L \prec \widetilde{s})) + \mu(N(L \prec \widetilde{s})) \cdot [1 - \mu(N(L \prec \widetilde{s} - \Delta)) + \mu(N(L \prec \widetilde{s} - \Delta)) \cdot [..... \cdot [1 - \mu(N(L \precsim \widetilde{s} - (n + 1)\Delta))]...]] \quad \text{(A1)}$$

Note that the last term on the RHS of exp.(A1) requires only that $\widetilde{s} - (n + 1)\Delta$ is weakly preferred to L. The reason is that $\widetilde{s} - (n + 1)\Delta$ might be smaller than EV_L for some $\widetilde{s}$.

Rewriting (A1) gives

$$\Pr\{CE > EV_L\}_{\widetilde{s}} \geq 1 - \prod_{j=1}^{n} \mu(N(L \prec \widetilde{s} - j\Delta)) \cdot \mu(N(L \precsim \widetilde{s} - (n + 1)\Delta))$$

Transitivity implies that

$$\begin{aligned} \mu(N(L &\prec \widetilde{s} - (j + 1)\Delta)) \leq \mu(N(L \prec 2EV_L - \underline{x}_L - j\Delta)) \\ \mu(N(L &\precsim \widetilde{s} - (j + 1)\Delta)) \leq \mu(N(L \precsim 2EV_L - \underline{x}_L - j\Delta)) \end{aligned}$$

Since $\mu(N(L \precsim z)) \geq \mu(N(L \prec z))$ for all $z \in X$ we have

$$1 - \prod_{j=1}^{n} \mu(N(L \prec \widetilde{s} - j\Delta)) \cdot \mu(N(L \precsim \widetilde{s} - (n + 1)\Delta)) \geq 1 - \prod_{j=1}^{n} \mu(N(L \precsim 2EV_L - \underline{x}_L - j\Delta))$$

Hence $\Pr\{CE > EV_L\}_{\widetilde{s}} > 0.5$ for $\widetilde{s} \in (2EV_L - \underline{x}_L, 2EV_L - \underline{x}_L + \Delta]$.

From Convexity, Transitivity and Assumption 5 follows that $\mu(N(L \prec \widetilde{s} - j\Delta)) = 1$ for all $\widetilde{s} \in (2EV_L - \underline{x}_L + j\Delta, 2EV_L - \underline{x}_L + (j+1)\Delta]$. Hence $\Pr\{CE > EV_L\}_{\widetilde{s}} = \Pr\{CE > EV_L\}_{\widetilde{s}-j\Delta}$ for all $\widetilde{s} \in (2EV_L - \underline{x}_L + j\Delta, 2EV_L - \underline{x}_L + (j+1)\Delta]$. Hence $\Pr\{CE > EV_L\}_{\widetilde{s}} > 0.5$ for all $\widetilde{s} \geq 2EV_L - \underline{x}_L$ and, therefore, $\Pr\{CE > EV_L\} > 0.5$ if $EV_L < \frac{\underline{x}_L + \overline{x}_L}{2}$.

Proof Lemma 2 (second price auction-ascending or descending)

Note that the stated certainty equivalent is equal to the price where the preferred alternative switches. Let n be the unique integer such that $\underline{x}_L + (n+0.5)\Delta \leq EV_L$ and $\underline{x}_L + (n+1.5)\Delta > EV_L$. If an ascending second-price auction starts at $\underline{x}_L$, the probability that the stated certainty equivalent is less than EV_L is:

$$\Pr\{CE < EV_L\} \geq 1 - \mu(N(L \prec \underline{x}_L + \Delta)) + \mu(N(L \prec \underline{x}_L + \Delta)) \cdot [1 - \mu(N(L \prec \underline{x}_L + 2\Delta)) + \mu(N(L \prec \underline{x}_L + 2\Delta)) \cdot [..... \cdot [1 - \mu(N(L \prec \underline{x}_L + n\Delta]...]]$$

Collecting terms gives $\Pr\{CE < EV_L\} \geq 1 - \prod_{j=1}^{n} \mu(N(L \prec \underline{x}_L + j\Delta))$. The result follows immediately from Assumption 6.

If a descending second-price auction starts at $\overline{x}_L$, the probability that the stated certainty equivalent is larger than EV_L is:

$$\Pr\{CE > EV_L\} \geq 1 - \mu(N(L \prec \overline{x}_L - \Delta)) + \mu(N(L \prec \overline{x}_L - \Delta)) \cdot [1 - \mu(N(L \prec \overline{x}_L - 2\Delta)) + \mu(N(L \prec \overline{x}_L - 2\Delta)) \cdot [..... \cdot [1 - \mu(N(L \prec \overline{x}_L - n\Delta]...]]$$

Collecting terms gives $\Pr\{CE > EV_L\} \geq 1 - \prod_{j=1}^{n} \mu(N(L \prec \overline{x}_L - j\Delta))$. The result follows immediately from Assumption 6.

Proof Lemma 3 (compare to list of amounts)

This is essentially the same situation as in Theorem 1 except that the amount for the first comparison is drawn from grid instead of an interval. Consider the case where $EV_L < \frac{\underline{x}_L + \overline{x}_L}{2}$. Define $\overline{s}$ with $\overline{s}(\underline{s}) = 2EV_L - \underline{s}$. Since by assumption EV_L lies on the grid, we have $\Pr\{CE_L < EV_L\}_{\underline{s}} = \Pr\{CE_L > EV_L\}_{\overline{s}(\underline{s})}$ for all $\underline{s} \in [\underline{x}_L, EV_L]$. The additional assumption that $\mu(N(L \sim EV_L)) = 0$ ensures that $\Pr\{CE_L = EV_L\}_s = 0$. The same argument as in the proof of Theorem 1 shows that $\Pr\{CE_L > EV_L\}_s > 0.5$ for all $s \in (2EV_L - \underline{x}_L, \overline{x}_L]$.