

TI 2013-065/VII
Tinbergen Institute Discussion Paper



Using Preferred Outcome Distributions to estimate Value and Probability Weighting Functions in Decisions under Risk

Bas Donkers¹

Carlos J.S. Lourenco²

Benedict G.C. Dellaert¹

Daniel G. Goldstein³

¹ *Erasmus School of Economics, Erasmus University Rotterdam;*

² *RSM, Erasmus University Rotterdam;*

³ *Microsoft Research.*

Tinbergen Institute is the graduate school and research institute in economics of Erasmus University Rotterdam, the University of Amsterdam and VU University Amsterdam.

More TI discussion papers can be downloaded at <http://www.tinbergen.nl>

Tinbergen Institute has two locations:

Tinbergen Institute Amsterdam
Gustav Mahlerplein 117
1082 MS Amsterdam
The Netherlands
Tel.: +31(0)20 525 1600

Tinbergen Institute Rotterdam
Burg. Oudlaan 50
3062 PA Rotterdam
The Netherlands
Tel.: +31(0)10 408 8900
Fax: +31(0)10 408 9031

Duisenberg school of finance is a collaboration of the Dutch financial sector and universities, with the ambition to support innovative research and offer top quality academic education in core areas of finance.

DSF research papers can be downloaded at: <http://www.dsf.nl/>

Duisenberg school of finance
Gustav Mahlerplein 117
1082 MS Amsterdam
The Netherlands
Tel.: +31(0)20 525 8579

Using Preferred Outcome Distributions to Estimate Value and Probability Weighting Functions in Decisions under Risk

Bas Donkers¹, Carlos J.S. Lourenço², Benedict G.C. Dellaert¹, Daniel G. Goldstein³

Abstract

In this paper we propose the use of *preferred outcome distributions* as a new method to elicit individuals' value and probability weighting functions in decisions under risk. Extant approaches for the elicitation of these two key ingredients of individuals' risk attitude typically rely on a long, chained sequence of lottery choices. In contrast, preferred outcome distributions can be elicited through an intuitive graphical interface, and, as we show, the information contained in two preferred outcome distributions is sufficient to identify non-parametrically both the value function and the probability weighting function in rank-dependent utility models. To illustrate our method and its advantages, we run an incentive-compatible lab study in which participants use a simple graphical interface – the Distribution Builder (Goldstein et al. 2008) – to construct their preferred outcome distributions, subject to a budget constraint. Results show that estimates of the value function are in line with previous research but that probability weighting biases are diminished, thus favoring our proposed approach based on preferred outcome distributions.

¹Erasmus School of Economics, Department of Business Economics, Marketing Section PO Box 1738, 3000 DR Rotterdam, The Netherlands, and Tinbergen Institute, ERIM and Netspar. ²Rotterdam School of Management, Marketing Management department, PO Box 1738, 3000 DR Rotterdam, The Netherlands, and ERIM and Netspar. ³Microsoft Research, New York City and London Business School. Corresponding author: Bas Donkers, email: donkers@ese.eur.nl, Tel.: +31 10 408 2411, Fax: +31 10 408 9169.

1. Introduction

In this paper, we introduce preferred outcome distributions as a novel risk preference elicitation method that ensures non-parametric identification of the value function (VF) and the probability weighting function (PWF) in rank dependent utility models. In a context of decision making under risk, we define preferred outcome distributions as individuals' preferred combinations of probabilities and outcomes – a distribution – subject to a budget constraint. As such, preferred outcome distributions are rich sources of information on individuals' risk preferences, because they instantaneously reveal the trade-offs that individuals make across the range of all possible outcomes and the full range of probabilities. Individuals can construct their preferred outcome distributions using an interactive graphical interface such as the Distribution Builder (Sharpe et al. 2001, Goldstein et al. 2008), and, as we show, measurement of individuals' risk preferences with our method requires only two of such independent (i.e., non-chained) preferred outcome distributions.

We build upon and extend the work by Goldstein et al. (2008), who used a single preferred outcome distribution to elicit risk preferences under the assumption of expected utility, i.e. without allowing decision weights to deviate from objective probabilities. As we explain, a single preferred outcome distribution does not contain enough information to disentangle the curvature of the value function from the curvature of the probability weighting function, but by combining the information of (at least) two preferred outcome distributions that differ in the underlying risk-return trade-offs, we can identify non-parametrically both the value function and the probability weighting function.

Our proposed approach also alleviates some of the weaknesses that traditional, lottery-based preference elicitation approaches suffer from. First, standard gamble tasks (chained or

otherwise) may suffer from biases resulting from a focus on either probabilities or outcomes (Hershey and Schoemaker 1985).¹ In our approach, individuals construct a preferred outcome distribution and distribute probability mass (i.e., risk) across the total range of outcomes, which balances attention over both probabilities and outcomes and is likely to improve the validity of estimated risk preferences (McCord and Neufville 1986). Furthermore, the interactive, ‘natural frequency’ representations of risk that we use have been shown to improve individuals’ understanding of risk preference elicitation tasks compared to the probability descriptions and graphs that are typically used in gamble tasks (see e.g. Ancker, Weber, and Kukafka 2011, Gigerenzer and Galesic 2012, Ahmed et al. 2012, Kaufmann, Weber and Haisley 2012).

Second, elicitation methods that rely on a chained sequence of gambles to ensure non-parametric identification of individuals’ value and probability weighting functions may suffer from error propagation because answers to initial questions carry over to the construction of subsequent lotteries (e.g., Wakker and Deneffe 1996, Abdellaoui 2000, Bleichrodt and Pinto 2000, Van de Kuilen and Wakker 2011). The use of preferred outcome distributions overcomes this problem since the method does not rely on dependent, sequential decision tasks. For this same reason, preferred outcome distributions do not create an incentive to overstate indifference payoffs between risky prospects in earlier response rounds that may raise the individual’s expected payoffs in later rounds (as identified by Harrison and Rutstrom 2008 in elicitation tasks with multiple rounds). As already mentioned, a preferred outcome distribution is constructed in one single task, over the full range of possible outcomes and the entire probability range.

Finally, construction of preferred outcome distributions stimulates the use of intuitive, interactive graphical interfaces in the elicitation of risk preferences. Tools such as the

¹ For instance, Holt and Laury’s (2002) classic experiment based on a menu of lotteries with fixed amounts and changing levels of risk, (over)emphasizes probabilities.

Distribution Builder easily and intuitively allow individuals to experience probabilities and associated outcomes (Goldstein et al. 2008), which is well in line with recent work in decision making under risk (e.g., Abdellaoui, L'Haridon and Paraschiv 2011 and Choi et al. 2007, Barron and Erev 2003, Hertwig, Barron, Weber, and Erev 2004, Camilleri and Newell 2011). Furthermore, the task of graphically constructing a preferred outcome distribution is much less repetitive in nature than a chained lotteries approach; studies estimating individuals' risk preferences using chained lotteries often requiring participants to make more than one hundred choices. When using a preferred outcome distributions approach, elicitation of two distributions is enough to identify individual-level estimates of the value and probability weighting function. These features of the preference elicitation task will be attractive to researchers and financial services firms interested in eliciting customers' risk preferences in an easy manner while accounting still for biases in probability weights.

The remaining of the manuscript is organized as follows. In the next section, we offer an intuitive overview of the basic structure of the Distribution Builder and of its main components. We explain risk-return tradeoffs and show how they are presented within the Distribution Builder. In the second and third sections we formally describe the decision problem underlying a typical Distribution Builder task and how this task can be used to elicit the value function and the probability weighting function of a model of rank-dependent utility (e.g. Tversky and Kahneman's (1992) Prospect Theory). Furthermore, we discuss the requirements for identification of these functions, both parametrically and non-parametrically. In the fourth section, we describe the design of an incentive-compatible lab study in which participants use the Distribution Builder to construct their preferred outcome distributions subject to a budget constraint. We present the results regarding the value function and the probability weighting

function, both at the individual-level and averaged across all subjects. The estimates of the value function are in line with previous research but the probability weighting biases are diminished, thus favoring our preferred outcome distributions based approach. We conclude the paper with a discussion of our results.

1. Basic Structure of the Distribution Builder

The Distribution Builder (Goldstein et al. 2008) draws on previous research on risk representations demonstrating that individuals are best capable of understanding probabilities when given a graphical presentation of the frequencies of occurrence of a risky event (Fagerlin, Zikmund-Fisher and Ubel, 2011). The main contribution of the Distribution Builder is the implementation of these ideas within a sound behavioral financial economics framework (e.g. modern portfolio theory and rank-dependent utility theory) and in a user-friendly online environment, thus making it attractive for widespread adoption by scholars and practitioners.

The Distribution Builder combines several decision variables interactively in a single tool, providing consumers with a simple and intuitive device to express their preferences over a large range of values of the decision variables used. In an investment context, consumers using the Distribution Builder are typically asked to determine their most preferred probability distribution for obtaining future revenues, given a budget constraint on today's investments. In other words, with the Distribution Builder, consumers make a joint decision on risk and returns of financial outcomes, subject to a budget constraint. In Figure 1, we present the initial layout of the Distribution Builder as seen by users on a computer screen, and we highlight and explain in detail six important features of the measurement tool that enable the reader to better follow our

examples and later sections. Figure 2 presents illustrative examples of the use of the Distribution Builder and of its main features.

At its most basic level, the Distribution Builder presents consumers with an interactive graph of 100 markers representing investment outcomes that occur with equal probability (see ‘Features 1 and 2’ in Figure 1), similar to tickets that are drawn in a lottery. In addition, the Distribution Builder graphically represents the fact that (1) not all investment outcomes have equal value, (2) an individual’s investments have to be made from a limited financial budget, (3) higher investment outcomes are more costly, and (4) by taking more risk a higher expected return can be obtained (this fourth aspect is implemented using unequal Arrow-Debreu state prices for each of the markers (Arrow 1964, Debreu 1959); more details are given below).




Figure 1.

The Distribution Builder: initial interface layout

Feature 1. One hundred 1-percent probability markers situated in the leftmost column at the start. **Feature 2.** Outcome levels and labels corresponding to just as many discrete outcome columns, to which any number of probability markers (individually or stacked) can be dragged-and-dropped with a computer mouse. In the present setting, there are 21 outcome levels, from zero to 2,000 euros, in increasing order from left to right and in equal steps of 100 euros. **Feature 3.** Interactive investment budget meter, set to zero value (0%) at the start. **Feature 4.** Interactive investment budget meter bar. This bar is colored gray at the start, corresponding to no investment budget spent; all markers are then situated in their starting position in the leftmost column. The bar changes color depending on the percentage of investment budget spent, which in turn depends on the outcome distribution being constructed by moving the markers to the outcome levels. The differently colored bar moves to the right (left) indicating an increase (decrease) in investment budget spent. When the investment budget is underused (overused), the budget meter bar will be blue (red); see Figure 2 (panels B and C). When between 99.5% and 100.5% of the investment budget is spent, the budget meter bar will be green (see Figure 2, panel D). The “Done” button can then be clicked to submit the outcome distribution constructed (see next feature). **Feature 5.** Users must click the rectangle-shaped “Done” button upon completion of (and when satisfied with) their most preferred outcome distribution. If the button is clicked when the investment budget is over- or underused, an automatic message will be displayed on screen, informing users to properly use the available budget. **Feature 6.** Period counter to identify the period or round that a user is in. Users can use the Distribution Builder from 1 to the maximum number of rounds. **Feature 7.** Circle-shaped question mark help button. When clicked, an automatic message will be displayed explaining what the user must do (e.g. move the markers and respect the budget meter).




Figure 2.

Examples of constructed outcome distributions in the Distribution Builder

Panel A. Initial Distribution Builder layout as seen by users on a computer screen (same as Figure 1). From here, users can drag-and-drop the green probability markers with a computer mouse into any possible outcome distribution. Panels B to D are examples of such outcome distributions. **Panel B.** An outcome distribution using less than 99.5% of the investment budget (hence the budget meter is blue). **Panel C.** An outcome distribution using more than 100.5% of the investment budget (hence the budget meter is red). **Panel D.** An outcome distribution using between 99.5% and 100.5% of the investment budget (hence the budget meter is green).

When using the Distribution Builder, these four aspects are experienced through two key interface features. First, potential outcomes are ordered from low to high in the graphical

interface. Second, consumers are presented with an interactive budget ‘meter’ that shows how much of their investment budget is spent depending on the potential outcome levels where they place the markers (see ‘Features 3 and 4’ in Figure 1). Thus, consumers cannot simply place all investment markers on the investment outcome that they prefer the most (i.e. the highest return). Instead, consumers are required to make trade-offs between a fairly safe but average payoff and the chance of a higher payoff in combination with the chance of a lower payoff (see Figure 3; Panels B to D). This latter downside risk cannot be avoided – the chance of winning a higher payoff is more expensive – so that expenditures for other markers will need to go down.

The Distribution Builder interface further assists individuals by automatically assigning the underlying stochastic or probabilistic ‘states of the world’ (e.g. degrees of economic recession or returns on the asset markets) and their corresponding state prices to the different outcomes in an optimal manner. Details will be provided in the next section, where we introduce the user’s decision problem more formally, but the essence is that the optimal assignment minimizes the budget required to achieve a specific distribution. An important consequence of this automatic optimal assignment is that to the user the presence of different Arrow-Debreu state prices is only noticeable through the implied trade-off between risk and return. The implementation of these trade-offs based on sound financial economics principles in an interactive graphical presentation of the frequencies of occurrence of risky outcomes is illustrated in the Distribution Builder examples in Figure 3.

A

B




Figure 3.

Example of the risk-return tradeoff and underlying state prices in the Distribution Builder

Panel A. An outcome distribution using between 99.5% and 100.5% of the investment budget (hence the green budget meter), where all probability markers are placed on two outcomes only (zero and 2,000 euros). With an investment budget of 1,000 euros, such preferred outcome distribution could indicate a *risk-seeking* individual.

Panel B. An outcome distribution using between 99.5% and 100.5% of the investment budget (hence the green budget meter), where all probability markers are placed on two outcomes only (1,000 and 1,100 euros). With an investment budget of 1,000 euros, such preferred outcome distribution could indicate a *risk-averse* individual.

Panels C and D. Starting from the outcome distribution of Panel B, moving three markers to the right to have a 3% chance of a return of 1,200 euros would mean spending more than the available investment budget, which is not feasible (in Panel C the budget meter is red). To stay on budget, one must trade the chance of a 1,200 euro return for the chance of a lower return. Moving one marker from the 1,000 euros pile to the left, to a 1% chance of a 900 euro return, would be a possibility (in Panel D the budget meter is green).

2. Eliciting preferred outcome distributions

In this section we discuss the use of the DB for the elicitation of preferred outcome distributions.

We define a preferred outcome distribution as an individual's preferred combination of probabilities and outcomes in a decision under risk, subject to a budget constraint. When Goldstein et al. (2008) introduced the DB, they assumed an expected utility model in the analysis of preferred outcome distributions. As a consequence, the decision weight assigned to an outcome was assumed to be equal to the objective probability associated with that outcome.

In this paper, we study the decision problem underlying preferred outcome distributions for the case of rank dependent preferences. Let $v(\cdot)$ denote the value function and $\pi(\cdot)$ denote the transformation that is applied to the cumulative probability distribution function (see Tversky and Kahneman 1992 and Quiggin 1981). Moreover, let $w=1, \dots, W$ denote the possible states of the world such that they are ordered in terms of their Arrow-Debreu state prices. In particular, s_w denotes the Arrow-Debreu state price of state w , which is the investment required to guarantee a payoff of one in state w . We assume states are ordered such that $s_1 > s_2 > \dots > s_W$. Finally, p_w denotes the probability that state of the world w realizes and o_w the desired outcome level for that state of the world. In order to construct a preferred outcome distribution, one has to maximize the subjective expected utility subject to the budget constraint. With B denoting the available budget, a formal representation of this decision problem is given by (we omit individual specific subscripts for clarity of exposition):

$$(1) \quad \max_{o_w, w=1, \dots, W} \sum_{w=1}^W \pi_w v(o_w), \quad s.t. \quad \sum_{w=1}^W o_w s_w \leq B,$$

with

$$(2) \quad \pi_w = \pi\left(\sum_{j=1}^w p_j\right) - \pi\left(\sum_{j=1}^{w-1} p_j\right).$$

Here π_w is the decision weight assigned to state w .² For Equation (2) to represent rank dependent preferences, the markers should be placed in order of increasing outcome levels as well. Therefore we first show that every optimal allocation of outcomes to states satisfies this property.

We assume, as it is common, that the value function does not depend on the state of the world w directly, but only on the outcome obtained in that state of the world, o_w . This assumption and the assumption of a non-decreasing value function – a higher utility can only be achieved when a larger budget is available – make the optimal assignment of a set of desired outcome levels to states of the world straightforward, and greatly simplify the solution of the complex decision problem outlined above. To see this, note that the cost of generating a specific outcome distribution is minimal if the states with the highest state prices (i.e., for states of the world in which it is very costly to guarantee a certain outcome level) are assigned to the lowest outcomes (i.e., investment returns). As a consequence, a utility maximizing outcome distribution assigns the states with high state prices to the lowest outcome levels. Otherwise, the outcome distribution could be improved upon by reassigning (higher) outcome levels to (lower priced) states of the world, which would require a smaller budget. Under utility maximization, the order (in terms of decreasing size) of the state prices is therefore perfectly aligned with the order in the outcome levels that the markers are assigned to. Hence, every solution to the problem in Equations (1) and (2) automatically satisfies the required link between the ordered state prices and the ordered outcomes for Equation (2) to represent rank dependent preferences.

This link between the rank of the outcome of a marker in a preferred outcome distribution and the corresponding rank of its state price is implemented in the DB graphical interface to

² If $\pi_w = p_w$, Equations (1) and (2) formalize the decision problem under expected utility as in Goldstein et al. (2008).

simplify the decision problem. Instead of a marker representing a particular state of the world, the tool automatically identifies the state of the world that a marker corresponds to *ex post*, based on its location in the desired distribution. These locations are updated every time a user moves some markers. Thus, without burdening the user, the tool itself ensures an optimal assignment of outcomes to states and corresponding state prices to minimize the budget needed to obtain a given outcome distribution.

3. Preference elicitation using preferred outcome distributions

3.1. Parametric identification

The structure of the decision problem that underlies preferred outcome distributions facilitates the analysis of underlying risk preferences by use of the first order conditions corresponding to the decision problem in Equations (1) and (2). With γ denoting the Lagrange multiplier of the constrained optimization problem,³ the first order conditions with respect to the decision variables o_w are:

$$(3) \quad \pi_w v'(o_w) - \gamma \pi_w = 0, \text{ for } w=1, \dots, W.$$

Therefore, one preferred outcome distribution provides W observations on the value function and the probability weighting function, allowing certain parametric versions of $\pi(\cdot)$ and $v(\cdot)$ to be estimated at the individual-level, as long as the markers (i.e., desired payoffs) are sufficiently spread over the range of outcome levels.⁴

³ Xia and Zhou (2012), indicate that, under mild conditions on the primitives of the economy, equilibrium state prices satisfy the conditions required for this approach to be valid. Given that the state prices we use are equilibrium prices, that is, they are calibrated to actual market data, we assume the necessary conditions, as listed in remark 3.6 of Xia and Zhou (2012), are satisfied in our setting.

⁴ This might not always be the case. For example, Goldstein et al. (2008) exclude 50 out of 620 observations, as they have all markers at the reference level. A Bayesian estimation method that aggregates parameter information across individuals could solve this identification problem.

3.2. Non-parametric identification using variation in state prices

While Equation (3) might provide enough information to enable parametric identification, it is desirable to make no prior assumptions regarding the functional forms of the functions of interest (Abdellaoui 2000, Bleichrodt and Pinto 2000) and identify $\pi(\cdot)$ and $v(\cdot)$ non-parametrically.

Given that there are W decision weights π_w , a greater number than W observations are needed for estimation. To achieve this, we propose to elicit preferred outcome distributions in multiple scenarios that vary in state prices s_w (and hence in the distribution of asset returns) and then combine the information contained in the resulting preferred outcome distributions for estimation.

In this paper we will elicit preferred outcome distributions in two scenarios, one with high average returns and one with low average returns, denoted by H and L , respectively. The corresponding first order conditions are represented by:

$$(5) \quad \pi_w v'(o_w^c) - \gamma^c s_w^c = 0 \text{ for } c=H,L \text{ and } w=1,\dots,W.$$

Combining the first order conditions for the w^{th} marker in the two scenarios provides the following equation that can be used straightforwardly to estimate the *value function*:

$$(6) \quad \frac{v'(o_w^H)}{v'(o_w^L)} = \frac{\gamma^H s_w^H}{\gamma^L s_w^L}.$$

In Equation (6) the probability weights have cancelled out, facilitating estimation of the value function. Further taking logarithms, we obtain

$$(7) \quad \ln(s_w^H) - \ln(s_w^L) = \ln(v'(o_w^H)) - \ln(v'(o_w^L)) - \ln(\gamma^H) + \ln(\gamma^L).$$

Estimates of $\ln(v'(o_w))$ can be obtained using linear regression based on Equation (7).⁵ Having established estimates of $v'(o_w)$ for all relevant outcome levels, identification of π_w is straightforward. Specifically, with $\pi_w = \gamma^c s_w^c / v'(o_w^c)$ from Equation (5), with s_w^c known, an estimate of $v'(o_w^c)$ available, and using $\sum \pi_w = 1$, we find $\gamma^c = 1 / \sum_w (s_w^c / v'(o_w^c))$ and hence π_w is identified. Having established identification, log-linearized versions of the original first order conditions in Equation (5) can be estimated using linear regression:

$$(8) \quad \ln(s_w^c) = \ln(\pi_w) + \ln(v'(o_w^c)) - \ln(\gamma^c), \text{ for } c=H,L \text{ and } w=1, \dots, W.$$

This enables estimation of $\ln(v'(o_w))$ and $\ln(\pi_w)$ jointly, based on the information contained in the two preferred outcome distributions. The analysis in Section 4 is performed using this estimation procedure.

3.3 State prices and preferred outcome distributions

For the construction of state prices we follow Sharpe et al. (2000) and Sharpe (2001); details are given in the appendix. Here we focus on the effect of varying state prices on individuals' preferred outcome distributions.

One of the characteristics of state prices is that under a risk free rate r , the average state price equals $1/(1+r)$ in any scenario, as a safe investment that pays off 1 in every possible state of the world costs $1/(1+r)$. When risk is introduced, state prices vary, and the effect of an increase in expected asset returns, while keeping the variance fixed, is to increase the differences between state prices. The prices for relatively bad states of the world become higher, while the prices for the good states become lower. As a result, the benefits of accepting lower payoffs are larger in

⁵ Note that the “intercept” in equation (7) is not identified. In combination with the log-transformation on the marginal utilities, this corresponds to a lack of identification of the scale of the utility function, which is not surprising.

case of high expected returns (as compared to low expected returns), because (i) accepting an increased risk of having a low payoff frees up more budget (as the state prices for bad states are higher) and (ii) it is less expensive to increase the odds of having a high outcome (as state prices for good states of the world are lower).

Equation (7) encapsulates this intuition. In the high returns scenario, the state prices for the bad states of the world are higher, and so $\ln(s_w^H) - \ln(s_w^L) > 0$. Assuming risk aversion, i.e. v' is decreasing, it is optimal to put the markers corresponding to these states at lower values in the high return scenario, since $o_w^H < o_w^L$ implies that $v'(o_w^H) > v'(o_w^L)$. The opposite holds for the markers that represent good states of the world, and so $\ln(s_w^H) - \ln(s_w^L) < 0$. Hence, compared to the low returns scenario, it is optimal to assign a marker for a good state of the world to a higher outcome (i.e. an outcome where marginal value is lower) when returns are high. As a result, individuals' preferred outcome distributions in a high returns scenario are expected to be more dispersed than in a low returns scenario, so more risk is accepted.

4. Empirical application: Individuals' investment decisions

4.1. Design of the lab study

To illustrate the use of our method, we set up a lab study in which preferred outcome distributions were used to measure individuals' value functions and probability weighting functions. Participants were told that the experiment concerned preferences for financial decisions and they were asked to consider they had an investment budget of 1,000 Euros. They were also instructed to consider investment decisions in which a person would expect relatively high or relatively low average returns and were asked to reveal their preferences under both scenarios; the order of the scenarios being randomized. Next, a video was shown that illustrated

the use of the DB. Test questions were asked to ensure participants had a proper understanding of outcome distributions and the use of the DB.⁶

The DB used to construct preferred outcome distributions was configured as follows. There were 100 markers that could be dragged and dropped onto various outcome levels one by one, or in blocks of up to 100 at a time. A total of 21 possible outcome levels were used, ranging from 0 Euros to 2,000 Euros, in steps of 100 Euros. The interactive graphical interface of the DB in Figure 1 corresponds to this set-up.

To induce variation in state prices, we used two investment scenarios, a high returns scenario and a low returns scenario. We generated state prices that correspond to a risk free rate of 2% and a mean return of 6% in the high returns scenario and 4% in the low returns scenario, in combination with the assumption of log-normally distributed returns and a standard deviation of 18%. These parameters are well within the range of values observed in the literature, see Van Ewijk et al. (2012) for a meta-analysis.

Participants received a fixed monetary award of 5 Euros upon participation. In addition, a BDM mechanism ensured the task was incentive-compatible. In particular, participants were told that one participant would be randomly selected to win an additional prize. Each of the preferred outcome distributions of this participant could be selected randomly and played out for real. In the context of the DB, this meant that one of the 100 markers was randomly selected and the corresponding prize was paid out. The participant who was selected to win the additional prize received the prize corresponding to the 22nd marker in her low return scenario preferred outcome distribution, which corresponded to a payoff of 800 Euros.

In the DB task itself, the procedure in Goldstein et al. (2008) was used as follows. First, participants constructed a distribution that satisfied their preferences and met the budget

⁶ The instruction video is available as part of the submitted materials.

constraint. In the next step, one of the states of the world was randomly drawn and the corresponding payoff was revealed to the participant erasing all other markers from the screen, one at a time, in random order. Three rounds followed in which participants were allowed to change their preferred outcome distributions, and experience the outcomes. Before the fifth round, they were told that they were about to *submit* their final preferred outcome distribution. We use this last submitted distribution as the preferred outcome distribution for the first scenario. Finally, participants were asked to reveal their preferences under the remaining scenario for investment returns (either low or high depending on what scenario they were exposed to first). The same procedure as above was repeated, but with a different set of underlying state prices, thus leading to two outcome distributions available per participant. A total of 137 undergraduate students participated in the study, from which 3 participants (2%) stopped before they constructed a preferred outcome distribution.

4.2 Results: Value function

We estimate the value function for each participant from the two preferred outcome distributions submitted in the final round of each scenario. Figure 4 illustrates the resulting average value function (which linearly extrapolates the value functions towards zero and assumes zero marginal value above the highest placed marker). We also add a linear value function for comparison that is represented by the dashed line. We find that on average individuals are fairly risk neutral for low outcomes and risk averse at the higher outcome levels.




Figure 4.
Mean of estimated value function

To obtain more detailed insights into individual differences in value functions we can characterize them by evaluating their shape, that is, whether they are convex, concave, first convex then concave or vice versa (Abdellaoui et al. 2007, Abdellaoui et al. 2008). From the DB, we are able to obtain estimates of marginal utilities and the question becomes whether they are decreasing or increasing, corresponding to a concave or a convex value function, respectively. We also consider the presence of reference points. A likely reference point in our study is the initial investment amount of 1,000 Euros, which is also the middle value in the range of possible

non-zero outcomes.⁷ This would suggest that outcomes below 1,000 Euro are perceived as losses while those above 1,000 Euro are perceived as gains. As a consequence, the shape of the value function might be different in the range between 0 Euros and 1,000 Euros, compared to the range of outcomes above 1,000 Euros.

We classify a value function as concave (convex) if marginal utility values show a decreasing (increasing) pattern. This procedure enables us to classify individual value functions as concave, convex, S-shaped or inverse S-shaped, depending on the shape of the value function below and above 1,000 Euros. In case the estimated value function does not deviate more than 0.05 (on a scale from 0 to 1) from a linear value function, we classify the participant as having a linear value function. We require 80% of the subsequent marginal utility levels to be increasing (or decreasing) for the corresponding classification to be assigned.

Based on this classification, we obtain the following results. A combination of a concave and a convex section in the value function characterizes 22% of all individuals, who either have an S-shaped (16%) or an inverse S-shaped value function (6%). The other participants' value functions can be classified as follows: the majority of the participants' value functions is concave (31%), a substantial number is close to linear (17%), and a minority as convex (10%). The remaining subjects (18%) could not be classified with this procedure.

4.3 Results for the probability weighting function

The estimate of the shape of the PWF is obtained from the estimated weights assigned to each of the 100 markers in the DB. Figure 5 depicts the resulting PWF, averaged across all participants (median values are similar). It shows that participants are somewhat pessimistic, assigning

⁷ In our study, the outcome of 0 Euro could also serve as a reference point, as individuals can incur losses only to the extent of foregone gains. In that case the shape of the utility function is not expected to change near the outcome level of 1000 Euro.

relatively low weights to the best outcomes and high weights to the worst outcomes, as reflected by the flatter slope for small probabilities and the steeper slope for large probabilities. The fact that the weights corresponding to expected utility do not lie between the dotted lines, which represent the 10th and 90th percentile of the PWF for each probability level, suggests that this pessimism is widespread.

It is worth noting that we find no evidence of overweighting of small probabilities for high outcomes, at least for the average individual. This is in line with the results of Van de Kuilen and Wakker (2011), who also find much less evidence for subadditivity than earlier studies. For comparison, the dash-dotted line in Figure 3 indicates the estimated probability weighting function from Van de Kuilen and Wakker, who used lotteries to elicit risk preferences. This figure shows that insensitivity to probability for high outcomes is much stronger in the lottery based approach than in the DB. For the worst outcomes, the decision weights applied in the DB-based approach are comparable to those of the lottery based approach. As a result, the average cumulative probability weights that we infer using the DB-based approach are much closer to their objective counterparts than those from the lottery based study of Van de Kuilen and Wakker (2011).




Figure 5.

Mean of estimated probability weighting functions (solid line)

The dashed line represents expected utility and the dotted lines are the 10th and 90th percentiles of the estimated probability weighting functions in the sample. The dash-dotted line is the estimate of Van de Kuilen and Wakker (2011).

To characterize the PWF at the individual level, we use the classification proposed by Bleichrodt and Pinto (2000) that builds on estimates of $\pi(p)$, for $p=0.1, 0.25, 0.5, 0.75,$ and 0.9 . The resulting classification is presented in Table 2, where totals do not add up to 100% because some individuals do not match a description, while others match multiple. The vast majority of individuals have a convex PWF (76.9%), while only a few individuals have a concave PWF

(6.0%) and no individual PWF could be classified as linear. Moreover, the majority of PWFs exhibit upper subadditivity, thus providing evidence for a “certainty effect.”

Shape	Proportion of subjects
Concave	6.0%
Linear	0.0%
Convex	76.9%
Lower subadditivity	22.4%
Upper subadditivity	88.8%
Lower and upper subadditivity	14.2%

Table 2.
Classification of subjects' probability weighting functions

Note: numbers do not add up to 100% because some individuals do not match any description, while others match multiple, for example concave and lower subadditivity.

Thus, in contrast to several previous findings (see e.g. Gonzalez and Wu 1999 and Prelec 1998), subjects who used the DB do not display the insensitivity to probabilities in the middle range. In general, the probability weights are close to the actual probabilities, so participants' behavior is close to normative behavior that is characterized by expected utility (cf. Bleichrodt et al 2001). This is a desirable aspect of the DB as a measurement tool and is in line with the fact that in the DB all markers – and hence the whole range of probabilities – are equally salient within the tool.

4.4 Individual level results

To further illustrate the tool and the type of individual-level results that can be obtained, Figure 6 illustrates the results for a single individual. The top panel shows the preferred outcome distributions as they have been constructed by this individual. On the left is the distribution for

the low returns scenario and on the right is the one for the high returns scenario. This latter distribution is more dispersed, which indicates that more risk is taken.

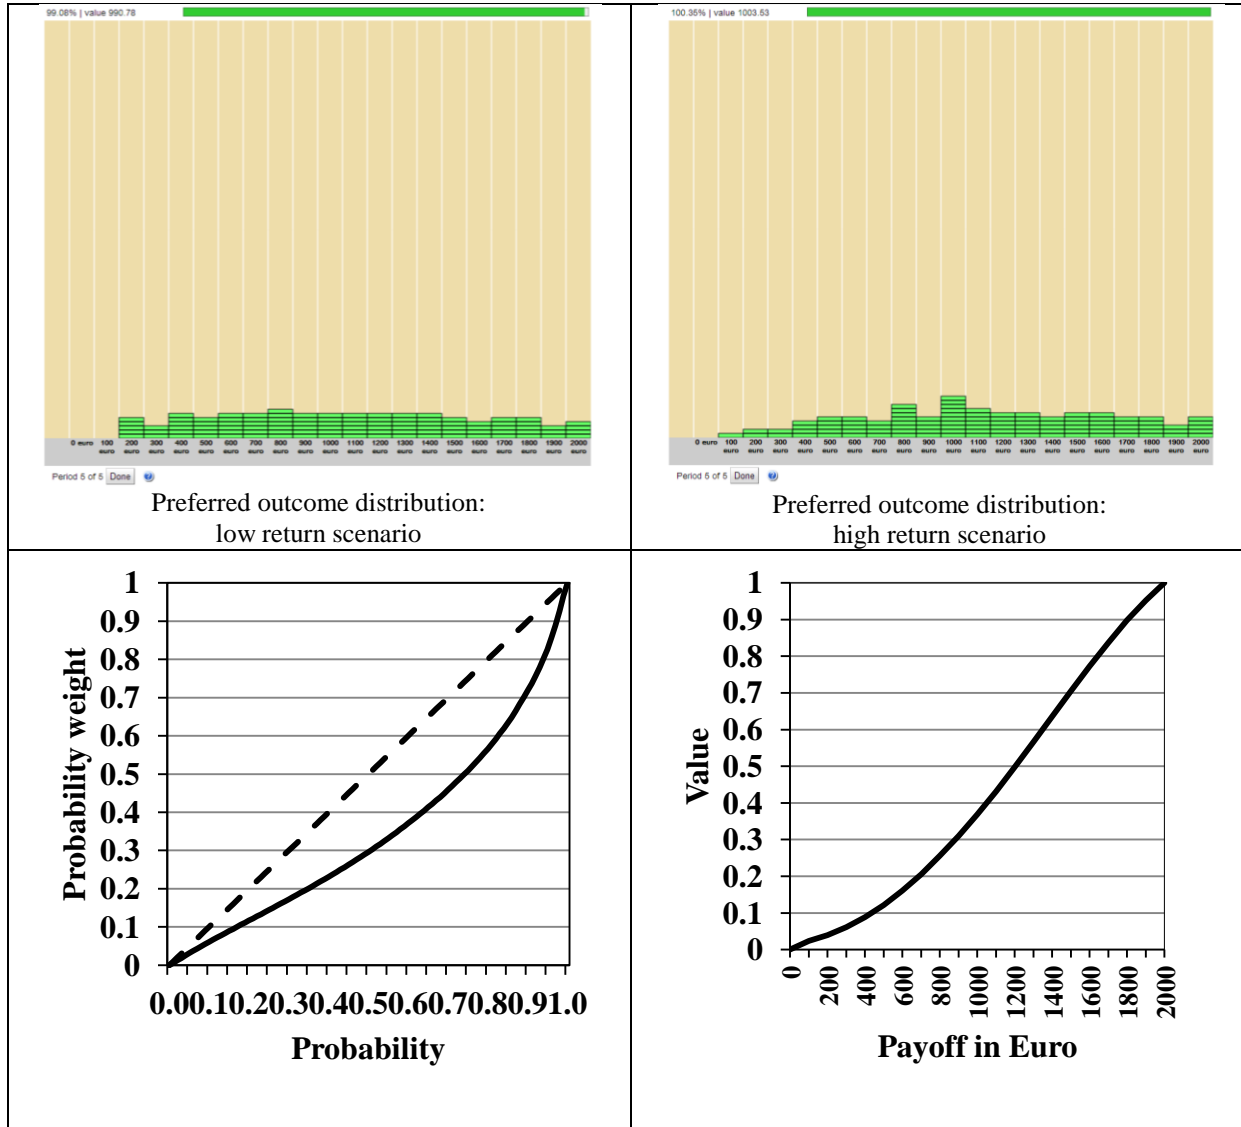


Figure 6.
Individual level analysis

The top panels show the preferred outcome distributions of one participant for the two scenarios. The bottom panels show the inferred probability weighting function and value function for that same participant.

The bottom panel of Figure 6 shows the estimated probability weighting function and value function for this individual. The PWF shows a strong degree of pessimism, although also the highest outcomes receive about the weight they should get. This is shown by the fact that the

estimated PWF is close to the normative weighting function in the bottom left of the graph. The value function shows risk seeking behavior for low outcomes and slightly risk-averse behavior for high outcomes.

4.5 Additional results on participants' usage of the DB

Because the DB is a new method of elicitation of risk preferences with several different features compared to existing tools, we report a few characteristics of the process participants went through. On average, participants took approximately 20 minutes to complete the entire DB task, which involved the construction of two preferred outcome distributions: 9 minutes watching the video and answering to questions to ensure understanding of the DB, 7 minutes completing the first scenario (one preferred outcome distributions and 4 preceding training rounds), and 4 minutes completing the second scenario (one preferred outcome distributions and 4 preceding training rounds).

To assess whether participants find it difficult to use the Distribution Builder, we used a three-item Likert scale. Specifically, participants stated their level of (dis)agreement with each of the following three statements: *“Using this tool to assess my risk preferences will be complicated,”* *“Using this tool to assess my risk preferences will be confusing,”* and *“Using this tool to assess my risk preferences will take a lot of effort”* (see Dabholkar 1994). The average rating on a five-point scale across participants and items was 2.9, suggesting the DB is not perceived to be a complicated tool in the context of decision making under risk.

5. Discussion and conclusion

This paper proposes and tests the use of preferred outcome distributions and the Distribution Builder to elicit individual-level value functions and probability weighting functions in the

context of rank-dependent utility models. A key benefit of our proposed approach is that it requires individuals to jointly take into account probabilities and outcomes when considering decisions under risk. The DB task is also far less repetitive than lottery-based tools.

The empirical results on the shape of the probability weighting function suggest that probability weighting functions appear to be more sensitive to framing and ways of measurement than the utility function, see also the discussion section in Van de Kuilen and Wakker (2011). In case one aims at estimating value functions while accounting for the impact of probability weighting, it will therefore be important to elicit the probability weighting function and the value function from very similar tasks, as we have done with the DB. Our method of elicitation of risk preferences always relies on the same task in terms of probabilities and outcomes, ensuring a single decision process. Methods that use chained lottery tasks that first keep outcomes fixed to learn the probability weighting function, and in a second stage keep probabilities fixed to learn the value function, might induce participants to shift their focus throughout the decision process.

A promising finding is the relatively small deviation of the estimated probability weights from the objective weights, at least when compared to many studies that have used lotteries to elicit probability weighting functions. Perhaps the frequency representation of probabilities in the DB, or the visually-experienced random drawing of events facilitates learning about and understanding of risk (Barron and Erev 2003, Kaufmann, Weber and Haisley 2012). The availability of a tool like the DB, with which individuals show only small deviations between estimated probability weights and objective weights, is important for many applications because it provides not only a risk elicitation tool but also a risk communication tool to support and improve individuals' decision making under risk. In these cases, normative, expected utility based advice on optimal investments or medical treatments that correct for distorted probability

weights (see Bleichrodt et al 2001 and references therein) would not be in conflict with stated preferences, which is likely to improve advice acceptance (Harvey and Fischer 1997).

The Distribution Builder is thus a promising tool not only for scholars but also for practitioners, as its interactive graphical interface makes the elicitation of risk preferences an intuitive, simple – even entertaining – task for participants. In many applications, practicality and ease of use are two crucial criteria for adoption of new elicitation techniques. In this way, the DB offers a promising link between academic research on decision making under risk and practitioners providing advice to help individuals make better investment decisions. These are potential contributions of our work that we leave for future research.

References

Abdellaoui, M (2000) Parameter-free elicitation of utility and probability weighting functions.

Management Science 46:1497–1512.

Abdellaoui, M, Bleichrodt H, Paraschiv C (2007) Loss Aversion Under Prospect Theory.

Management Science 53(10):1659–1674.

Abdellaoui, M, Bleichrodt H, L’Haridon O (2008) A tractable method to measure utility and loss aversion under prospect theory. *Journal of Risk and Uncertainty* 36:245–266.

Ahmed, H, Naik G, Willoughby H, Edwards AGK (2012) Communicating risk. *British Medical Journal: Clinical Review* 1-7.

Ancker, JS, Weber EU, Kukafka R (2011) Effects of game-like interactive graphics on risk perceptions and decisions. *Medical Decision Making* 130-142.

Arrow KJ (1964) The Role of Securities in the Optimal Allocation of Risk Bearing. *Review of Economic Studies* 31:91-96.

- Barron G, Erev I (2003) Small Feedback-based Decisions and Their Limited Correspondence to Description-based Decisions. *Journal of Behavioral Decision Making* 16:215-233.
- Bleichrodt H, Pinto JL (2000) A parameter-free elicitation of the probability weighting function in medical decision analysis. *Management Science* 46:1485–1496.
- Bleichrodt H, Pinto JL, Wakker PP (2001) Making Descriptive Use of Prospect Theory to Improve the Prescriptive Use of Expected Utility. *Management Science* 47(11):1498-1514.
- Camilleri AR, Newell BR (2011) When and why rare events are underweighted: A direct comparison of the sampling, partial feedback, full feedback and description choice paradigms. *Psychonomic Bulletin Review* 18:377-384.
- Choi S, Fisman R, Gale D, Kariv S (2007) Consistency and Heterogeneity of Individual Behavior under Uncertainty. *American Economic Review* 97(5):1921-1938.
- Dabholkar PA (1994) Incorporating Choice Into an Attitudinal Framework: Analyzing Models of Mental Comparison Processes. *Journal of Consumer Research* 21 June:100-118.
- Debreu G (1959) *The Theory of Value*. New York: John Wiley.
- Fagerlin, A., Zikmund-Fisher BJ, Ubel PA (2011) Helping patients decide: Ten steps to better risk communication *Journal of the National Cancer Institute* 103:1-8.
- Gigerenzer G, Galesic M (2012) Why do single event probabilities confuse patients? *British Medical Journal: Observations* 1-3.
- Goldstein DG, Johnson EJ, Sharpe WF (2008) Choosing Outcomes versus Choosing Products: Consumer-Focused Retirement Investment Advice. *Journal of Consumer Research* 35:440-456.

- Gonzalez R, Wu G (1999) On the Shape of the Probability Weighting Function. *Cognitive Psychology* 38:129-166.
- Harrison GW, Rutstrom EE (2008) Risk Aversion in the Laboratory. *Research in Experimental Economics* 12:41–196.
- Harvey N, Fischer I (1997) Taking Advice: Accepting Help, Improving Judgment and Sharing Responsibility. *Organizational Behavior and Human Decision Processes* 70(2):117–133.
- Holt CA, Laury SK (2002) Risk Aversion and Incentive Effects. *American Economic Review* 92(5):1644-1655.
- Hershey JC, Schoemaker PJH (1985) Probability versus Certainty Equivalence Methods in Utility Measurement: Are They Equivalent? *Management Science* 31(10):1213-1231.
- Hertwig R, Barron G, Weber EU, Erev I (2004) Decisions from experience and the effect of rare events in risky choice. *Psychological Science* 15(8):534-539.
- Kaufmann C, Weber M, Haisley E (2013) The role of experience sampling and graphical displays on one's investment risk appetite. *Management Science* 59(2):323-340.
- McCord M, De Neufville R (1986) "Lottery Equivalents": Reduction of the Certainty Effect Problem in Utility Assessment. *Management Science* 31(1):56-60.
- Prelec D (1998) The Probability Weighting Function. *Econometrica* 66:497–527.
- Quiggin J (1981) Risk perception and risk aversion among Australian farmers. *Australian J. Agricultural Economics* 25:160-169.
- Sharpe WF (2001) Individual Risk and Return Preferences: A Preliminary Survey. Working paper available at www.stanford.edu/~wfsarpe/art/rrsurvey/vienna2001.htm.

- Sharpe WF, Goldstein DG, Blythe PW (2000) The Distribution Builder: A Tool for Inferring Investor Preferences. Working paper at www.stanford.edu/~wfsharpe/art/qpaper/qpaper.html.
- Tversky A, Kahneman D (1992) Advances in Prospect Theory: Cumulative Representation of Uncertainty. *Journal of Risk and Uncertainty* 5:297-323.
- Van de Kuilen G, Wakker PP (2011) The Midweight Method to Measure Attitudes Toward Risk and Ambiguity. *Management Science* 112:582-598.
- Van Ewijk C, De Groot HLF, Santing C (2012) A Meta-Analysis of the Equity Premium. *Journal of Empirical Finance* 19:819-830.
- Wakker PP, Deneffe D (1996) Eliciting von Neumann-Morgenstern utilities when probabilities are distorted or unknown. *Management Science* 42(8):1131–1150.

Appendix A. State price calculation

To induce variation in state prices we use state prices corresponding to the following two sets of economic parameters. In the *High* returns scenario, the expected return on risky assets (stocks) is 6% and the standard deviation of the risky return is 18%, while the risk free rate is 2%, so the equity premium is 4%. In the *Low* returns scenario, the expected return is only 4% while all other parameters are kept the same. This low return might be more in line with the recent economic downturn and at the time of the experiment. The exposition that follows is based on Sharpe (2001) and the reader is referred to that work for an extensive motivation and more detailed exposition of the approach.

In a world with only one risky asset and a risk-free asset, the wealth level in the economy – and hence the state price – is determined by the payoff of the risky asset. To obtain 100 states of the world, we selected the 5, 15, ..., 99.5 percentiles from a log-normal distribution with parameters that correspond to the return process of the risky asset outlined above. For the corresponding state prices we relied on a specification that is log-linear in prices and in the wealth level in each state of the world.

In particular, we assume

$$\ln(s_w) = \alpha + \beta \ln(W_w).$$

Or equivalently

$$s_w = \exp(\alpha)(W_w)^\beta.$$

In addition, the state prices need to satisfy a number of restrictions, namely:

1. An investment that has a payoff of 1 unit in every state of the world will cost $1/(1 + r^f)$ with r^f denoting the risk free rate, so $\sum s_w = 1/(1 + r^f)$.

2. A 1 unit investment in the stock should cost its (risk) discounted present value, so

$$\sum s_w W_w = 1.$$

From restriction 1, one obtains $\exp(\alpha) \sum (W_w)^\beta = 1/(1 + r^f)$. From restriction 2 one obtains $\exp(\alpha) \sum (W_w)^{\beta+1} = 1$. These can be combined into a single equation with one unknown, β , which is solved numerically. Substituting this back into one of the two restrictions provides the value for α .