

Searching for the Reference Point

AURÉLIEN BAILLON, HAN BLEICHRODT, AND VITALIE SPINU

Erasmus School of Economics,
Erasmus University Rotterdam, The Netherlands

July 2016

Abstract

Reference-dependence plays a key role in explaining people's choices. A fundamental problem of reference-dependent theories, like prospect theory, is that they do not tell us what the reference point is. Our paper addresses this problem. We assume a comprehensive reference-dependent model that includes the main reference-dependent theories as special cases and that allows isolating the reference point rule from other behavioral parameters. We use Bayesian hierarchical modeling to estimate the (posterior) probability that a specific reference point rule was used. In a high-stakes experiment with payoffs up to a weekly salary, we found that most subjects used the status quo or MaxMin (the maximum of the minimal outcomes of the prospects in a choice) as their reference point. Twenty percent of the subjects used an expectations-based reference point as in the influential model of Köszegi and Rabin (2006, 2007).

Key words: reference point formation, reference-dependence, Bayesian hierarchical modeling, large-stake experiment.

JEL code: D81, C91.

1. Introduction

A key insight of behavioral economics is that people evaluate outcomes as gains and losses from a reference point. Reference-dependence is central in prospect theory, currently the most influential theory of decision under risk, and it plays a crucial role in explaining people's attitudes towards risk (Rabin 2000; Wakker, 2010). Evidence abounds, from both the lab and the field, that preferences are reference-dependent.¹

A fundamental problem of prospect theory and other reference-dependent theories is that they are silent about how reference points are formed. Back in 1952, Markowitz (1952, p.157) already remarked about customary wealth, which plays the role of the reference point in his analysis, that "It would be convenient if I had a formula from which customary wealth could be calculated when this was not equal to present wealth. But I do not have such a rule and formula." This silence is undesirable as it creates too much freedom in deriving predictions, making it impossible to rigorously test reference-dependent theories empirically.² In his recent review of the literature, more than 60 years after Markowitz, Barberis (2013) still concludes that addressing the formation of the reference point is a key challenge to apply prospect theory to economics (p.192).

The leading theory of reference point formation was proposed by Köszegi and Rabin (2006, 2007). Köszegi and Rabin argue that the reference point is determined by

¹ Examples of real-world evidence for reference-dependence are the equity premium puzzle, the finding that stock returns are too high relative to bond returns (Benartzi and Thaler 1995), the disposition effect, the finding that investors hold losing stocks and property too long and sell winners too early (Odean 1998, Genesove and Mayer 2001), default bias in pension and insurance choice (Samuelson and Zeckhauser 1988, Thaler and Benartzi 2004) and organ donation (Johnson and Goldstein 2003), the excessive buying of insurance (Sydnor 2010, the annuitization puzzle, the fact that at retirement people allocate too little of their wealth to annuities (Benartzi et al. 2011), the behavior of professional golf players (Pope and Schweitzer 2011) and poker players (Eil and Lien 2014), and the bunching of marathon finishing times just ahead of round numbers (Allen et al. forthcoming).

² For example, different assumptions about the reference point have to be made to explain two well-known anomalies from finance: the equity premium puzzle demands that the reference point adjusts over time, whereas the disposition effect demands that the reference point remains constant over time at the purchase price.

people's (rational) expectations. Their model for the first time made the reference point operational and gave testable implications. It is close in spirit to the disappointment models of Bell (1985), Loomes and Sugden (1986), Gul (1991), and Delquié and Cillo (2006) in which decision makers also form expectations about uncertain prospects and experience elation or disappointment depending on whether the actual outcome is better or worse than those expectations.³

Empirical evidence on the formation of reference points is scarce and what is available gives mixed conclusions. Some evidence is consistent with Köszegi and Rabin's model of expectations-based reference points (Abeler et al. 2011, Card and Dahl 2011, Crawford and Meng 2011, Gill and Prowse 2012, Bartling et al. 2015), but other is not (e.g. Baucells et al. 2011, Allen et al. forthcoming, and Lien and Zheng 2015). Moreover, the evidence that is consistent with Köszegi and Rabin's model does not necessarily exclude other reference point rules. To illustrate, in Appendix A we show that the data of Abeler et al. (2011) are also consistent with MaxMin, a security-based rule where subjects adopt the minimum outcome that they can reach for sure as their reference point. Barberis (2013) concludes that in finance there are "natural reference points other than expectations." Evidence from medical decision making suggests that, instead of using an expectations-based reference point, people adopt the MaxMin rule described above to determine their reference point (Bleichrodt et al. 2001, van Osch et al. 2004, van Osch et al. 2006).

This paper addresses the question which reference point people adopt in decision under risk. We performed an experiment in Moldova, an Eastern European country, with large stakes up to a weekly salary. Guided by the available literature, we specified

³ Other models of reference point formation were proposed by Heath et al. (1999), who suggested that people use goals as their reference points and by Diecidue and Van de Ven (2008), who presented a model with an aspiration level, which is a form of reference dependence.

several reference point rules, including Köszegi and Rabin's expectations-based reference point rule, MaxMin, and the status quo, which is often used as a reference point in experiments. The selected rules vary depending on whether they are choice-specific (the reference points is determined by the choice set) or prospect-specific (the reference point is determined by the prospect itself), stochastic or deterministic, and on whether they are defined only by the outcome dimension or by both the outcome and the probability dimension.

All the reference points that we consider can be identified through choices. Hence, we work within the revealed preference paradigm and do not require introspective data. In this we follow Rabin's (2013) approach to develop more realistic theories that are maximally useful to core economic research. Rabin argues that new models should be "portable" and use the same independent variables as existing models. In decision under risk, economic models, like expected utility, use probabilities and outcomes as independent variables. Tversky and Kahneman's (1992) prospect theory is not portable because it leaves the reference point unspecified. By contrast, all our reference point rules can be derived from probabilities and outcomes and are portable.

We define a comprehensive reference-dependent model that includes the main reference-dependent theories as special cases. This makes it possible to compare reference point rules *ceteris paribus*, i.e. to isolate the reference point rule from the specification of the other behavioral parameters in the model. We use a Bayesian hierarchical model to estimate each subject's reference point rule. Bayesian hierarchical modeling estimates the parameters of each individual separately, but accounts for their similarities in the population. This leads to more precise estimates and prevents inference from being dominated by outliers (Rouder and Lu 2005, Nilsson et al. 2011).

Our results indicate that two reference point rules stand out: the status quo and MaxMin. Together these two reference points account for the behavior of over sixty percent of our subjects. Around twenty percent of our subjects use an expectations-based reference point, as suggested by Köszegi and Rabin (2006, 2007).

2. Theoretical background

A *prospect* is a *probability distribution* over money amounts. *Simple prospects* assign probability 1 to a finite set of outcomes. We denote these simple prospects as $(p_1, x_1; \dots; p_n, x_n)$, which means that they pay $\text{€}x_j$ with probability $p_j, j = 1, \dots, n$. We identify simple prospects with their cumulative distribution functions and denote them with capital Roman letters (F, G) . The decision maker has a weak preference relation \succsim over the set of prospects and, as usual, we denote strict preference by \succ , indifference by \sim , and the reversed preferences by \preccurlyeq and \prec . The function V defined from the set of simple prospects to the reals *represents* \succsim if for all prospects $F, G, F \succsim G \Leftrightarrow V(F) \geq V(G)$.

Outcomes are defined as gains and losses relative to a *reference point* r . An outcome x is a *gain* if $x > r$ and a *loss* if $x < r$.

2.1. Prospect theory

Under prospect theory (Tversky and Kahneman 1992), there exist *probability weighting functions* w^+ and w^- for gains and losses and a non-decreasing *gain-loss utility function* $U: \mathbb{R} \rightarrow \mathbb{R}$ with $U(0) = 0$ such that preferences are represented by

$$F \rightarrow PT_r(F) = \int_{x \geq r} U(x - r) dw^+(1 - F) + \int_{x \leq r} U(x - r) dw^-(F). \quad (1)$$

The integrals in Eq. (1) are Lebesgue integrals with respect to distorted measures $w^+(1 - F)$ and $w^-(F)$. For losses, the weighting applies to the cumulative distribution (F) but for gains it is applied to the decumulative distribution ($1 - F$).

The functions w^+ and w^- are non-decreasing and map probabilities into $[0,1]$ with $w^i(0) = 0, w^i(1) = 1, i = +, -$. When the w^i are linear, *PT* reduces to *expected utility* with referent-dependent utility:

$$F \rightarrow EU_r(F) = \int U(x - r)dF . \quad (2)$$

Equation (2) shows that reference-dependence by itself does not violate expected utility as long as the reference point is held fixed.

Based on empirical observations, Tversky and Kahneman (1992) hypothesized specific shapes for the functions U, w^+ , and w^- . The gain-loss utility U is S-shaped, concave for gains and convex for losses. It is steeper for losses than for gains to capture loss aversion, the finding that losses loom larger than gains. The probability weighting functions are inverse S-shaped, reflecting overweighting of small probabilities and underweighting of middle and large probabilities.

2.2. Köszegi and Rabin's model

Tversky and Kahneman (1992) defined prospect theory for a riskless reference point r . Köszegi and Rabin (2006, 2007) added two elements to prospect theory. First, they distinguished the economic concept of consumption utility and the psychological concept of gain-loss utility and, second, they allowed for random reference points. Let R be the random reference point. In Köszegi and Rabin's model preferences over prospects F are represented by

$$F \rightarrow KR_R(F) = \int v(x)dF + \int \int U(v(x) - v(r))dFdR. \quad (3)$$

In Eq. (3), v represents consumption utility and U is the gain-loss utility function, which reflects the psychological part of utility. Consumption utility does not depend on the reference point, but only on the absolute size of the payoffs.

Kőszegi and Rabin (2007, p.1052) argue that “for modest-scale risk, such as \$100 or \$1000,[...] consumption utility can be taken to be approximately linear”. As the incentives in our experiment did not exceed \$100 (in PPP) and the prospects in the different choice sets had approximately equal expected value, we concentrate on the gain-loss function U and take $v(x) = x$:

$$KR_R(F) = \int x dF + \int EU_r(F) dR. \quad (4)$$

The assumption of linear consumption utility is commonly made in empirical applications of Kőszegi and Rabin’s model (e.g. Heidhues and Kőszegi 2008, Abeler et al. 2011, Gill and Prowse 2012, Eil and Lien 2014).

There is no probability weighting in Eq. (4). It is unclear how the rational expectations reference point should be defined in the presence of probability weighting. Kőszegi and Rabin (2006, 2007) do not address this problem and leave out probability weighting, even though they acknowledge its relevance (Kőszegi and Rabin 2006, footnote 2, p. 1137).

While prospect theory is silent about the formation of the reference point, Kőszegi and Rabin (2007) assume that it is based on the decision maker’s rational expectations. They distinguish two specifications, one prospect-specific and one choice-specific. In a “*choice-acclimating personal equilibrium*” (CPE), the reference point is the prospect itself. This prospect-specific reference point gives:

$$KR(F) = \int x dF + \int EU_r(F) dF. \quad (5)$$

A choice-specific variant is Köszegi and Rabin's model with an "unacclimating personal equilibrium" (UPE), which is Eq. (4) with the reference point equal to the preferred prospect in the choice set.

2.3. Disappointment models

Köszegi and Rabin's (2006, 2007) model is close to the disappointment models of Bell (1985), Loomes and Sugden (1986), Gul (1991), and Delquié and Cillo (2006). Bell's model is equivalent to Eq. (3) with $v(r)$ replaced by the expected value of the prospect (although Bell remarks that this may be too restrictive and also presents a more general model), Loomes and Sugden's model (1986) is equivalent to Eq.(3) with $v(r)$ replaced by the expected utility of the prospect,⁴ and Gul's (1991) model is equivalent to Eq.(3) with $v(r)$ replaced by the certainty equivalent of the prospect. Delquié and Cillo's (2006) model is identical to Köszegi and Rabin's (2007) CPE model (Eq. 5). Masatlioglu and Raymond (forthcoming) formally characterize the link between Köszegi and Rabin's (2007) CPE model, the disappointment models, and other generalizations of expected utility.

2.4 General reference-dependent specification

This paper compares the performance of different reference point rules in explaining behavior. To isolate the reference point from the rest of the model, we must use the same specification across all reference point rules. In other words, in the estimation of the reference point, all other behavioral parameters must enter the same way. To address this *ceteris paribus* principle we adopt the following general reference-dependent model:

⁴ With consumption utility v .

$$F \rightarrow RD(F) = \int x dF + \int PT_r(F) dR. \quad (6)$$

Eq. (6) contains prospect theory (Eq. 1), Köszegi and Rabin's (2006, 2007) model (Eq. 4) and the disappointment models as special cases. In Eq. (6), probability weighting plays a role in the psychological part of the model (the second term of the addition), but it does not affect consumption utility (the first term). This seems reasonable as consumption utility reflects the "rational" part of utility and probability weighting is usually considered irrational. Adjusting the model to include probability weighting in consumption utility is straightforward.

Probability weighting does not affect the formation of reference points either. In this we follow the literature on stochastic reference points (Sugden 2003, Delquié and Cillo 2006, Köszegi and Rabin 2006, Köszegi and Rabin 2007, Schmidt et al. 2008). We will consider alternative specifications in Section 6.4.

3. Reference point rules

A reference point rule specifies for each choice situation which reference point is used. Table 1 summarizes the reference point rules that we studied. We distinguish reference point rules along three dimensions. First, whether they are *prospect-specific* and determine a reference point for each prospect separately, or *choice-specific* and determine a common reference point for all prospects within a choice set. Second, whether they determine a random or a deterministic reference point. Third, whether they use only the payoffs to determine the reference point or both payoffs and probabilities.

	Prospect Specific	Random	Uses probability
Status Quo	Choice	No	No
MaxMin	Choice	No	No
MinMax	Choice	No	No
X at Max P	Choice	No	Yes
Expected Value	Prospect	No	Yes
Prospect Itself	Prospect	Yes	Yes

Table 1. The reference point rules studied in this paper

The first reference point is the *Status Quo*, which is often used in (experiments on) reference-dependence. In our experiment, subjects knew they would receive a participation fee for sure. We took the participation fee as the status quo reference point and any extra money that subjects could win if one of their choices was played out for real as a gain.

MaxMin, the second reference point rule, is based on Hershey and Schoemaker (1985). They found that when subjects were asked for the probability p that made them indifferent between outcome z for sure and a prospect $(p, x_1; 1 - p, x_2)$, they took z as their reference point and perceived $x_1 - z$ as a gain and $x_2 - z$ as a loss. Bleichrodt et al. (2001) and van Osch et al. (2004, 2006, 2008) found similar evidence for such a strategy in medical decisions. For example, van Osch et al. (2006) asked their subjects to think aloud while choosing. The most common reasoning in a choice between life duration z for sure and a prospect $(p, x_1; 1 - p, x_2)$ was: "I can gain $x - z$ years if the gamble goes well or lose $z - y$ if it doesn't."

MaxMin generalizes the above reasoning to the choice between any two prospects.⁵ In a comparison between two prospects, people look at the minimum outcomes of the two prospects and take the maximum of these as their reference point. This reference point is the amount they can obtain for sure. For example, in a comparison between

⁵ To the best of our knowledge this rule was first proposed in Bleichrodt and Schmidt (2005). See also Birnbaum and Schmidt (2010) and Schneider and Day (forthcoming).

(0.50,100; 0.50,0) and (0.25,75; 0.75,25), the minimum outcomes are 0 and 25, and because 25 exceeds 0 the rule implies that subjects take 25 as their reference point.

MaxMin is a cautious rule and implies people are looking for security. *MinMax* is the bold counterpart of MaxMin. A MinMax decision maker looks at the maximal opportunities and takes the minimum of the maximum outcomes as his reference point. Hence, MinMax predicts that the decision maker takes 75 as his reference point when choosing between (0.50,100; 0.50,0) and (0.25,75; 0.75,25).

The MaxMin and the MinMax rules both look at the extreme outcomes. One reason is that these outcomes are salient. Another salient outcome is the payoff with the highest probability and our next rule, *X at MaxP*, uses this outcome as the reference point. The importance of salience is widely-documented in cognitive psychology (Kahneman 2011). Barber and Odean (2008) and Chetty et al. (2009) show the effect of salience on economic decisions. Bordalo et al. (2012) present a theory of salience in decision under risk.

The final two reference points that we considered are the *expected value* of the prospect, as in the disappointment models of Bell (1985) and Loomes and Sugden (1986)⁶ and the *prospect itself* as in Köszegi and Rabin's (2007) CPE model and Delqu   and Cillo's (2006) disappointment model. In contrast with the other reference points considered so far, these reference points are prospect-specific. The prospect itself is the only rule that specifies a random reference point. We do not consider the rule that specifies that the preferred prospect in a choice is used as the reference point, as in K  szegi and Rabin's (2007) UPE model, because the model in Eq. (6) is then defined recursively and could not be estimated.

⁶ The equivalence with Loomes and Sugden (1986) follows because we assume $v(x) = x$.

4. Experiment

Subjects and payoffs

The subjects were 139 students and employees from the Technical University of Moldova (49 females, age range 17-47, average age 22 years). They received a 50 Lei participation fee (about \$4, which was \$8 in PPP at the time). In addition, to incentivize the experiment, each subject had a one third chance to play out one of their choices for real.

The payoffs were substantial. The subjects who played out their choices for real earned 330 Lei on average, which was more than half the average weekly salary in Moldova at the time of the experiment. Two subjects won about 600 Lei, the average weekly salary.

Procedure

The experiment was computer-run in group sessions of 10 to 15 subjects. Subjects took 30 minutes on average to complete the experiment including instructions.

Subjects made 70 choices in total. Each choice involved two options, Option 1 and Option 2. The options had between one and four possible outcomes. We randomized the order of the choices and we also randomized whether a prospect was presented as Option 1 or as Option 2.

Because we are interested in the MaxMin and the Minmax rules, we constructed choices with different maxima of the minima and with different minima of the maxima. Choices were created by an optimal design procedure that minimized their joint correlation. Using minimally correlated choices gives more precise and more robust estimates of the behavioral parameters, just like orthogonal covariates in linear regression do.

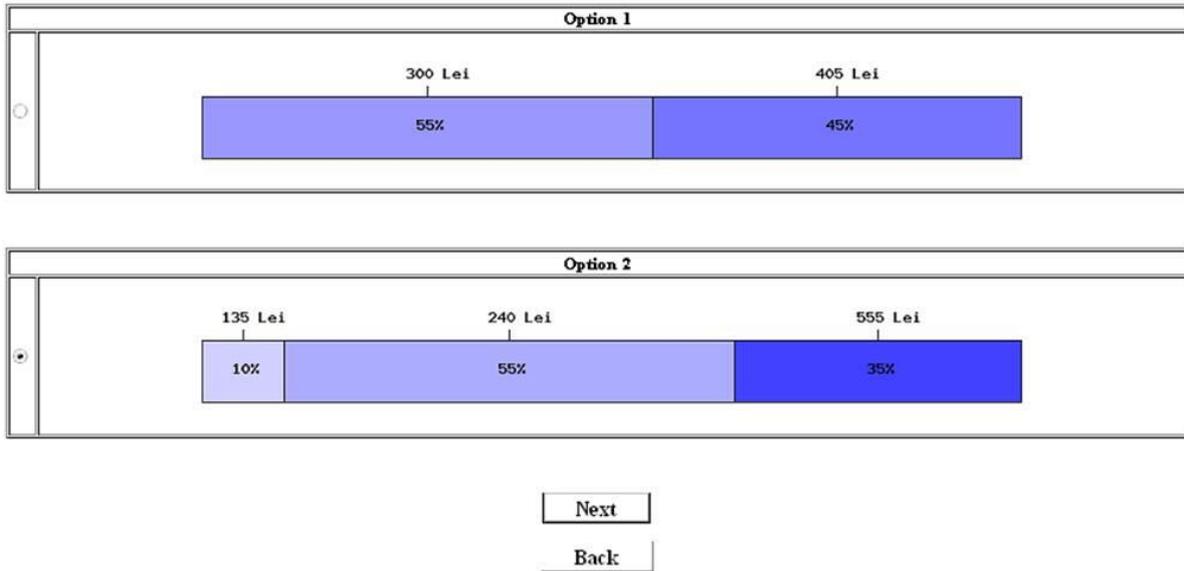


Figure 1. Presentation of the choices in the experiment

Figure 1 shows the display of the choices. Prospects were presented as horizontal bars with as many parts as there were different payoffs. The size of each part corresponded with the probability of the payoff and the intensity of the color (blue) was proportional to the size of the payoff. The payoffs were presented in increasing order. Subjects were asked to click on a bullet to indicate their preferred option (Figure 1 illustrates a choice for Option 2).

5. Bayesian hierarchical modeling

We analyzed the data using Bayesian hierarchical modeling. This is an appealing compromise between assuming a representative agent (thus ignoring individual heterogeneity) and treating all individuals as independent. Hierarchical models estimate the model parameters for each individual separately, but assume that they share similarities and that their individual parameter values come from a common

(population-level) distribution. Hence, the parameter estimates for one individual benefit from the information that is obtained from all others.

The common approach in behavioral decision making is to assume that all subjects are independent and to estimate the individual parameters by only using the choice data of that individual. Hierarchical modeling has two advantages over this approach. First, the data from the other respondents are used in the estimation of the individual parameters, which improves the precision with which these parameters can be estimated. In Bayesian statistics this is known as *collective inference*. A second advantage is that outliers have less impact on the results. Each individual parameter is shrunk towards the group mean, an effect that is stronger for individuals with noisier behavior, thus making the overall estimation more robust. This is particularly true for parameters that are estimated with lower precision. An example is the loss aversion coefficient in prospect theory, for which the standard deviation of the parameter estimates is usually high. Nilsson et al. (2011) illustrate that Bayesian hierarchical modeling leads to more accurate and more efficient estimates of loss aversion than the commonly-used maximum likelihood estimation.

Figure 2 shows a schematic representation of our statistical model. The model consists of two main parts: the specification of the core behavioral parameters and the specification of the reference point rule.

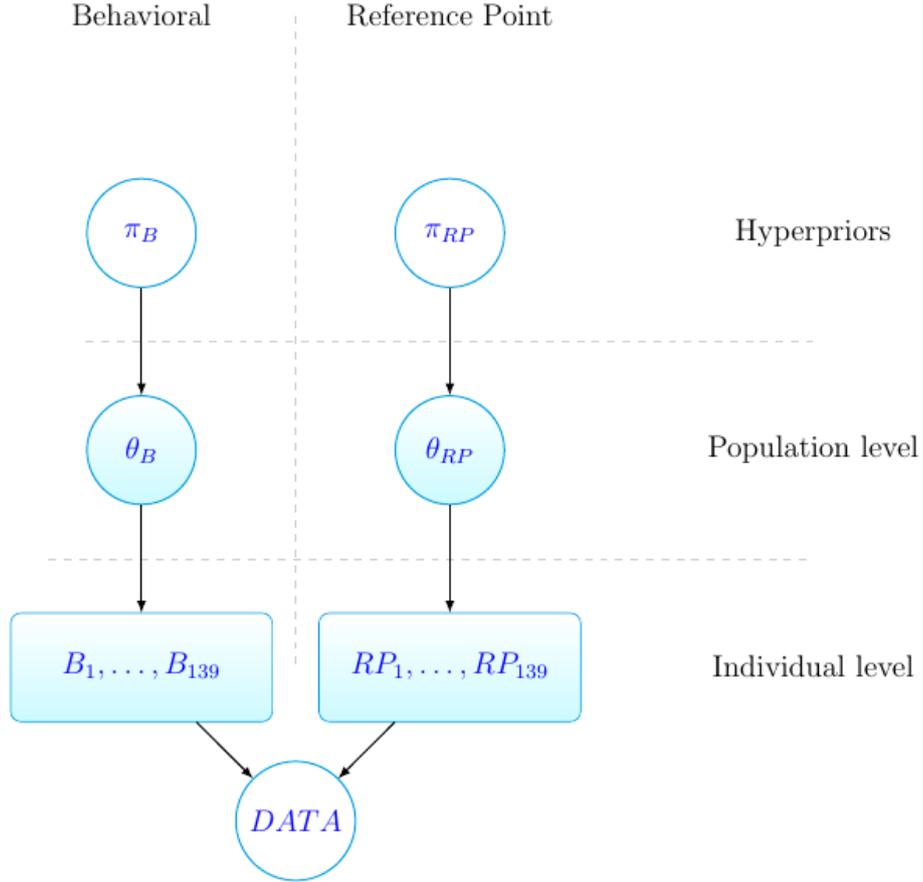


Figure 2. Graphical representation of our model. Non-shaded nodes are known or predefined quantities, shaded nodes are the unknown latent parameters.

We adopt the following mnemonic conventions. For individual $i \in \{1, \dots, 139\}$ the vector of core behavioral parameters is denoted B_i and his reference point rule is denoted RP_i . We assume that a subject uses the same reference point rule in all choices, where the reference point rule is one of the candidates listed in Table 1. The distribution of the core behavioral parameters in the population is parameterized by an unknown vector θ_B , and the distribution of reference point rules in the population is parameterized by an unknown parameter θ_{RP} which is also estimated from the data. These parameters also follow a distribution, but with a known shape. This final layer in the hierarchical specification is commonly referred to as a hyper-prior. The hyper priors

are denoted by π_B and π_{RP} respectively. The vector of the observed choices (data) of the individual i is denoted by $D_i = (D_{i1}, \dots, D_{i70})$. We will now describe our estimation procedure.

5.1. Specification of the behavioral parameters

We assume that the utility function U in Eq. (6) is a power function:

$$U(x) = \begin{cases} (x - r)^\alpha & \text{if } x \geq r \\ -\lambda(r - x)^\alpha & \text{if } x < r \end{cases} \quad (7)$$

In Eq. (7) the α -parameter reflects the curvature of utility and the λ -parameter indicates loss aversion. We assumed the same curvature for gains and losses. It is hard to estimate loss aversion when utility curvature for gains and for losses can both vary freely (Nilsson et al. 2011).

For probability weighting, we assumed Prelec's (1998) one-parameter specification:

$$w(p) = \exp(-(-\ln p)^\nu). \quad (8)$$

We used the same probability weighting for gains and losses. Empirical studies usually find that the differences in probability weighting between gains and losses are relatively small (Tversky and Kahneman 1992, Abdellaoui 2000, Kothiyal, Spinu, and Wakker, 2014).

To account for the probabilistic nature of people's choices we used Luce's (1959) logistic choice rule. Let $RD(F)$ and $RD(G)$ denote the respective values of prospects F and G according to our general reference-dependent model, Eq. (6). Luce's rule says that the probability $P(F, G)$ of choosing prospect F over prospect G equals

$$P(F, G) = \frac{1}{1 + e^{\xi[RD(G) - RD(F)]}} \cdot \quad (9)$$

In Eq. (9), $\xi > 0$ is a precision parameter that measures the extent to which the decision maker's choices are determined by the differences in value between the

prospects. In other words, the ξ -parameter signals the quality of the decision. Larger values of ξ imply that choice is driven more by the value difference between prospects F and G . If $\xi = 0$, choice is random and if ξ goes to infinity choice essentially becomes deterministic. In his comprehensive exploration of prospect theory specifications, Stott (2006) concluded that power utility, the Prelec one-parameter probability weighting function, and Luce's choice rule gave the best fit to his data. We, therefore, selected these specifications.

To test for robustness, we also ran our analysis with exponential utility, Prelec's (1998) two-parameter specification of the weighting function, and an alternative, incomplete beta (IBeta) probability weighting function (Wilcox 2012). IBeta is a flexible, two parameter family that can accommodate many shapes (convex, concave, s-shaped and inverse s-shaped) (see Appendix B and the online Appendix).

Each of the 139 subjects in the experiment had his own parameter vector $B_i = (\alpha_i, \gamma_i, \lambda_i, \xi_i)$. We assumed that each parameter in B_i comes from a lognormal distribution: $\alpha_i \sim \log N(\mu_\alpha, \sigma_\alpha^2)$, $\lambda_i \sim \log N(\mu_\lambda, \sigma_\lambda^2)$, $\gamma_i \sim \log N(\mu_\gamma, \sigma_\gamma^2)$, and $\xi_i \sim \log N(\mu_\xi, \sigma_\xi^2)$. Thus, the complete vector of unknown parameters at the population-level is $\theta_G = (\mu_\alpha, \mu_\lambda, \mu_\gamma, \mu_\xi, \sigma_\alpha^2, \sigma_\lambda^2, \sigma_\gamma^2, \sigma_\xi^2)$. For the hyper-priors, $\pi_* = (\mu_*, \sigma_*^2), * \in \{\alpha, \lambda, \gamma, \xi\}$ of the parent distributions we made the usual assumption that the μ_* follow a lognormal distribution and that the σ_*^2 follow an inverse Gamma distribution. We centered the hyper-priors at linearity (expected value) and chose the variances such that the hyper-priors were diffuse and would have a negligible impact on the posterior estimation.

Denote the set of hyper-priors of the core behavioral parameters by π_B . Then the joint probability distribution of the behavioral parameters $\mathbf{B} = (B_1, \dots, B_{139})$ and θ_B is

$$P(\mathbf{B}, \theta_B | \pi_B) = \left(\prod_{i=1}^{139} P(B_i | \theta_B) \right) P(\theta_B | \pi_B). \quad (10)$$

Given reference point rule RP_i , the likelihood of subject i 's responses is

$$P(D_i|B_i, RP_i) = \prod_{q=1}^{70} P(D_{i,q}|G_i, RP_i). \quad (11)$$

The probability of each choice $D_{i,q}$ is computed using Luce's rule, Eq.(9). From Eqs. (10) and (11), it follows that the joint probability distribution of all the unknown behavioral parameters \mathbf{B} and θ_B and all the observed choices $\mathbf{D} = (D_1, \dots, D_{139})$ is

$$P(\mathbf{D}, \mathbf{B}, \theta_B | \mathbf{RP}, \pi_B) = \left(\prod_{i=1}^{139} \prod_{q=1}^{70} P(D_{i,q}|B_i, RP_i) \right) \left(\prod_{i=1}^{139} P(B_i|\theta_B) \right) P(\theta_B|\pi_B). \quad (12)$$

In Eq.(12), $\mathbf{RP} = (RP_1, \dots, RP_{139})$ is the vector of individual reference point rules.

5.2. Specification of the reference point rule

We assume that subjects use one of the six reference point rules in Table 1. For each of these rules, estimated the posterior probability that an individual uses it given the data: $P(RP_i|\mathbf{D})$. RP_i is a (six-dimensional) categorical variable for which it is common to use the Dirichlet distribution: $\theta_{RP} \sim \text{Dirichlet}(\pi_{RP})$, where θ_{RP} is a probability vector in a six-dimensional simplex and π_{RP} is a diffuse hyper-prior parameter for the Dirichlet distribution. Then the joint probability density of \mathbf{RP} and θ_{RP} becomes:

$$P(\mathbf{RP}, \theta_{RP} | \pi_{RP}) = \left(\prod_{i=1}^{139} P(RP_i|\theta_{RP}) \right) P(\theta_{RP}|\pi_{RP}). \quad (13)$$

Substituting Eq.(13) into Eq.(12) gives the complete specification of our statistical model:

$$P(\mathbf{D}, \mathbf{B}, \theta_B, \mathbf{RP}, \theta_{RP} | \pi_B, \pi_{RP}) = \left(\prod_{i=1}^{139} \prod_{q=1}^{70} P(D_{i,q}|B_i, RP_i) \right) \left(\prod_{i=1}^{139} P(B_i|\theta_B) \right) \left(\prod_{i=1}^{139} P(RP_i|\theta_{RP}) \right) P(\theta_B|\pi_B) P(\theta_{RP}|\pi_{RP}). \quad (14)$$

5.3. Estimation

To compute the marginal posterior distributions $P(B_i|\mathbf{D}, \pi_B, \pi_{RP})$, $P(RP_i|\mathbf{B}, \pi_B, \pi_{RP})$, $P(\theta_B|\mathbf{D}, \pi_B, \pi_{RP})$, and $P(\theta_{RP}|\mathbf{D}, \pi_B, \pi_{RP})$, we used Markov Chain Monte Carlo (MCMC)

sampling (Gelfand and Smith 1990) with blocked Gibbs sampling.⁷ We first used 10,000 burn-in iterations with adaptive MCMC and then 20,000 standard MCMC burn-in iterations. The results are based on the subsequent 50,000 iterations.

6. Results

6.1. Consistency

To test for consistency, five choices were asked twice. In 68.7% of these repeated choices, subjects made the same choice. Reversal rates up to one third are common in experiments (Stott 2006). Moreover, our choices were complex, involving more than two outcomes and with expected values that were close.

6.2. Reference points

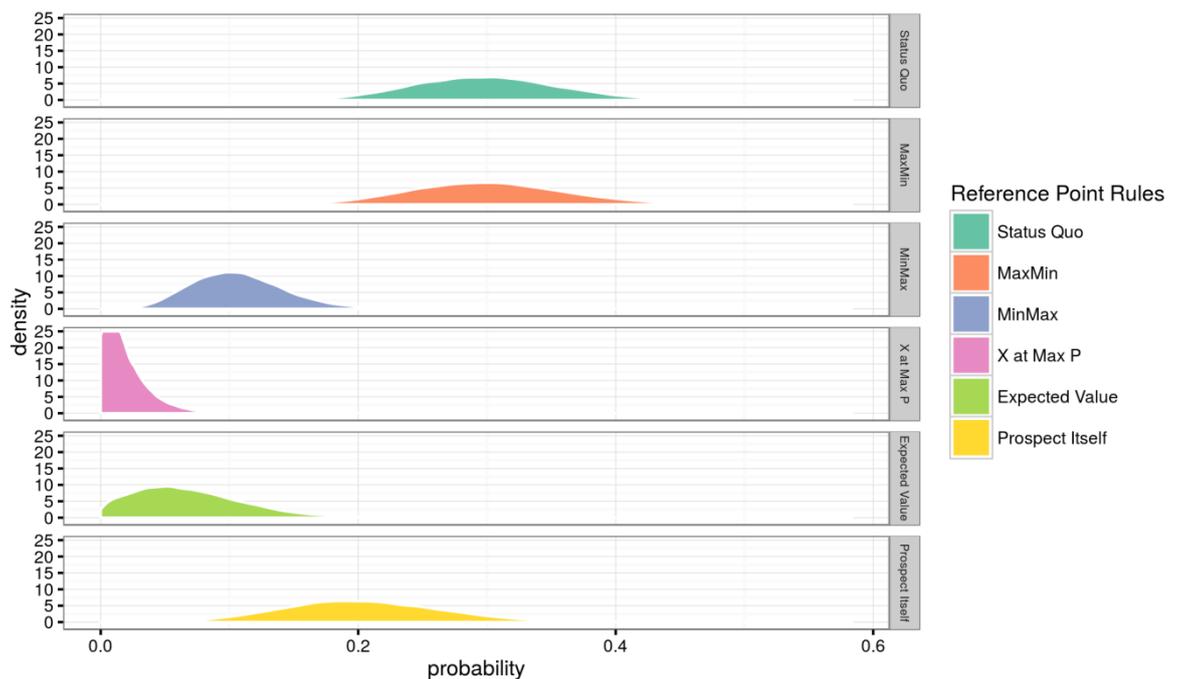


Figure 3. Marginal posterior distributions of each reference point rule

⁷ For the behavioral parameters B_1, \dots, B_{139} we used Metropolis-Hasting MCMC with symmetric normal proposal on the log-scale, for the block RP_1, \dots, RP_{139} we used Metropolis-Hasting MCMC with uniform proposal, and the group-level blocks θ_G and θ_{RP} were sampled directly from the conjugate Gamma-Normal and Dirichlet-Categorical distributions, respectively.

We first report our estimates of θ_{RP} , which indicate for each reference point rule the probability that a randomly chosen subject behaved in agreement with it. Figure 3 shows the marginal posterior distributions of θ_{RP} in the population, for each RP rule. Table 2 reports the medians and standard deviations of these distributions.⁸

	Median	SD
Status Quo	0.30	0.06
MaxMin	0.30	0.06
MinMax	0.10	0.04
X at Max P	0.01	0.02
Expected Value	0.06	0.04
Prospect Itself	0.20	0.06

Table 2. Medians and standard deviations of the marginal posterior distributions of the reference point rules in the population.

The reference points that were most likely to be used were the status quo and MaxMin. According to our median estimates, each of these two rules was used by 30% of the subjects. The prospect itself (the rule suggested by Köszegi and Rabin (2006, 2007) and Delquié and Cillo (2006)) was used by 20% of the subjects. The other three rules were used rarely.

At the individual subject level, we can also assess the likelihood that a subject uses a specific reference point by looking at that individual's posterior distributions. Figure 4 shows, for example, the posterior distributions of subjects 17, 50, and 100. Subject 17 has about 60% probability to use the prospect itself as his reference point and 25% probability to use the minimum of the maximums. Subject 50 almost surely uses MaxMin and subject 100 almost surely uses the status quo as his reference point.

⁸ Note that the medians need not add to 100%.

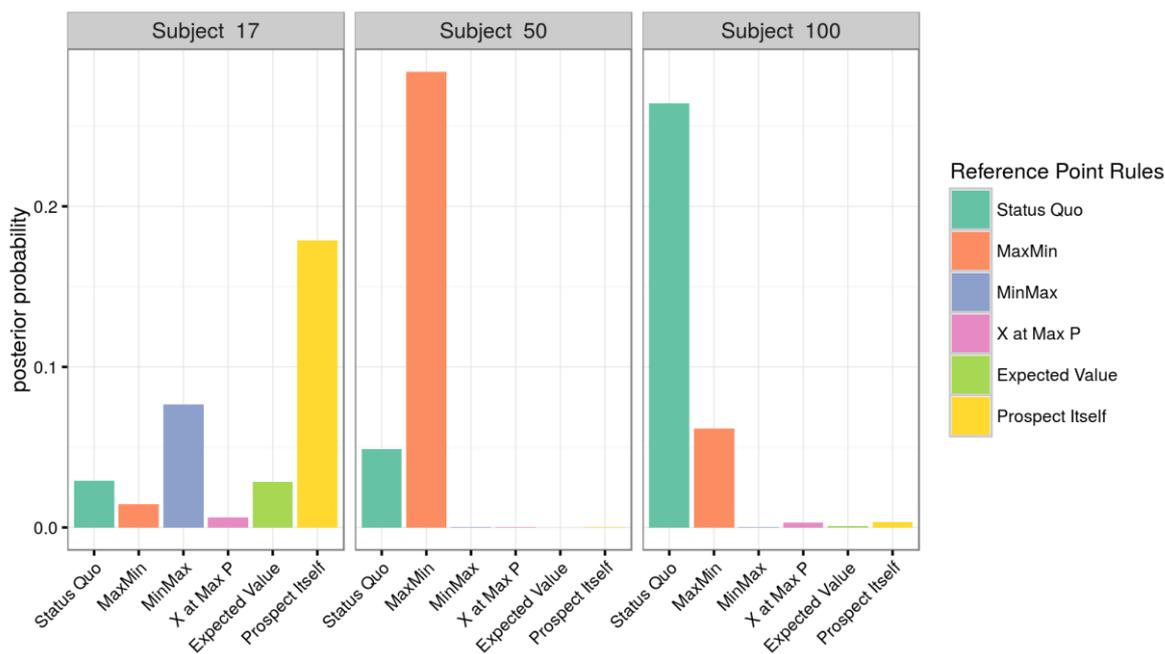


Figure 4. Posterior distributions of subjects 17, 50, and 100.

We say that a subject is classified *sharply* if he has a posterior probability of at least 50% to use one of the six reference point rules. For example, subjects 17, 50, and 100 were all classified sharply. Out of the 139 subjects, 107 could be classified sharply. Figure 5 shows the distribution of the sharply classified subjects over the six reference point rules. The dominance of the Status Quo and MaxMin increased further and around 70% of the sharply classified subjects used one of these two rules.

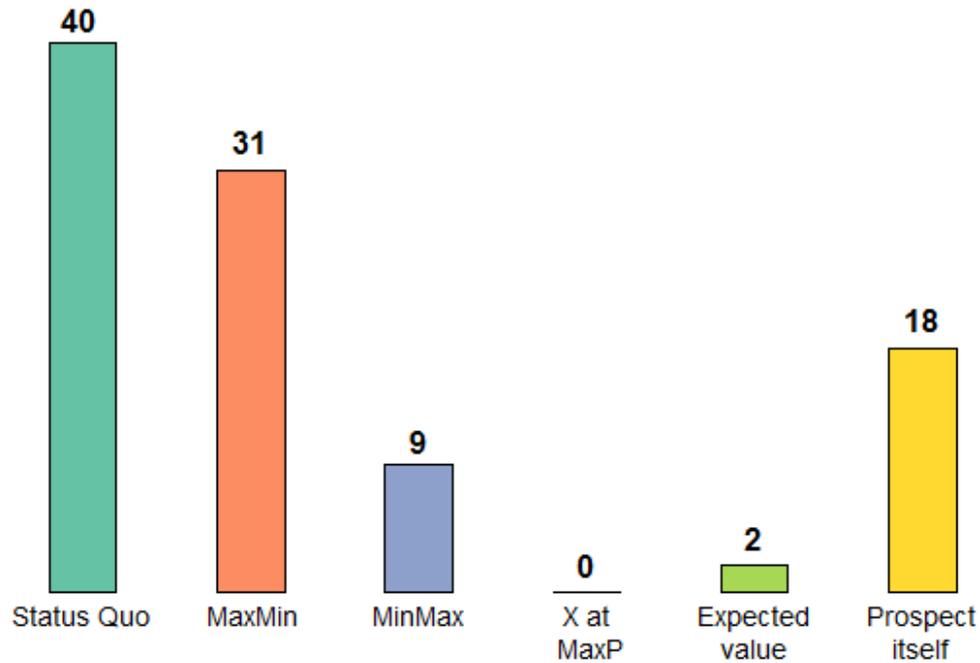


Figure 5. Proportion of sharply classified respondents satisfying a particular reference point rule (percent)

6.3. Behavioral parameters

Figure 6 shows the gain-loss utility function in the psychological (PT) part of Eq. (6) based on the estimated behavioral population level parameters (θ_B). The utility function is S-shaped: concave for gains and convex for losses. We found more utility curvature than most previous estimations of gain-loss utility (for an overview see Fox and Poldrack 2014), but our estimated utility function is close to the functions estimated by Wu and Gonzalez (1996), Gonzalez and Wu (1999), and Toubia et al. (2013). The loss aversion coefficient was equal to 2.34, which is consistent with other findings in the literature.

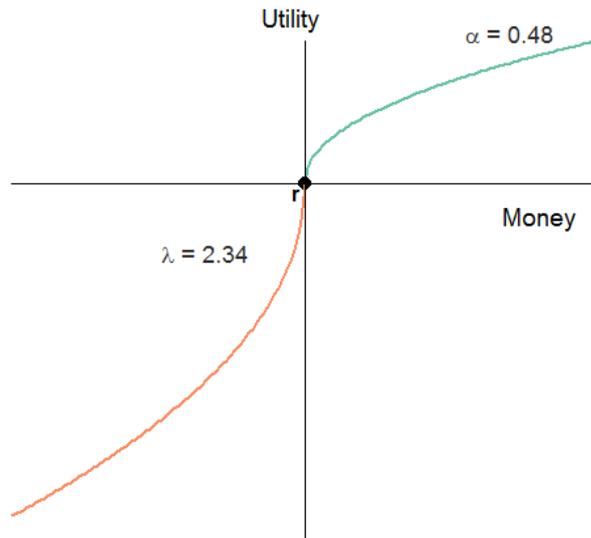


Figure 6. The gain-loss utility function based on the estimated group parameters

Figure 7 shows the estimated probability weighting function in the population. The function has the commonly observed inverse S-shape, which reflects overweighting of small probabilities and underweighting of intermediate and large probabilities.⁹ Our estimated probability weighting function is close to the functions that were estimated by Gonzalez and Wu (1999), Bleichrodt and Pinto (2000) and Toubia et al. (2013).

⁹The Prelec one-parameter probability weighting function only allows for inverse- or S-shaped weighting. However, the two-parameter Prelec function and the IBeta function allow for all shapes and their estimated shapes were also inverse S.

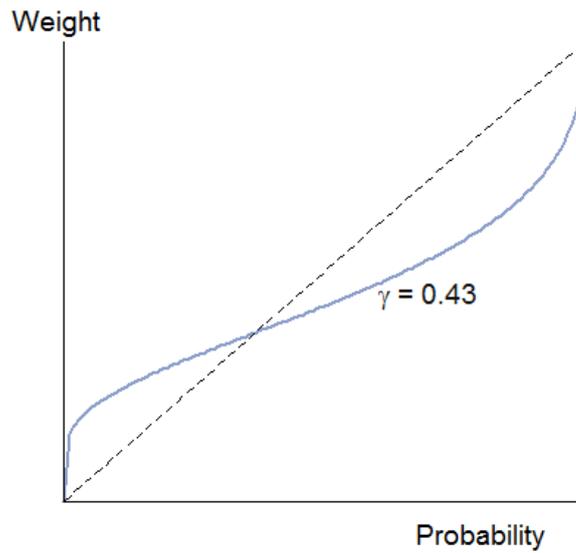


Figure 7. The probability weighting function based on the estimated group parameters

Bayesian hierarchical modeling permits expressing uncertainty in the individual parameter estimates by means of the posterior densities. To illustrate, Figure 8 shows the posterior densities of subject 17. As the graph shows, subject 17's parameter estimates varied considerably, although it is safe to say that he had concave utility and inverse S-shaped probability weighting.

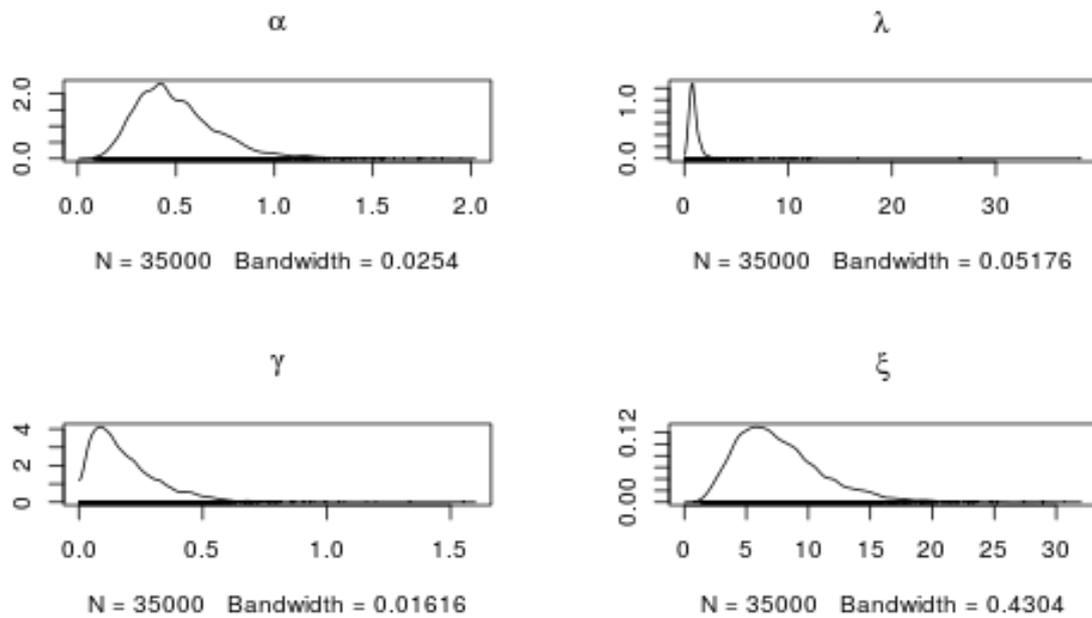


Figure 8. Posterior densities of the behavioral parameters for subject 17.

Table 3 shows the quantiles of the posterior point estimates of all 139 subjects. The table shows that utility curvature and, to a lesser extent, probability weighting were rather stable. Loss aversion varied much more although the estimates of more than 75% of the subjects were consistent with loss aversion.

	2.5%	25%	50%	75%	97.5%
α	.31	.40	.44	.50	.60
γ	.09	.14	.24	.44	1.66
λ	.36	1.19	1.59	2.25	4.63
ξ	6.11	8.26	10.89	14.41	25.76

Table 3. Quantiles of the point estimates of the behavioral parameters of the 139 subjects

Table 4 shows the median behavioral parameters of the sharply classified subjects in each group.¹⁰ A priori, it seemed plausible that subjects who used different rules might

¹⁰ X at MaxP is not in the table as there were no sharply classified subjects who behaved according to this rule.

also have different behavioral parameters, in particular loss aversion, and the table confirms this conjecture. While utility curvature and probability weighting were rather stable across the groups, the loss aversion coefficients varied from 0.50 in the MinMax group to 2.44 in the Expectation group. The loss aversion coefficient of 0.50 in the MinMax group has the interesting interpretation that these optimistic subjects weight gains twice as much as losses and they exhibit what might be seen as the reflection of the preferences of the cautious MaxMin subjects who weight losses more than twice as much as gains.

	α	γ	λ	ξ
Status Quo	.42	.28	1.51 ¹¹	11.75
MaxMin	.46	.24	2.24	10.30
MinMax	.40	.15	.50	14.34
Expectation	.36	.25	2.44	6.14
Prospect Itself	.45	.16	2.23	10.89

Table 4: Median individual level parameters for the sharply classified subjects in each group.

Table 4 also shows that subjects who used the status quo as their reference point were typically no expected utility maximizers as there was substantial probability weighting in this group. Table 5 gives more detailed evidence. It shows the subdivision of the subjects who used the status quo as their reference point based on the 95% Bayesian credible intervals of their estimated utility curvature and probability weighting parameters. Twelve subjects (those with $\gamma = 1$) behaved according to expected utility, three of whom (those with $\alpha = 1$ and $\gamma = 1$) were expected value maximizers. Thus, less than 10% of our subjects were expected utility maximizers.

¹¹ The reason that λ is not equal to 1 for subjects who were sharply classified as using the status quo rule is that a subject's behavioral parameters stayed the same for all reference point rules. Consequently, even when a subject was (sharply) classified as a status quo type, there was still a non-negligible probability that he used any of the other reference point rules.

		Probability weighting			Total
		$\gamma < 1$	$\gamma = 1$	$\gamma > 1$	
Utility	$\alpha < 1$	28	9	0	37
	$\alpha = 1$	3	3	0	9
	$\alpha > 1$	0	0	0	0
	Total	31	12	0	43

Table 5. Behavioral parameters of the subjects using the status quo as their reference points (classification into groups is based on the 95% Bayesian credible intervals).

6.4. Robustness

Throughout the main analysis, we used power utility (Eq. (7)) and Prelec’s (1998) one-parameter probability weighting function (Eq. (8)). We performed three robustness checks, replacing power utility by exponential utility, Prelec’s (1998) one-parameter weighting function by his two-parameter function, and finally using another much more flexible IBeta weighting function.¹² In all three cases, the status quo and MaxMin remained the most commonly-used reference points. They always captured the behavior of more than 60% of the subjects. Details are in the online appendix.

In the main analysis, we assumed Eq. (6) for all reference point rules, allowing us to keep all behavioral parameters constant when comparing reference point rules. We also tried several other specifications, which are summarized in Table 6. Model 1 corresponds to the results reported in Sections 6.2 and 6.3. The two variables we varied in the robustness checks were the inclusion of consumption utility and probability weighting. While models with prospect-specific reference points need consumption utility to exclude implausible choice behavior,¹³ models with a choice-specific reference point do not. Prospect theory, for example, does not include consumption utility.

¹² We used the incomplete Beta function because it allows for a wide variety of shapes (Wilcox 2012). See Appendix B.

¹³ For example, in Köszegi and Rabin’s (2007) CPE model without consumption utility any prospect that gives x with probability 1 has a value of 0, regardless of the size of x . So the decision maker should be indifferent between \$1 for sure and \$1000 for sure. Consumption utility prevents this.

Consequently, we estimated the models with a choice-specific reference point both with and without consumption utility.

Model	Choice-specific reference point		Prospect-specific reference point	
	Consumption utility	Probability weighting	Consumption utility	Probability weighting
1	Yes	Yes	Yes	Yes
2	No	Yes	Yes	Yes
3	Yes	Yes	Yes	No
4	No	Yes	Yes	No
5	Yes	No	Yes	No
6	No	No	Yes	No

Table 6: Estimated models

In Eq. (6) we assumed that subjects weight probabilities when they evaluate prospects relative to a reference point, but, following the literature on stochastic reference points, we abstracted from probability weighting in the determination of the stochastic reference point. This may be arbitrary and we, therefore also estimated the models without probability weighting. We performed two sets of estimations: one in which the models with a choice-specific reference point included probability weighting, but the models with a prospect-specific reference point did not (models 3 and 4) and one in which no model had probability weighting (models 5 and 6).

The results of the robustness checks were as follows. First, our main conclusion that the status quo and MaxMin were the dominant reference points remained valid. The behavior of 60% to 75% of the subjects was best described by a model with one of these two reference points. Second, excluding consumption utility from models with a choice-specific reference point (models 2 and 4) led to a substantial increase in the precision parameter ξ . This suggests that there is no need to include consumption utility in models like prospect theory. Third, probability weighting played a crucial role.

Excluding probability weighting from the models with a prospect-specific reference point (model 3) decreased the share of the *prospect itself* as a reference point to 10% (8% if we only include the sharply classified subjects) and increased the share of the MaxMin reference point to 44% (52% if we only include the sharply classified subjects). The other shares changed only little. Hence, prospect-specific models like Köszegi and Rabin's (2006, 2007) benefit from including probability weighting. Ignoring probability weighting altogether, as in models 5 and 6, led to unstable estimation results.

The behavioral parameters were comparable across all models that we estimated. The power utility coefficient was approximately 0.50 in all models, the probability weighting parameter varied between 0.40 and 0.60 (except, of course, when no probability weighting was assumed), and the loss aversion coefficient varied between 2 and 2.50. Full results of the robustness analysis are in the online appendix.

7. Discussion

Experiments in decision under risk often assume that subjects take the status quo (0) as their reference point. Our data show that this assumption is justified for 30-40% of the subjects, but that a majority uses a different reference point. Our data also suggest how researchers can increase the likelihood that subjects use 0 (or the participation fee) as their reference point. For example, in choosing between mixed prospects, researchers could include a prospect with 0 as its minimum outcome in each choice. This ensures that MaxMin subjects will also use 0 as their reference point and, as our results suggest, that a substantial majority of the subjects will use 0 as their reference point. Our results can also help to assess the validity of empirical studies that take the status quo as the reference point.

We tested the reference point rules in an experiment with large incentives (subjects could win up to a weekly salary) and used a Bayesian hierarchical approach to analyze the data. Economic researchers usually estimate models either by pooling all data or by individual estimation. Both approaches have their limitations. Pooling ignores individual heterogeneity and may result in estimates that are not representative of any individual in the sample. Individual estimation relies on relatively few data points, which may lead to unreliable results. Bayesian analysis strikes a nice balance between pooling and individual estimation and it leads to more precise parameter estimates. A potential limitation of Bayesian analysis is that the selected priors can affect the estimations, but in our analysis the choice of priors had negligible impact on the estimates.

To make inferences about the different reference point rules, we used a comprehensive model which allowed isolating the impact of the reference point rule from the other behavioral parameters (*ceteris paribus* principle). This approach is cleaner and better interpretable than the common practice in mixture modeling where each model in the mixture is specified separately and parameterizations can differ across models and than horse races between models based on criteria like the Akaike Information Criterion. Using a Bayesian model has the additional advantage that we could obtain the parameter estimates for both the distribution of reference point rules in the population and each subject separately.

Our robustness tests have two interesting implications for the modeling of reference-dependent preferences. First, they indicate that models with a choice-specific reference point do not benefit from including consumption utility. Kahneman and Tversky (1979, p.277) point out that even though an individual's attitudes to money depend both on his asset position and on changes from his reference point, a utility

function that is only defined over changes from the reference point generally provides a satisfactory approximation. Our results provide support for their argument.

Second, we concluded that probability weighting played an important role and could not be ignored. Expectation-based prospect-specific models like Köszegi and Rabin's (2006, 2007) model, which originally don't assume probability weighting, clearly benefited from including probability weighting. Further research on the role of probability weighting in such models, particularly how it can be incorporated in the determination of the reference point, is needed.

We did not test all reference points that have been proposed in the literature. As we explained in the introduction, we followed Rabin's (2013) approach and studied reference point rules that used the same independent variables as the core economic theory of decision under risk, expected utility. This implied, for example, that we did not test explicitly for subjects' goals. On the other hand, subjects may have had few goals for the current experiment and it is also possible that their goals were equal to one of the reference points that we used (e.g. expected value or the security level). A similar remark applies to aspiration levels. Testing for goals and aspiration levels seems difficult within Rabin's approach and may require other data inputs.

8. Conclusion

Reference-dependence is a key concept in explaining people's choices. A fundamental problem of reference-dependent theories is that they are silent about how reference points are formed. This paper has tried to break this silence. We have estimated the prevalence of six reference point rules using a unique data set in which we could use stakes up to a weekly salary. We modeled the reference point rule as a latent categorical variable, which we estimated using Bayesian hierarchical modeling. Our results indicate

that the status quo and MaxMin are the most commonly used reference points. Around twenty percent of the subjects used the prospect itself, as in Köszegi and Rabin's (2006, 2007) CPE model.

Appendix A: Derivation that Abeler et al.'s (2011) results are consistent with the MaxMin rule.

In the experiment of Abeler et al. (2011), the subject could perform a task many times and received either a fixed payment f with probability 50% or we with probability 50% where e is the effort (the number of tasks performed) and $w > 0$ the piece rate per task. There were two treatments defined by the value of f : the low treatment (LO) and the high treatment (HI), with $f_{LO} < f_{HI}$. In a model without a reference point or if the reference point is the status quo, the subjects should perform the same number of tasks in both treatments. Abeler et al. (2011) derived two hypotheses when the prospect itself is the reference point (the rule suggested by Köszegi and Rabin's (2007) CPE model):

- Hypothesis 1: Average effort in the HI treatment is higher than in the LO treatment.
- Hypothesis 2: The probability to stop at $we = f_{LO}$ is higher in the LO treatment than in the HI treatment; the probability to stop at $we = f_{HI}$ is higher in the HI treatment than in the LO treatment.

The results of their experiment were in line with these hypotheses. We will now show that we can derive the same hypotheses with the MaxMin rule. First, subjects' choice sets consisted of binary prospects $(50\%, f; 50\%, we)$ for all e . The minimum outcome is we if $we < f$ and f otherwise. Hence the maximum of the minimum outcomes is f and the MaxMin rule predicts f to be the reference point. We adopt now the same model as Abeler et al. (2011) but with the MaxMin rule. They consider a piecewise linear gain-loss utility with $U(x) = \eta x$ for gains and $U(x) = \eta \lambda x$ for losses with $\eta \geq 0$ and $\lambda > 1$. The cost of performing effort e is $c(e)$.

The subjects will maximize if $we < f$:

$$\frac{we + f}{2} - c(e) + \frac{1}{2}\eta\lambda(we - f)$$

Or, if $we \geq f$:

$$\frac{we + f}{2} - c(e) + \frac{1}{2}\eta(we - f)$$

The first-order conditions (FOCs) give:

if $we < f$:

$$c'(e^*) = \frac{1}{2}w(1 + \eta\lambda)$$

if $we \geq f$:

$$c'(e^*) = \frac{1}{2}w(1 + \eta)$$

Hence, the marginal utility of providing an effort (the second member of the FOC equations) is higher when accumulated earnings we are below the reference point f than when they exceed it, because of the loss aversion coefficient ($\lambda > 1$). Increasing the fixed payment from f_{LO} to f_{HI} increases the marginal utility of providing an effort when the accumulated earnings are within this range. Hence, Hypothesis 1 follows.

To show that MaxMin also predicts the second hypothesis, there is a discrete drop of the marginal utility of efforts when $we = f$, which will cause some subjects to stop their effort at precisely this level, in line with Hypothesis 2.

Appendix B: IBeta

The *incomplete regularized beta function* (*IBeta*), is a very flexible monotonically increasing $[0,1] \rightarrow [0,1]$ function. It can capture a wide range of convex, concave, S-shape and inverse S-shape functions without favoring specific shapes or inflection points. The family is symmetric in the sense that $IBeta(x; a, b) = 1 - IBeta(1 - x; a, b)$. Various shapes of *IBeta* function are illustrated in Figure B.1.

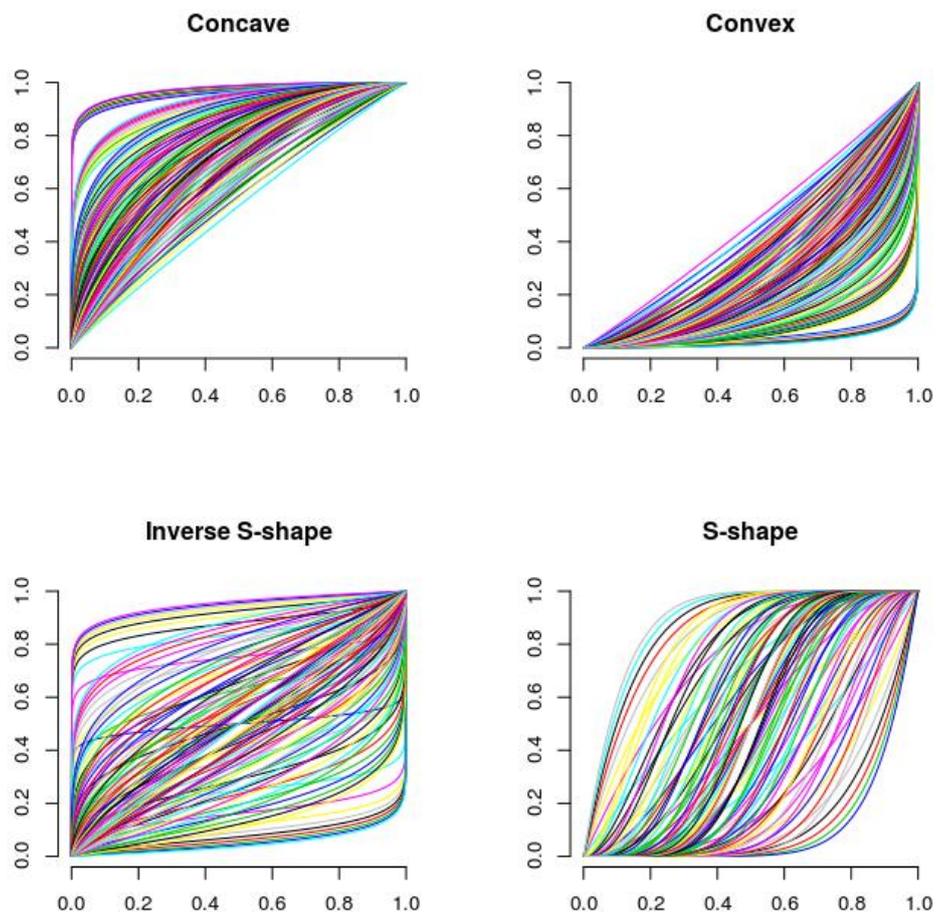


Figure B.1. Various shapes of the *IBeta* function

References

- Abdellaoui M (2000) Parameter-free elicitation of utility and probability weighting functions. *Management Science* 46: 1497-1512.
- Abeler J, Falk, A, Goette, L, Huffman, D (2011) Reference points and effort provision. *American Economic Review* 101: 470-492.
- Allen EJ, Dechow, PM, Pope, DG, Wu, G (forthcoming) Reference-dependent preferences: Evidence from marathon runners. *Management Science* .
- Barber BM, Odean, T (2008) All that glitters: The effect of attention and news on the buying behavior of individual and institutional investors. *Review of Financial Studies* 21: 785-818.
- Barberis NC (2013) Thirty years of prospect theory in economics: A review and assessment. *Journal of Economic Perspectives* 27: 173-196.
- Bartling B, Brandes, L, Schunk, D (2015) Expectations as reference points: Field evidence from professional soccer. *Management Science* 61: 2646-2661.
- Baucells M, Weber, M, Welfens, F (2011) Reference-point formation and updating. *Management Science* 57: 506-519.
- Bell DE (1985) Disappointment in decision making under uncertainty. *Operations Research* 33: 1-27.
- Benartzi S, Previtro, A, Thaler, RH (2011) Annuitization puzzles. *Journal of Economic Perspectives* 25: 143-164.
- Benartzi S, Thaler, RH (1995) Myopic loss aversion and the equity premium puzzle. *Quarterly Journal of Economics* 110: 73-92.
- Birnbaum MH, Schmidt, U (2010) Testing transitivity in choice under risk. *Theory and Decision* 69: 599-614.
- Bleichrodt H, Schmidt, U (2005) Context- and reference-dependent utility: A generalization of prospect theory. Working Paper.
- 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: 1498-1514.
- Bordalo P, Gennaioli, N, Shleifer, A (2012) Salience theory of choice under risk. *Quarterly Journal of Economics* 127: 1243-1285.

- Card D, Dahl, GB (2011) Family violence and football: The effect of unexpected emotional cues on violent behavior. *Quarterly Journal of Economics* 126: 103-143.
- Chetty R, Looney, A, Kroft, K (2009) Salience and taxation: Theory and evidence. *American Economic Review* 99: 1145-1177.
- Crawford VP, Meng, J (2011) New York city cab drivers' labor supply revisited: Reference-dependent preferences with rational expectations targets for hours and income. *American Economic Review* 101: 1912-1932.
- Delqu   P, Cillo, A (2006) Disappointment without prior expectation: A unifying perspective on decision under risk. *Journal of Risk and Uncertainty* 33: 197-215.
- Diecidue E, Van de Ven, J (2008) Aspiration level, probability of success and failure, and expected utility. *International Economic Review* 49: 683-700.
- Eil D, Lien, JW (2014) Staying ahead and getting even: Risk attitudes of experienced poker players. *Games and Economic Behavior* 87: 50-69.
- Fox CR, Poldrack, RA (2014) Prospect theory and the brain. Glimcher Paul, Fehr Ernst, eds. *Handbook of Neuroeconomics (2nd Ed.)* (Elsevier, New York), 533-567.
- Gelfand AE, Smith, AF (1990) Sampling-based approaches to calculating marginal densities. *Journal of American Statistical Association* 85: 398-409.
- Genesove D, Mayer, CJ (2001) Loss aversion and seller behavior: Evidence from the housing market. *Quarterly Journal of Economics* 116: 1233-1260.
- Gill D, Prowse, V (2012) A structural analysis of disappointment aversion in a real effort competition. *American Economic Review* 102: 469-503.
- Gonzalez R, Wu, G (1999) On the form of the probability weighting function. *Cognitive Psychology* 38: 129-166.
- Gul F (1991) A theory of disappointment aversion. *Econometrica* 59: 667-686.
- Heath C, Larrick, RP, Wu, G (1999) Goals as reference points. *Cognitive Psychology* 38: 79-109.
- Heidhues P, K  szegi, B (2008) Competition and price variation when consumers are loss averse. *American Economic Review* 1245-1268.
- Hershey JC, Schoemaker, PJH (1985) Probability versus certainty equivalence methods in utility measurement: Are they equivalent? *Management Science* 31: 1213-1231.
- Johnson EJ, Goldstein, D (2003) Do defaults save lives? *Science* 302: 1338-1339.
- Kahneman D, Tversky, A (1979) Prospect theory: An analysis of decision under risk. *Econometrica* 47: 263-291.

- Kőszegi B, Rabin, M (2007) Reference-dependent risk attitudes. *American Economic Review* 97: 1047-1073.
- Kőszegi B, Rabin, M (2006) A model of reference-dependent preferences. *Quarterly Journal of Economics* 121: 1133-1166.
- Lien JW, Zheng, J (2015) Deciding when to quit: Reference-dependence over slot machine outcomes. *American Economic Review: Papers & Proceedings* 105: 366-370.
- Loomes G, Sugden, R (1986) Disappointment and dynamic consistency in choice under uncertainty. *Review of Economic Studies* 53: 271-282.
- Luce RD (1959) On the possible psychophysical laws. *Psychological Review* 66: 81.
- Masatlioglu Y, Raymond, C (forthcoming) A behavioral analysis of stochastic reference dependence. *American Economic Review*.
- Nilsson H, Rieskamp, J, Wagenmakers, E (2011) Hierarchical Bayesian parameter estimation for cumulative prospect theory. *Journal of Mathematical Psychology* 55: 84-93.
- Odean T (1998) Are investors reluctant to realize their losses? *Journal of Finance* 53: 1775-1798.
- Pope DG, Schweitzer, ME (2011) Is Tiger Woods loss averse? persistent bias in the face of experience, competition, and high stakes. *American Economic Review* 101: 129-157.
- Prelec D (1998) The probability weighting function. *Econometrica* 66: 497-528.
- Rabin M (2013) An approach to incorporating psychology into economics. *American Economic Review* 103: 617-622.
- Rabin M (2000) Risk aversion and expected-utility theory: A calibration theorem. *Econometrica* 68: 1281-1292.
- Rouder JN, Lu, J (2005) An introduction to Bayesian hierarchical models with an application in the theory of signal detection. *Psychonomic Bulletin & Review* 12: 573-604.
- Samuelson W, Zeckhauser, R (1988) Status quo bias in decision making. *Journal of Risk and Uncertainty* 1: 7-59.
- Schmidt U, Starmer, C, Sugden, R (2008) Third-generation prospect theory. *Journal of Risk and Uncertainty* 36: 203-223.
- Schneider M, Day, R (forthcoming) Target-adjusted utility functions and expected-utility paradoxes. *Management Science*.
- Stott HP (2006) Cumulative prospect theory's functional menagerie. *Journal of Risk and Uncertainty* 32: 101-130.

- Sugden R (2003) Reference-dependent subjective expected utility. *Journal of Economic Theory* 111: 172-191.
- Sydnor J (2010) (Over) insuring modest risks. *American Economic Journal: Applied Economics* 2: 177-199.
- Thaler RH, Benartzi, S (2004) Save more tomorrow™: Using behavioral economics to increase employee saving. *Journal of Political Economy* 112: S164-S187.
- Toubia O, Johnson, E, Evgeniou, T, Delquié, P (2013) Dynamic experiments for estimating preferences: An adaptive method of eliciting time and risk parameters. *Management Science* 59: 613-640.
- Tversky A, Kahneman, D (1992) Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty* 5: 297-323.
- Van Osch S, Stiggelbout, AM (2008) The construction of standard gamble utilities. *Health Economics* 17: 31-40.
- van Osch S, van den Hout, WB, Stiggelbout, AM (2006) Exploring the reference point in prospect theory: Gambles for length of life. *Medical Decision Making* 26: 338-346.
- van Osch SMC, Wakker, PP, van den Hout, WB, Stiggelbout, AM (2004) Correcting biases in standard gamble and time tradeoff utilities. *Medical Decision Making* 24: 511-517.
- Wilcox, NT (2012) Plenary lecture FUR Atlanta.
- Wu G, Gonzalez, R (1996) Curvature of the probability weighting function. *Management Science* 42: 1676-1690.