

Third-generation prospect theory

Ulrich Schmidt

Department of Economics, Christian-Albrechts-Universität zu Kiel, D-24098 Kiel, and Kiel Institute for the World Economy, 24105 Kiel, Germany.

Chris Starmer

School of Economics, University of Nottingham, Nottingham, NG7 2RD, UK.

Robert Sugden

School of Economics, University of East Anglia, Norwich, NR4 7TJ, UK.

Abstract

We present a new theory of decision under uncertainty: third-generation prospect theory (PT³). This retains the predictive power of previous versions of prospect theory, but extends that theory by allowing reference points to be uncertain while decision weights are specified in a rank-dependent way. We show that PT³ preferences respect a state-conditional form of stochastic dominance. The theory predicts the observed tendency for willingness-to-accept valuations of lotteries to be greater than willingness-to-pay valuations. When PT³ is made operational by using simple functional forms with parameter values derived from existing experimental evidence, it predicts observed patterns of the preference reversal phenomenon.

Acknowledgements

The authors thank Serge Blondel, Louis Lévy-Garboua, William Neilson, Peter Wakker, Horst Zank and anonymous referees for helpful comments on earlier drafts. We have also benefitted from discussions with participants at various conferences and seminars where we have presented this work. Sugden's work was supported by the Economic and Social Research Council (award no. RES 051 27 0146).

Keywords: Prospect theory, preference reversal, reference dependence

JEL classifications: D81 (Criteria for decision-making under risk and uncertainty)

In this paper we present a new theory of decision under uncertainty: *third-generation prospect theory* (PT³ for short). The motivation for the theory is empirical: our model is presented as a descriptive theory intended to outperform the current ‘best buys’ in the literature. PT³ has three key features: *reference dependence*, *decision weights* and *uncertain reference points* (i.e. reference points that can be lotteries). The first two features are the common characteristics of different versions of prospect theory, including the original (or first-generation) version (Kahneman and Tversky, 1979) and the later *cumulative* (or second-generation) versions featuring rank-dependent decision weights (e.g. Starmer and Sugden, 1989; Luce and Fishburn 1991; Tversky and Kahneman, 1992; Wakker and Tversky 1993). Variants of cumulative prospect theory are increasingly widely applied in both theoretical and empirical work (recent examples are Davies and Satchell 2004; Trepel, Fox and Poldrack, 2005; Wu, Zhang and Abdellaoui, 2005; Baucells and Heukamp, 2006; Schmidt and Zank 2008) and some have argued that such theories may be serious contenders for replacing expected utility theory at least for specific purposes (see Camerer, 1989). No doubt this is partly because there is considerable empirical support for both reference-dependence and decision weights (see Starmer, 2000).

First- and second-generation prospect theory have a common limitation: the reference points from which prospects are evaluated are assumed to be certainties. If reference points are interpreted as endowments or status quo positions, these theories cannot be applied to problems in which a decision maker is endowed with a lottery and has the opportunity to sell or exchange it. Such problems are common in real economic life – for example, buying insurance or selling stocks. They also feature in many experimental designs and, in consequence, a good deal is known about how in fact decision-makers respond to them. This evidence shows two particularly well-established and robust patterns of deviation from the predictions of conventional expected utility theory.

The first is that willingness-to-accept (WTA) valuations of lotteries tend to be greater than willingness-to-pay (WTP) valuations (e.g. Knetsch and Sinden, 1984; Loomes, Starmer and Sugden, 2003). Intuitively, one might expect that loss aversion, as modelled in prospect theory, would explain this effect; but because a WTA valuation of a lottery is made from a reference point at which the decision-maker is endowed with that lottery, that intuition cannot be expressed in existing versions of prospect theory.

The second deviation is *preference reversal* (PR). The classic instances of PR involve decisions relating to pairs of gambles. In the simplest cases, gambles are binary lotteries with

just one positive outcome (the *prize*); the other outcome is zero. One of the lotteries, usually called the ‘P bet’, gives the better chance of winning a prize while the other, the ‘\$ bet’, has the larger prize. In a typical experiment, agents’ preference orderings over pairs of such bets are elicited in two ways: in a pairwise choice task, and by comparing WTA valuations of lotteries elicited separately for P and \$ bets. There is a widely observed tendency for agents to reveal a preference for the P bet in choice but the \$ bet in valuation. We will call this pattern *standard PR*. Such inconsistencies between choice and valuation might arise through chance or error. But the opposite inconsistency, in which the \$ bet is chosen but the P bet is given a higher value (*non-standard PR*), is much less frequently observed. It is this asymmetry between the two types of reversal which constitutes the puzzle of PR. Because existing versions of prospect theory cannot deal with WTA valuations of lotteries, they cannot be applied to classic PR problems.

We have set ourselves the task of generalising cumulative prospect theory so that it can encompass uncertain reference points. Thus, PT³ inherits all the descriptive success (and, of course, any descriptive failure¹) of cumulative prospect theory. As we shall show, it can also explain both PR and the WTA/WTP disparity for lotteries.

In all versions of prospect theory, the subjective value of a given monetary lottery, viewed from a given certain reference point, is a weighted average of the subjective values of the monetary gains or losses associated with the lottery outcomes. Gains and losses are defined relative to the reference point; an increasing *value function* (analogous with a utility function in expected utility theory) is used to transform these into positive or negative indices of subjective ‘value’. These indices are then aggregated by using *decision weights*, which are determined by using a *probability weighting function* (a strictly increasing mapping from the interval [0, 1] onto itself) which assigns to each probability p a *transformed* probability $w(p)$. In cumulative prospect theory, this function is used in a rank-dependent way, which we describe in more detail in Section 2. A zero gain is defined to have zero value, and so has no impact on the overall value of the lottery. Decision weights for gains are assigned cumulatively, beginning with the largest gain and working downwards. A mirror-image method is used to assign weights to losses. As an illustration of this operation applied to gains, consider a lottery which gives monetary gains of zero, £5 and £10 with respective probabilities 0.5, 0.3 and 0.2. The weight for the largest positive gain, £10, is determined by a direct transformation of the relevant probability, giving the weight $w(0.2)$. The decision weight for the £5 gain is then defined as $w(0.2 + 0.3) - w(0.2)$. Notice that the sum of the decision weights for the £5 and £10 gains is $w(0.5)$, i.e. the transformation of the sum of the

corresponding probabilities. This construction ensures that, for any two lotteries A, B, if A stochastically dominates B, then A is preferred to B, viewed from any certain reference point.

A generalisation of this theory to uncertain reference points requires two components. The first is a definition of ‘gain’ and ‘loss’ relative to such reference points. Our approach, in the tradition of Savage (1954), is to define preferences over *acts*, where an act is an assignment of consequences to *states of the world*. A stochastic reference point is modelled as a particular act, the *reference act*. Gains and losses are defined separately for each state of the world. For example, consider a lottery with tickets numbered 1, ..., 100; one ticket will be drawn at random and its owner will win £100. Consider an agent who is endowed with ticket 1 and who treats this as her reference act. She is offered the opportunity to exchange ticket 1 for ticket 2. If she consents to this exchange, there is a 0.01 probability that ticket 1 will be drawn, in which case she will be £100 worse off than if she had kept her initial endowment. There is also a 0.01 probability that ticket 2 will be drawn, in which case she will be £100 better off. If any other ticket is drawn, she is neither better off nor worse off. Thus, the option of taking ticket 2 in exchange for ticket 1, viewed from the reference act, gives £100 gain with probability 0.01, £100 loss with probability 0.01, and no change in wealth with probability 0.98.

The second component for a generalisation is a method for determining decision weights when reference points are uncertain. Given our Savage-style framework, we need a rank-dependent method of assigning decision weights to any act f , viewed relative to any reference act h . For any state of the world s , we use $f(s)$ and $h(s)$ to denote the outcomes of f and h in that state. Our approach is to rank states in terms of the ex post net gain from choosing f rather than h – that is, by the value of $f(s) - h(s)$; separate rankings are constructed for ‘gain’ states (for which $f(s) - h(s)$ is positive) and ‘loss’ states (for which it is negative). These rankings are then used to determine decision weights, just as rankings of outcomes are used in second-generation prospect theory. This approach ensures that preferences respect a state-conditional form of stochastic dominance, even when the reference point is uncertain.

By using these two components, any parameterised form of second-generation prospect theory can be generalised to the case of uncertain reference points. No additional parameters are required. This is significant, because the empirical literature provides a lot of information about the parameterisations of prospect theory that are most successful in organising experimental data, and about the values that the various parameters typically take. Thus, we can investigate whether PT³ explains observations of behaviour in decision

problems with uncertain reference points, *using parameterisations whose empirical validity has already been established.*

We show that PT³ predicts the observed WTA/WTP disparity for lotteries, given only the standard assumption of loss aversion in the value function. Perhaps more surprisingly, we show that, when PT³ is configured with parameterisations typical of those already established in the empirical literature, it predicts standard patterns of PR. Consequently we suggest that PT³ is a best buy theory: it offers the predictive power of previous variants of prospect theory and adds to that an explanation of PR. The latter comes ‘free of charge’ since it involves no extra parameters and no re-parameterisation.

2. PT³ in its Most General Form

In this section we introduce PT³. In this theory, preferences are defined over Savage acts. Consider a finite state space S , consisting of the states s_i , $i = 1, \dots, n$, and a set of consequences X given by an interval of the real line. In interpreting the theory, we treat consequences as levels of wealth. As in the case of second-generation prospect theory, PT³ can be formulated either for risk or for uncertainty. In this paper, we restrict attention to risk. The extension to uncertainty is straightforward and is exactly as in second generation prospect theory. For decision making under risk each state s_i has an objective probability $\pi_i \geq 0$, with $\sum_i \pi_i = 1$. F is the set of all acts. A particular act $f \in F$ is a function from S to X , i.e. an act f specifies for each state s_i the resulting consequence $f(s_i) \in X$.

As in other versions of prospect theory, preferences over acts are reference-dependent. We formalise this following the approach of reference-dependent subjective expected utility theory (RDSEU), as proposed by Sugden (2003). For any three acts f, g, h , $f \succsim_h g$ denotes that f is weakly preferred to g viewed from h , the *reference act*. (The corresponding relations of strict preference and indifference are written as $f \succ_h g$ and $f \sim_h g$.) For present purposes the reference act can be interpreted as the status quo position. We depart from earlier versions of prospect theory by not requiring h to be a constant act (i.e. an act which gives the same consequence in every state). Instead, we adopt a key innovation of RDSEU.

Sugden’s axiom system implies maximisation of the function:

$$(1) \quad V(f, h) = \sum_i v(f[s_i], h[s_i])\pi_i$$

In this expression, $v(f[s_i], h[s_i])$ is a *relative value* function. It can be interpreted as the desirability of the consequence of act f in state s_i relative to the consequence of a reference act h in the same state. This function is strictly increasing in its first argument, with $v(f[s_i], h[s_i]) = 0$ when $f(s_i) = h(s_i)$. The function $V(f, h)$ is the expectation of relative value. It assigns a real value to any act $f \in F$ viewed from any reference act $h \in F$ (i.e. $V: F \times F \rightarrow \mathbf{R}$). It is a preference representation in the sense that, for all f, h, g in F , $f \succcurlyeq_h g \Leftrightarrow V(f, h) \geq V(g, h)$.

In opting for the RDSEU approach, we have made a significant modelling decision. This approach is based on a *state-contingent* conception of reference-dependence. That is, gain/loss comparisons are made separately for each state of the world; thus, the pattern of gains and losses associated with any act f , viewed from any reference act h , depends on the state-contingent juxtaposition of consequences in f and h . An alternative approach, proposed by Kőszegi and Rabin (2006, 2007), takes no account of state-contingency. Instead, reference-dependent preferences are defined over *prospects* (i.e. probability distributions over consequences). Notice that two acts can have different assignments of consequences to states of the world, and thus be distinct acts, while inducing the same probability distribution over consequences and so being represented by a single prospect. Implicitly, the Kőszegi–Rabin (KR) approach treats such acts as equivalent to one another. It specifies the expected relative value of any prospect A , viewed from any (certain or uncertain) reference prospect R , in such a way as to coincide with the implications of the RDSEU specification (1) in the special case in which A and R are stochastically independent.² Since this condition holds trivially when R is a degenerate prospect, the KR and RDSEU approaches coincide when reference points are certain.

However, the KR approach has paradoxical implications. Consider any two non-degenerate acts f and h that induce the same probability distribution of consequences A . What is the preference ranking of f and h , viewed from h ? The KR approach treats this problem *as if* f and h were stochastically independent. If those acts were in fact independent, moving from h to f would induce gains in some states and losses in others. These gains and losses exactly offset one another in money terms, but if preferences are loss-averse (as is assumed in most versions of prospect theory), the subjective impact of the losses outweighs that of the gains. Expressed in the language of RDSEU, we have $h \phi_h f$, which makes intuitive sense *if the acts really are perceived as stochastically independent*. In fact, there is experimental evidence of such preferences: subjects who have been endowed with numbered lottery tickets are typically unwilling to exchange these for equivalent but differently

numbered tickets, even when offered small payments for doing so (Bar-Hillel and Neter, 1996).³ But in the KR framework, this result has to be expressed in terms of reference-dependent preferences *over prospects*. If A is the prospect that represents the acts h and f, we reach the conclusion that A is strictly preferred to A, viewed from A, with the apparent implication that the agent has a strict preference for ‘keeping’ some lottery rather than ‘exchanging’ it *for exactly the same lottery*. One way to avoid such paradoxical conclusions would be to interpret the KR theory as assuming that each lottery under evaluation is perceived by the agent as stochastically independent of the reference point. On this interpretation, the KR theory is less general than RDSEU, but yields the same conclusions in those cases to which both apply.

The preference representation in (1), like that proposed by KR, is linear in probabilities. PT³ relaxes this restriction of RDSEU by generalising (1) to:

$$(2) \quad V(f, h) = \sum_i v(f[s_i], h[s_i]) W(s_i; f, h)$$

where $W(s_i; f, h)$ is the *decision weight* assigned to state s_i when f is being evaluated from h . In principle, decision weights could be determined by a simple transformation of state probabilities (i.e. $W(s_i; f, h) = w(\pi_i)$) as in Handa (1977). However, in the contemporary literature on prospect theory it has become conventional to construct decision weights cumulatively using a rank-dependent transformation (Quiggin, 1982; Starmer and Sugden, 1989; Tversky and Kahneman, 1992). One of the key theoretical rationales for the cumulative construction is that, unlike first-generation prospect theory (which included some elements of Handa’s approach), it results in *monotonic* preferences. That is, if one prospect stochastically dominates another, the first is preferred to the second when viewed from any certain reference point. In PT³ we retain the rank-dependent approach, but reconfigure it to work with state-contingent reference-dependence.

Given that we have chosen to follow the RDSEU approach to the definition of reference-dependence, there is little remaining freedom of manoeuvre for choosing how to specify rank-dependent decision weights. In order to construct cumulative weights for a given f, h pair, states must be ordered according to the ‘attractiveness’ of f ’s consequences in each state. In a cumulative construction, the weight attached to a given state depends not only on the probability of that state but also on the position of its consequence in the ranking of all consequences associated with f . To ensure that monotonicity is preserved when probabilities are transformed, it is essential that the ‘attractiveness’ ranking is determined by the function

whose decision-weighted expected value represents preferences. In RDSEU, this is the relative value function. Thus, the position of each state s_i in the ranking must be determined by the ranking of $v(f[s_i], h[s_i])$ values. This implies that the ordering of states must be constructed separately for each f, h pair. Further, if we are to generalise cumulative prospect theory, we must have separate rank-dependent transformations of probability for gains and losses, and these transformations must be mirror-images of one another.

Consider any f, h pair. Relative to that pair, there is a *weak gain* in a state s_i if $f(s_i) \geq h(s_i)$, and a *strict loss* if $f(s_i) < h(s_i)$. Let m^+ be the number of states in which there are weak gains and let $m^- = n - m^+$ be the number of states in which there are strict losses. We re-assign subscripts so that, for all subscripts i, j , we have $i > j$ if and only if $v(f[s_i], h[s_i]) \geq v(f[s_j], h[s_j])$, and so that the states with weak gains are indexed $m^+, \dots, 1$ and the states with strict losses are indexed $-1, \dots, -m^-$.

Cumulative decision weights are then defined as follows:

$$(3) \quad W(s_i; f, h) = \begin{cases} w^+(\pi_i) & \text{if } i = m^+, \\ w^+(\sum_{(j \geq i)} \pi_j) - w^+(\sum_{(j > i)} \pi_j) & \text{if } 1 \leq i \leq m^+ - 1, \\ w^-(\sum_{(j \leq i)} \pi_j) - w^-(\sum_{(j < i)} \pi_j) & \text{if } -m^- + 1 \leq i \leq -1, \\ w^-(\pi_i) & \text{if } i = -m^-, \end{cases}$$

where w^+ and w^- are, respectively, probability weighting functions for the gain and loss domains (i.e. w^+, w^- are strictly increasing mappings from $[0, 1]$ onto $[0, 1]$). This specification has two important implications. First, if h is a constant act, (3) implies the same assignment of decision weights as in cumulative prospect theory. Thus, the monotonicity property of cumulative prospect theory carries over to PT^3 . That is, for all acts f, g , and for all *constant* acts h , if the probability distribution of consequences induced by f stochastically dominates that induced by g , then f is strictly preferred to g , viewed from h . Second, a state-contingent form of monotonicity holds for *all* reference acts. For any acts f, g , we will say that f *statewise dominates* g if $f(s_i) \geq g(s_i)$ for all states s_i , with a strict inequality for at least one state with non-zero probability. It follows from (2) and (3) that if f statewise dominates g , then for all reference acts h , $f \phi_h g$.

PT^3 , as specified by (2) and (3), straightforwardly captures several models as special cases. RDSEU is the special case in which decision weights are untransformed state

probabilities (i.e. $w^+(\pi_i) = w^-(\pi_i) = \pi_i$ for all i). Cumulative prospect theory is the special case in which the relative value function takes the form $v(f[s_i], h[s_i]) = u(f[s_i] - h[s_i])$, where $u(\cdot)$ is a ‘value’ function,⁴ and in which reference acts are constrained to be certainties (i.e. $h(s_i) = h(s_j)$ for all i, j). Expected utility theory is the special case in which decision weights are untransformed state probabilities, as in RDSEU, and relative value is independent of the reference outcome (i.e. $v(f[s_i], h[s_i]) = u(f[s_i])$ where $u(\cdot)$ is a von Neumann-Morgenstern utility function).

3. A Parameterised Form of PT³

One way of evaluating PT³ as a descriptive theory is to test its novel predictions – that is, the predictions it makes about behaviour in decision problems with uncertain reference points – using parameterisations whose validity has been established for earlier versions of prospect theory. This approach requires that we select specific functional forms for our general model. In doing this we are guided by three criteria. First, we seek a model flexible enough to allow us to investigate how decision-making behaviour varies with three key aspects of the agent’s preferences: *attitudes to consequences*, *attitudes to probability*, and *attitudes to gain and loss*. Second, subject to that constraint, we seek to use the simplest model possible – that is, a model with just one parameter for each of the three attitudes we consider. Third, for comparability with existing evidence, we use wherever possible the functional forms that are most common in previously published research.

We begin by imposing the restriction that the relative value function takes the form $v(f[s_i], h[s_i]) = u(z)$, where $z = f(s_i) - h(s_i)$. When h is a constant act, this special case of state-contingent reference-dependence is equivalent to that built into earlier generations of prospect theory; $u(\cdot)$ is the counterpart of the value function in those theories.

Next, we specify a functional form for $u(\cdot)$. We adopt the power function which has been widely used in recent empirical literature (see Starmer 2000). Specifically,

$$(4) \quad u(z) = \begin{cases} z^\alpha & \text{if } z \geq 0 \\ -\lambda |z|^\alpha & \text{if } z < 0. \end{cases}$$

The parameters α and λ are required to be strictly positive. The first of these parameters controls the curvature of the value function. If $\alpha < 1$, this function is concave in the domain

of gains and convex in the domain of losses (the property of *diminishing sensitivity*). Diminishing sensitivity imparts a tendency for risk aversion with respect to gains and risk-loving with respect to losses. While the empirical literature has suggested some differences in the exponents of the value function between the domains of gains and losses, in the interests of parsimony we will apply the same exponent in both domains. The parameter λ controls attitudes to gain and loss. With $\lambda = 1$ there is loss neutrality. For λ values above unity, there is loss aversion: losses are weighted more heavily than gains. For values below unity, the opposite is the case.

We model decision weights by means of a single-parameter probability weighting function. Again, for reasons of parsimony we impose the restriction of identical weighting functions for gains and losses (i.e. $w^+(\pi) = w^-(\pi)$). Hence for the purpose of parameterisation the probability weighting function is denoted simply by $w(\pi)$; it takes the form

$$(5) \quad w(\pi) = \pi^\beta / (\pi^\beta + [1 - \pi]^\beta)^{1/\beta}$$

with $\beta > 0$. This type of weighting function has been discussed by Tversky and Kahneman (1992) and Prelec (1998); variants of it have been widely used in the empirical literature. With $\beta = 1$, decision weights are linear (i.e. $w(\pi) = \pi$) but for values of β less than 1 and going down to around 0.4, the function generates an inverse-S pattern of weights with over-weighting (under-weighting) of probabilities below (above) some critical probability π^* . Inverse-S weighting has been reported across a wide range of empirical studies (Wu and Gonzalez, 1996, 1999; Abdellaoui, 2000; Bleichrodt and Pinto, 2000; Abdellaoui, Vossman and Weber, 2005).

That completes the specification of the generic model to be used in the calibrations. We will refer to this specification as *parameterised PT*³.

Notice that, when applied to cases in which reference acts are certainties, parameterised PT³ can also be interpreted as a parameterisation of cumulative prospect theory. In fact, models of this kind have already been estimated using experimental data (e.g. Tversky and Kahaneman, 1992; Loomes, Moffatt and Sugden, 2002). Thus, parameter values from these estimations are applicable to our model.

We take the following to be relatively well-established stylised facts concerning the median values of the three parameters for experimental subjects. First, many studies suggest the existence of loss aversion, while its opposite is almost unknown; values of the loss aversion parameter in the range $1 \leq \lambda \leq 2.5$ would capture a reasonably wide range of

evidence. Studies fitting variants of prospect theory with power utility almost invariably find diminishing sensitivity. Although some studies have found values of α as low as 0.22 (Loomes, Moffatt and Sugden, 2002, note 17), values in the range $0.5 \leq \alpha \leq 1$ are typical. Inverse-S probability weighting, while not universal, is a very common finding; it would be reasonable to expect values of β in the range $0.5 \leq \beta \leq 1$. These ranges of values will be the focus for evaluating the predictions of our model.

4. Explaining WTA/WTP Disparities for Lotteries

We now show that parameterised PT^3 implies that WTA valuations of lotteries exceed WTP valuations whenever $\lambda > 1$. This is not a surprising prediction: intuitively, one might expect the WTA/WTP disparity to be explained by loss aversion, which is represented in our model by the parameter restriction $\lambda > 1$. Nevertheless, as we pointed out in Section 1, previous versions of prospect theory have not been able to make this prediction. Further, it is significant that our model makes a clear-cut prediction about WTA/WTP disparities despite the complications introduced by decision weights.

The hypothesis that individuals tend to prefer to retain status quo positions rather than to move away from them can be expressed in terms of the *non-reversibility* of reference-dependent preferences.⁵ We will say that an agent's preferences are non-reversible if, for all acts f and g , $f \bullet_g g \Rightarrow f \bullet_f g$; they are *strictly* non-reversible if $f \bullet_g g \Rightarrow f \phi_f g$. Intuitively, suppose the agent is endowed with g and is willing to exchange this for f . Non-reversibility implies that if she makes this exchange, and if her reference point adjusts to f , she does not then have a strict preference for reversing the exchange. Given the assumptions built into parameterised PT^3 , strict non-reversibility implies the WTA/WTP disparity. To see why, let g be any act, let h be a constant act which gives consequence z in every state, and suppose that $h \sim_g g$. Then z is the WTA valuation of g . Let h' be a constant act, giving z' in every state, such that $h' \sim_{h'} g$. If income effects are zero, as implied by the assumptions of parameterised PT^3 , h' is the WTP valuation of g . If preferences are strictly non-reversible, the supposition $h \sim_g g$ implies $h \phi_h g$: the agent would not be willing to pay z to get g . Thus, by virtue of the monotonicity properties of PT^3 , the WTP valuation of g is less than the WTA valuation.

We now show that parameterised PT^3 satisfies strict (weak) non-reversibility if $\lambda > 1$ (if $\lambda \geq 1$). Consider any two acts f and g . Let F be the set of states s_i such that $f(s_i) > g(s_i)$ and

let G be the set of states s_j such that $g(s_j) > f(s_j)$. Combing (2) with the parameterisation of the relative value function:

$$(6) \quad V(f, g) = \sum_{i \in F} (f[s_i] - g[s_i])^\alpha W(s_i; f, g) - \sum_{j \in G} \lambda (g[s_j] - f[s_j])^\alpha W(s_j; f, g)$$

$$(7) \quad V(g, f) = \sum_{j \in G} (g[s_j] - f[s_j])^\alpha W(s_j; g, f) - \sum_{i \in F} \lambda (f[s_i] - g[s_i])^\alpha W(s_i; g, f).$$

The specification of decision weights in (3) implies $W(s_i; f, g) = W(s_i; g, f)$ for all $i \in F$ and $W(s_j; f, g) = W(s_j; g, f)$ for all $j \in G$. Hence:⁶

$$(8) \quad V(f, g) + V(g, f) = (1 - \lambda) (\sum_{i \in F} [f(s_i) - g(s_i)]^\alpha W[s_i; f, g] + \sum_{j \in G} [g(s_j) - f(s_j)]^\alpha W[s_j; g, f]).$$

The right-hand side of (8) is positive, zero, or negative according to whether λ is less than, equal to, or greater than unity. Thus, the parameter restriction $\lambda > 1$ implies $V(f, g) + V(g, f) < 0$ or, equivalently, $f \bullet_g g \Rightarrow f \phi_f g$, i.e. strict non-reversibility. Similarly, $\lambda \geq 1$ implies weak non-reversibility.

5. Explaining Preference Reversal

In this section, we apply parameterised PT^3 to preference reversal experiments.⁷ For simplicity, we restrict attention to P and \$ bets that give either a positive payoff or zero. This case has been widely studied in the empirical literature. Consider two acts with this structure. Specifically, let f^P represent an act giving an increment of wealth x with probability p and a zero increment otherwise, and let f^S be an act giving an increment of wealth y with probability q and a zero increment otherwise, with $y > x > 0$ and $1 > p > q > 0$. As a normalisation, we define consequences as increments or decrements of wealth relative to the agent's wealth (treated as a certainty) prior to the PR experiment. This pre-experimental endowment is denoted by the constant act h , where $h(s_i) = 0$ for all i .

A feature of the power utility function used in parameterised PT^3 is that model predictions are unchanged if all outcomes are multiplied by any positive constant. Exploiting this property, we may normalise the expected value of the P bet to unity by setting its payoff $x = 1/p$. Given this normalisation, we can characterise any pair of P and \$ bets by a three-parameter vector (p, q, r) , where p is the probability of winning the prize in the P bet, q is the corresponding probability for the \$ bet, and r is the expected value of the \$ bet as a ratio of the expected value of the P bet (implying that the positive payoff of the \$ bet is $y = r/q$). Notice that the condition $y > x$ (i.e. the \$ bet has the higher prize) implies $rp > q$.

Consider an agent choosing between the two bets in the PR experiment. Her reference point is the constant act h . Using (2), (3) and (4), the agent's choice between the two bets is determined by:⁸

$$(9) \quad f^P \succ_h f^S \Leftrightarrow w(p) (1/p)^\alpha \geq w(q) (r/q)^\alpha;$$

the values of $w(p)$ and $w(q)$ are given by (5).

Now consider the agent's WTA valuations of the bets. Take the case of the P bet. The agent is endowed with this bet and is asked to consider selling it. Her situation is depicted by the following matrix (in which the columns are states, with probabilities shown at the top, the rows are acts, and the entries in the cells are consequences):

	p	1 - p
h^P	1/p	0
g^P	z_P	z_P

Her reference act, denoted h^P , is the P bet. Her WTA valuation of this bet, denoted z_P , is the increment of wealth such that she is indifferent between retaining h^P or giving up h^P in exchange for the certainty of that increment. Hence, we define z_P as the sure payoff of some constant act g^P defined such that $V(g^P, h^P) = 0$. With z_S defined in an analogous way, the values of z_P and z_S are then determined, respectively, by the solutions to equations

$$(10) \quad w(1-p)z_P^\alpha - w(p) \lambda ([1/p] - z_P)^\alpha = 0 \quad \text{and}$$

$$(11) \quad w(1-q)z_S^\alpha - w(q) \lambda ([r/q] - z_S)^\alpha = 0.$$

Substituting in the parameterisation of $w(\cdot)$, the solutions of these equations are:

$$(12) \quad z_P = (1/p) / ([(1-p)/p]^{\beta/\alpha} (1/\lambda)^{1/\alpha} + 1) \quad \text{and}$$

$$(13) \quad z_S = (r/q) / ([(1-q)/q]^{\beta/\alpha} (1/\lambda)^{1/\alpha} + 1).$$

It is easy to see that the preference ranking of the bets, as given by (9), need not be the same as the ranking of their WTA valuations, as given by (12) and (13). A sufficient condition for the two rankings to be the same, and hence for PR *not* to occur, is the combination of parameter values $\lambda = 1$, $\alpha = 1$ and $\beta = 1$, implying that both choice and valuation are determined by the maximisation of expected monetary value. In general,

however, differences between the two rankings can be induced by asymmetric attitudes to gain and loss (i.e. $\lambda \neq 1$), non-linearity of the value function (i.e. $\alpha \neq 1$), or non-linearity of the probability weighting function (i.e. $\beta \neq 1$).

To see that attitudes to gain and loss are relevant, it is sufficient to notice that λ appears in both (12) and (13) but not in (9). To see that the shape of the probability weighting function is relevant, it is sufficient to notice that the terms $w(1 - p)$ and $w(1 - q)$ appear in (10) and (11) respectively, but not in (9). To see that the curvature of the value function is relevant, notice that the choice and valuation expressions use different points on this function. For example, in the case of the P bet (for which the prize is $1/p$), the choice expression includes the term $(1/p)^\alpha$, whereas the valuation expression includes the terms z_P^α and $([1/p] - z_P)^\alpha$.

Given all this, it would not be surprising to find (as is in fact the case) that *some* combinations of parameter values induce standard PR, that is, the conjunction of $f^P \phi f^S$ and $z_P < z_S$. But our objective is to do much more than this: it is to explore whether our model provides an empirically convincing account of observed instances of PR. We therefore need to ask whether standard PR is predicted when the model is calibrated using empirically-established values of λ , α and β .

It is convenient to explore the implications of the model graphically in (α, λ) space. This space is divided into quadrants by the lines $\alpha = 1$ and $\lambda = 1$. The stylised facts presented in Section 3 suggest that we should focus on the north-west quadrant, in which the value function is either linear or exhibits diminishing sensitivity (i.e. $\alpha \leq 1$) and in which there is either loss neutrality or loss aversion (i.e. $\lambda \geq 1$). We call this the *empirically plausible quadrant*.

For any given pair of bets and any given value of the decision weight parameter β , (α, λ) space can be divided into four regions by identifying two boundaries. One boundary – *the choice boundary* – identifies the locus of (α, λ) pairs along which the P and \$ bets are indifferent in choice (i.e. $f^P \sim_h f^S$). A second boundary – *the valuation boundary* – is the locus of (α, λ) pairs along which the P and \$ bets have equal WTA valuations.

The following property of the choice boundary is an immediate implication of (9):

Property 1: The choice between P and \$ is independent of the value of λ . For any given value of β , there is a critical value of α at which the two bets are indifferent. At lower values of α , P is chosen; at higher values, \$ is chosen.

This property reflects the fact that, in the choice task, all consequences are positive or zero. Thus, loss aversion has no role in determining choice. Because the negative domain of the value function is not relevant for this task, diminishing sensitivity (i.e. $\alpha < 1$) plays essentially the same role in PT³ as diminishing marginal utility does in expected utility theory: the lower the value of α , the greater the attractiveness of the safer P bet relative to the riskier \$ bet.

An examination of the expression for z_P/z_S , derived from (12) and (13), yields the following property of WTA valuations:

Property 2: As λ increases, the value of z_P/z_S falls; in the limit, as $\lambda \rightarrow \infty$, this value tends to q/rp , where $q/rp < 1$.

In other words, increases in the loss aversion parameter λ increase z_S relative to z_P ; at sufficiently high values of λ , we have $z_S > z_P$. Intuitively, this is because the act of selling a bet carries the risk of losing the prize of that bet in the state in which the bet wins; since the \$ bet has the higher prize, the potential for loss in selling it is greater. Thus, loss aversion induces a particular reluctance to sell the \$ bet.

Properties 1 and 2 are enough to give a preliminary sense of some of the combinations of parameter values that will induce PR. In order for the P bet to be selected in the choice task, α must be lower than some critical value. Given any such value of α , the \$ bet will have the higher WTA valuation if the value of λ is sufficiently high. Thus, standard PR is induced by the combination of sufficiently low α and sufficiently high λ .

Figure 1 plots the choice and valuation boundaries for a typical pair of bets, defined by $(p, q, r) = (0.8, 0.2, 1)$, with $\beta = 1$. This particular combination of parameters will be called the *benchmark case*. Standard PR occurs in the region above the valuation boundary and to the left of the choice boundary; non-standard PR occurs in the region below the valuation boundary and to the right of the choice boundary. The two boundaries intersect at $(1, 1)$. (This reflects the fact that, when $\alpha = 1$, $\lambda = 1$ and $\beta = 1$, our model reduces to the maximisation of expected value; since the two bets have equal expected value, they are equally preferred and have equal valuations.) In this benchmark case, the empirically plausible quadrant is made up of two sub-regions, separated by the valuation boundary. Above this boundary there is standard PR. Below it, the P bet is both preferred in the straight

choice and valued more highly. Thus, for the benchmark case, our model predicts the classic asymmetry between standard and non-standard reversals: the former occur at parameter values within the empirically plausible quadrant, while the latter do not.

It is easy to see that this qualitative conclusion is unaffected by changes in the values of the probability parameters p and q (given the defining condition $p > q$). Whenever $r = 1$ and $\beta = 1$, the choice and valuation boundaries intersect at $(1, 1)$. Thus, P is chosen everywhere in the empirically plausible quadrant, ruling out the possibility of non-standard PR. There is always a non-empty region of this quadrant, to the left of the choice boundary and above the valuation boundary, in which standard PR occurs.

We now consider the effect of changes in the value of r (the expected value of the \$ bet as a ratio of that of the P bet). Figure 2 plots the choice and valuation boundaries for $r = 0.8$, $r = 1.2$ and $r = 1.4$ when the other parameters take their benchmark values (i.e. $p = 0.8$, $q = 0.2$, $\beta = 1$). As r increases, both boundaries shift to the left, expanding both the region in which the \$ bet is chosen and the region in which it has the higher valuation. The intuition for this is straightforward: an increase in r increases the \$-bet prize relative to the P-bet prize, and so makes the \$ bet relatively more attractive. Notice that, at all values of r , standard PR occurs at some points in the empirically plausible quadrant, while non-standard PR occurs only outside that quadrant. If $r > 1$, the empirically plausible quadrant is made up of *three* sub-regions. Below the valuation boundary, the P bet is favoured in both choice and valuation. To the right of the choice boundary, the \$ bet is favoured in both choice and valuation. Above the valuation boundary and to the left of the choice boundary, there is standard PR.⁹

In the cases we have considered so far, our model has consistently predicted the classic asymmetry between standard and non-standard PR in the empirically plausible quadrant. However, it has failed to predict another stylised fact about PR experiments: that, even when the two bets have equal expected value, a significant proportion of subjects not only value the \$ bet more highly but also choose it in preference to the P bet. Because of the role of diminishing sensitivity in choice, our model predicts that P will be chosen whenever $r = 1$, $\alpha < 1$, and $\beta = 1$. To show that this is not a problem for our approach, we note that the benchmark assumption $\beta = 1$ is an extreme case – the case in which the probability weighting function is linear. We now consider the implications of assuming lower values of β , that is, an inverse-S function.

Figure 3 plots the choice and valuation boundaries for three empirically plausible values of β , namely 0.9, 0.75 and 0.6. The other parameters take their benchmark values (i.e. $p = 0.8, q = 0.2, r = 1$). Essentially, the effect of reducing the value of β is to shift both boundaries to the left, expanding the regions in which the \$ bet is favoured. The intuition for this is that, as the value of β falls, small probabilities (such as 0.2, the probability that the \$ bet wins) are increasingly overweighted while large probabilities (such as 0.8, the probability that the P bet wins) are increasingly underweighted. The resulting configurations of choice and valuation boundaries are similar to those generated by setting $r > 1$. Again, the empirically plausible quadrant of (α, λ) space is made up of three sub-regions. In one, the P bet is favoured in both choice and valuation; in another, the \$ bet is favoured in both choice and valuation; in the third, there is standard PR.¹⁰

All the diagrams we have presented so far have the common feature that standard PR occurs only when $\lambda > 1$, and non-standard PR occurs only when $\lambda < 1$. The reader should not infer from this that loss aversion is essential if PT^3 is to predict standard PR. To the contrary, both standard and non-standard PR are compatible with $\lambda = 1$ for some pairs of bets. Figures 4 and 5 illustrate these possibilities for, respectively, the pairs of bets (0.8, 0.4, 1) and (0.6, 0.2, 1) with $\beta = 0.7$.¹¹ Conversely, PT^3 can predict standard PR in cases in which the model differs from expected utility theory *only* in respect of loss aversion. (Consider the case in which $p = 0.8, q = 0.2, r = 0.8$ and $\beta = 1$, shown in Figure 2. Notice that standard PR occurs at some points at which $\alpha = 1$ and $\lambda > 1$.) But these cases depend on special assumptions about the characteristics of the two bets. In contrast, the classic PR phenomenon occurs across a wide range of values of p, q and r . Because our model includes the effects of loss aversion, diminishing sensitivity and inverse-S probability weighting, it is able to explain PR across the range in which it has been observed.

A further feature of the PR phenomenon has been identified by Tversky, Slovic and Kahneman (1990). Consider any pair of P and \$ bets faced by any given agent. Suppose that standard PR occurs: the P bet is strictly preferred in the choice task, but the \$ bet has the higher WTA valuation. Let g^P and g^S be the constant acts defined so that, for all states s_i , $g^P(s_i) = z_P$ and $g^S(s_i) = z_S$ (where, as before, z_P and z_S are the WTA valuations of the respective bets). Now consider the ranking of the four acts f^P, f^S, g^P, g^S , viewed from the constant reference act h , where $h(s_i) = 0$ for all s_i . The occurrence of standard PR implies $f^P \phi_h f^S$ and (on the assumption that preferences are monotonic) $g^S \phi_h g^P$. How can these apparently conflicting preferences be reconciled? One possibility is that $f^P \phi_h g^P$. That is, given a

pairwise choice between the P bet and a certainty equal to its WTA valuation, the agent would strictly prefer the bet. If this is the case, the agent is said to have *underpriced* the P bet. A second possibility (not incompatible with the first) is that $g^{\$} \phi_h f^{\$}$. That is, in a pairwise choice between the \$ bet and a certainty equal to its WTA valuation, the agent would strictly prefer the certainty. If this is so, the \$ bet has been *overpriced*. If neither of these possibilities is the case, we have $f^P \phi_h g^P$, $g^P \bullet_h f^P$, $f^P \phi_h f^{\$}$, and $f^{\$} \bullet_h g^{\$}$, which is a violation of transitivity. Tversky et al. report an experiment designed to discriminate between these possibilities; they find that standard PR is most commonly associated with overpricing of the \$ bet.

This observation is consistent with PT³. At the parameter values that induce standard PR, our model predicts overpricing of both the \$ and P bets. We show this for the \$ bet; the analysis for the P bet is essentially the same. The \$ bet is overpriced if it is strictly preferred to the certainty of $z_{\$}$, viewed from h. This is the case if and only if $z_{\$}^{\alpha} > w(q) (r/q)^{\alpha}$, which can be rearranged as

$$(14) \quad (1 - w[q])z_{\$}^{\alpha} - w(q)([r/q]^{\alpha} - z_{\$}^{\alpha}) > 0.$$

Combining this inequality with (10), the equation which defines $z_{\$}$, the condition for overpricing is:

$$(15) \quad (1 - w[q])z_{\$}^{\alpha} - w(q)([r/q]^{\alpha} - z_{\$}^{\alpha}) > w(1 - q)z_{\$}^{\alpha} - w(q)\lambda([r/q] - z_{\$})^{\alpha}.$$

As one would expect, the LHS and RHS of (15) are equal (and equal to zero) if $\lambda = 1$, $\alpha = 1$ and $\beta = 1$ (implying $w(q) = q$ and $w(1 - q) = 1 - q$). In this case, $z_{\$} = r$, i.e. the WTA valuation of the bet is its expected monetary value, and there is neither underpricing nor overpricing.

Now consider the implications of variations in the agent's attitude to gain and loss. It is easy to see that, if α and β are held constant at the values $\alpha = 1$ and $\beta = 1$, the inequality (15) is satisfied if and only if $\lambda > 1$. In other words, loss aversion induces overpricing. Alternatively, consider the implication of variations in the agent's attitude to consequences. If λ and β are held constant at $\lambda = 1$ and $\beta = 1$, (15) is satisfied if and only if $\alpha < 1$.¹² Thus, diminishing sensitivity induces overpricing. Finally, consider variations in the agent's attitude to probability. If λ and α are held constant at $\lambda = 1$ and $\alpha = 1$, (15) is satisfied if and only if $\beta < 1$.¹³ Thus, inverse-S probability weighting induces overpricing.

Bringing these results together, PT³ explains PR as a result of the interaction of empirically plausible degrees of loss aversion, diminishing sensitivity and probability weighting. The same attitudes induce the ‘overpricing’ effect observed by Tversky et al.

6. Conclusions

We have presented a new theory of choice under uncertainty: third generation prospect theory (PT³). PT³ retains the empirically grounded features of previous variants of prospect theory (loss aversion, diminishing sensitivity and non-linear probability weighting), but extends that theory by allowing reference points to be uncertain. The resulting theory retains all the predictive power of those previous variants, but in addition provides a framework for determining the money valuation that an agent places on a lottery. We have shown that PT³ predicts the observed tendency for willingness-to-accept (WTA) valuations of lotteries to be greater than willingness-to-pay (WTP) valuations. More surprisingly, when PT³ is made operational by using simple functional forms with parameter values derived from existing experimental evidence, it predicts observed patterns of preference reversal (PR) across a wide range of specifications of P and \$ bets, consistent with the range in which that phenomenon has in fact been observed.

PR is one of the most notorious anomalies in individual decision making, but despite the large volume of literature it has generated, no satisfactory preference-based account of it has thus far been produced. In the psychology literature it has been common to interpret PR as evidence that preferences do not satisfy procedural invariance but, instead, depend upon the method used to elicit them. On this view, if preferences are to be invoked at all in explaining PR, those preferences must be context-sensitive: that is, they must allow *different* preferences to govern decisions in choice and valuation tasks. In the economics literature, various models of context-free preferences have been proposed as possible accounts of PR. One approach is to relax the independence and/or reduction axioms of expected utility theory (Holt, 1986; Karni and Safra, 1987; Segal, 1988). Subsequent studies, however, have generated strong PR in experimental designs implementing controls for the explanations postulated in these theories (Tversky, Slovic and Kahneman, 1990; Cubitt, Munro and Starmer, 2004). Another possible explanation is that PR arises as a consequence of context-free, but non-transitive preferences. Persistent non-transitive cycles of choice analogous to PR have been observed in experimental studies (Loomes, Starmer and Sugden, 1989, 1991; Humphrey 2001), but the only preference theory that has been put forward to explain such behaviour is regret theory

(Bell, 1982; Loomes and Sugden, 1983), which has failed other tests (Starmer and Sugden 1998). A third approach is to explain PR by supplementing a conventional theory of preferences with some mechanism of stochastic error; but this has been shown to be empirically unconvincing (Schmidt and Hey, 2004). Against this unpromising background, it is remarkable to find that observed patterns of PR are predicted by a simple extension of an existing and empirically well-supported preference-based theory.

We do not claim that PT^3 provides a *complete* explanation of PR. We recognise that psychologists have proposed credible non-preference mechanisms of context-sensitive choice and valuation behaviour that are consistent with observations of PR. Predictions based on those mechanisms have been tested and confirmed in experimental tasks other than PR and, in some cases, outside the domain of theories of choice under uncertainty (Slovic, Griffin and Tversky, 1990). This evidence clearly suggests that non-preference mechanisms contribute to PR. We assert only that PT^3 has a similar claim to be a model of mechanisms which contribute to that phenomenon. It too is based on psychologically credible hypotheses – loss aversion, diminishing sensitivity, the overweighting of small probabilities and the underweighting of large ones. It too is consistent with observations of PR. It too has been tested and confirmed in experimental tasks other than PR – namely, pairwise choices between lotteries involving gains and losses. If one accepts prospect theory as an explanation of observed regularities in choice among lotteries, it seems reasonable to infer that the mechanisms modelled by PT^3 play a significant role in the explanation of PR.

We should also acknowledge that we have treated the use of WTA valuations as one of the defining characteristics of a PR experiment. Insofar as it relies on the assumption of loss aversion, the explanation of PR provided by PT^3 is specific to such valuations. In fact, there have been surprisingly few experiments in which WTP valuations of P and \$ bets have been used. Such experiments have produced mixed results, but asymmetric PR is generally less pronounced than in WTA experiments, and sometimes is not present at all. It seems that WTP treatments tend to reduce the frequency of standard reversals and to increase the frequency of non-standard ones (Lichtenstein and Slovic, 1971; Knez and Smith, 1987; Casey, 1991). These findings are compatible with the hypothesis that PR is the product of several causal mechanisms, at least one of which is in some way linked to WTA valuations. We suggest that PT^3 captures a mechanism of the latter kind.

More generally, we offer PT^3 as a natural extension of prospect theory – a theory that is already widely accepted as empirically successful. By allowing reference points to be

uncertain, we have filled a major gap in that theory's previous domain of application. The result is a flexible and parsimonious model of choice under uncertainty which organises a large body of experimental evidence. We hope that it will find fruitful applications in future work.

References

- Abdellaoui, Mohammed. (2000). "Parameter-Free Elicitation of Utilities and Probability Weighting Functions," *Management Science* 46, 1497-1512.
- Abdellaoui, Mohammed, Frank Vossman and Martin Weber. (2005). "Choice-based Elicitation and Decomposition of Decision Weights for Gains and Losses under Uncertainty," *Management Science* 51, 1384-1399.
- Bar-Hillel, Maya, and Efrat Neter. (1996). "Why Are People Reluctant to Exchange Lottery Tickets?" *Journal of Personality and Social Psychology* 70, 17-27.
- Baucells, Manel, and Franz H. Heukamp. (2006). "Stochastic Dominance and Cumulative Prospect Theory," *Management Science* 52, 1409-1423.
- Bell, David. (1982). "Regret in Decision Making under Uncertainty," *Operations Research* 20, 961-81.
- Birnbaum, Michael, H. and Jeffrey P. Bahra. (2007). "Gain-Loss Separability and Coalescing in Risky Decision Making," *Management Science* 53, 1016-1028.
- Bleichrodt, Han, and Jose. L. Pinto. (2000). "A Parameter-Free Elicitation of the Probability Weighting Function in Medical Decision Analysis," *Management Science* 46, 1485-1496.
- Camerer, Colin F. (1989). "An Experimental Test of Several Generalized Utility Theories," *Journal of Risk and Uncertainty* 2, 61-104.
- Casey, Jeff. (1991). "Reversal of the Preference Reversal Phenomenon," *Organizational Behavior and Human Decision Processes* 48, 224-251.
- Cubitt, Robin P., Alistair Munro and Chris Starmer. (2004). "Testing Explanations of Preference Reversal," *Economic Journal* 114, 709-726.
- Davies, Greg B. and Stephen, E. Satchell. (2004). "Continuous Cumulative Prospect Theory and Individual Asset Allocation," Cambridge Working Papers in Economics (CWPE 0467).
- Handa, Jagdish. (1977). "Risk, Probabilities, and a New Theory of Cardinal Utility," *Journal of Political Economy* 85, 97-122.
- Holt, Charles, A. (1986). "Preference Reversals and the Independence Axiom," *American Economic Review* 76, 508-515.

- Humphrey, Stephen J. (2001). "Non-Transitive Choice: Event-Splitting Effects or Framing Effects?," *Economica*, 68, 77-96.
- Kahneman, Daniel, and Amos Tversky. (1979). "Prospect Theory: an Analysis of Decision under Risk," *Econometrica* 47, 263-291.
- Karni, Edi, and Zvi Safra. (1987). "Preference Reversals and the Observability of Preferences by Experimental Methods," *Econometrica* 55, 675-685.
- Knetsch Jack, and John A. Sinden. (1984). "Willingness to Pay and Compensation Demanded: Experimental Evidence of an Unexpected Disparity in Measures of Value," *Quarterly Journal of Economics* 99: 507-521.
- Knez, Marc, and Vernon Smith. (1987). "Hypothetical Valuations and Preference Reversals in the Context of Asset Trading," in Roth, A. (ed.), *Laboratory Experimentation in Economics: Six Points of View*, Cambridge, Cambridge University Press, pp. 131-154.
- Kőszegi, Botond, and Matthew Rabin. (2006). "A Model of Reference-Dependent Preferences," *Quarterly Journal of Economics* 121: 1133-1165.
- Kőszegi, Botond, and Matthew Rabin. (2007). "Reference-Dependent Risk Attitudes," *American Economic Review* 97: 1047-1073.
- Lichtenstein, Sarah, and Paul Slovic. (1971). "Reversals of Preferences between Bids and Choices in Gambling Decisions," *Journal of Experimental Psychology* 89, 46-55.
- Loomes, Graham, Peter Moffatt and Robert Sugden. (2002). "A Microeconomic Test of Alternative Stochastic Theories of Risky Choice," *Journal of Risk and Uncertainty*, 24, 103-130.
- Loomes, Graham, Chris Starmer and Robert Sugden. (1989). "Preference Reversal: Information Processing or Rational Non-Transitive Choice?," *Economic Journal* 99, Supplement, 140-151.
- Loomes, Graham, Chris Starmer and Robert Sugden. (1991). "Observing Violations of Transitivity by Experimental Methods," *Econometrica* 59, 425-439.
- Loomes, Graham, Chris Starmer and Robert Sugden. (2003). "Do anomalies disappear in repeated markets?" *Economic Journal* 113, C153-166.
- Loomes, Graham, and Robert Sugden. (1983). "A Rationale for Preference Reversal," *American Economic Review* 73, 428-432.

- Munro, Alistair, and Robert Sugden. (2003). "On the Theory of Reference-Dependent Preferences," *Journal of Economic Behavior and Organization* 50, 407-28.
- Luce, Duncan, and Peter C. Fishburn. (1991). "Rank- and Sign-Dependent Linear Utility Models for Finite First-Order Gambles," *Journal of Risk and Uncertainty* 4, 29-59.
- Prelec, Drazen. (1998). "The Probability Weighting Function," *Econometrica* 66, 497-527.
- Quiggin, John. (1982). "A Theory of Anticipated Utility," *Journal of Economic Behavior and Organization*, 3, 323-343.
- Savage, Leonard, J. (1954). *The Foundations of Statistics*. New York: Wiley.
- Schmidt, Ulrich, and John D. Hey. (2004). "Are Preference Reversals Errors? An Experimental Investigation," *Journal of Risk and Uncertainty* 29, 207-218.
- Schmidt, Ulrich, and Horst Zank. (2008). "Risk Aversion in Cumulative Prospect Theory," *Management Science* 54, 208-216.
- Segal, Uzi. (1988). "Does the Preference Reversal Phenomenon Necessarily Contradict the Independence Axiom?," *American Economic Review* 78, 233-236.
- Slovic, Paul, Dale Griffin and Amos Tversky. (1990). "Compatibility Effects in Judgment and Choice," in Hogarth, Robin., (ed.) *Insights in Decision-Making*, Chicago, University of Chicago Press, pp. 5-27.
- Starmer, Chris. (2000). "Developments in Non-expected Utility Theory: The Hunt for a Descriptive Theory of Choice under Risk," *Journal of Economic Literature* 38, 332-382.
- Starmer, Chris, and Robert Sugden. (1989). "Probability and Juxtaposition Effects: An Experimental Investigation of the Common Ratio Effect," *Journal of Risk and Uncertainty* 2, 159-178.
- Starmer, Chris and Robert Sugden. (1998). "Testing Alternative Explanations of Cyclical Choices," *Economica* 65, 347-361.
- Sugden, Robert. (2003). "Reference-dependent Subjective Expected Utility," *Journal of Economic Theory* 111, 172-191.
- Trepel, Christopher, Craig R. Fox and Russell A. Poldrack. (2005). "Prospect Theory on the Brain? Toward a Cognitive Neuroscience of Decision under Risk," *Cognitive Brain Research* 23, 34-50.

- Tversky, Amos, and Daniel Kahneman. (1991). “Loss Aversion in Riskless Choice: A Reference-Dependent Model,” *Quarterly Journal of Economics* 106, 1039-1061.
- Tversky, Amos, and Daniel Kahneman. (1992). “Advances in Prospect Theory: Cumulative Representation of Uncertainty,” *Journal of Risk and Uncertainty* 5, 297-323.
- Tversky, Amos, Paul Slovic and Daniel Kahneman. (1990). “The Causes of Preference Reversal,” *American Economic Review* 80, 204-217.
- Wakker, Peter P., and Amos Tversky. (1993). “An Axiomatization of Cumulative Prospect Theory,” *Journal of Risk and Uncertainty* 7, 147-176.
- Wu, George and Richard Gonzalez. (1996). “Curvature of the Probability Weighting Function,” *Management Science* 42, 1676-1690.
- Wu, George and Richard Gonzalez. (1999). “Nonlinear Decision Weights in Choice under Uncertainty,” *Management Science* 45, 74-85.
- Wu, George, Jiao Zhang, and Mohammed Abdellaoui. (2005). “Testing Prospect Theory using Probability Tradeoff Consistency,” *Journal of Risk and Uncertainty* 30, 107-131.

Notes

¹ Some descriptive limitations of prospect theory are discussed by Birnbaum and Bahara (2007).

² Compare (1) above with equation (2) in Köszegi and Rabin (2007).

³ In Bar-Hillel and Neter’s ‘Experiment 2’, subjects are unwilling to exchange lottery tickets even if they know that the number of the winning ticket will not be announced. This suggests that subjects use a state-contingent conceptualisation of gains and losses even if they will never know which state is realised.

⁴ The restriction $v(f[s_i], h[s_i]) = u(f[s_i] - h[s_i])$ prevents cumulative prospect theory from taking account of income effects. Kahneman and Tversky (1979, pp. 277-278) comment that ‘strictly speaking’, value should be defined as a function in two arguments – changes in wealth relative to a current asset position, and that position itself. The simpler function they use is presented as ‘a satisfactory approximation’.

⁵ This idea was introduced by Tversky and Kahneman (1991), and developed by Munro and Sugden (2003), to characterise reference-dependence in preferences over multi-dimensional consumption bundles.

⁶ Notice that the derivation of this result uses two symmetries between the treatments of gains and losses that are specific to our parameterisation: the weighting function is the same for gains and losses, and the exponent of $u(\cdot)$ is the same for gains and losses.

⁷ Our analysis of PR does not depend on the cumulative transformation of probabilities. The acts that we analyse have no more than one strictly positive consequence and no more than

one strictly negative one. For such acts, the cumulative transformation is observationally equivalent to Handa's simple transformation.

⁸ In addition, $f^P \sim_h f^S \Leftrightarrow w(p)/w(q) = (pr/q)^\alpha$. From now on, to avoid cluttering the exposition, we will not state conditions for indifference explicitly. In all cases, the condition for indifference can be constructed from the condition for weak preference by substituting an equality for a weak inequality.

⁹ The reader may wonder why, in all the cases represented in Figure 2, the valuation boundary intersects the choice boundary at $\lambda = 1$. It can be shown that this property is induced by a special feature of the benchmark case, namely that the P and S bets are *symmetrical* in the sense that $q = 1 - p$. If the bets are symmetrical and if $\beta = 1$, the valuation boundary passes through the point $(\ln[p/q] / \ln[rp/q], 1)$. Whether the bets are symmetrical or not, this point lies on the choice boundary. (Since this result has little substantive significance, we leave the proof to sufficiently curious readers.) If the assumption of symmetry is relaxed, the two boundaries may intersect at positive or negative values of λ . If $r > 1$ and if the intersection is at a positive value of λ , there is a (typically small) region of the empirically plausible quadrant at which non-standard PR occurs.

¹⁰ The reader may have noticed that, in all three cases, the choice and valuation boundaries intersect at $(\beta, 1)$. On the assumption that $r = 1$, it can be shown that, for all admissible values of p and q , the valuation boundary passes through $(\beta, 1)$. The choice boundary passes through the same point if and only if $q = 1 - p$ (compare note 9). Figures 4 and 5, discussed in the next paragraph, illustrate some cases in which this symmetry condition does not hold.

¹¹ In each case, the valuation boundary passes through $(\beta, 1)$, illustrating the general result stated in note 10. In general, if $\beta < 1$ and $r = 1$, the choice boundary lies to the right of (respectively passes through, lies to the left of) the point $(\beta, 1)$ if $p + q$ is greater than (equal to, less than) 1. The proof of this result is omitted for brevity.

¹² If $\lambda = 1$ and $\beta = 1$, (11) implies $0 < z_S < r/q$, while (15) reduces to $([r/q] - z_S)^\alpha > (r/q)^\alpha - z_S^\alpha$.

¹³ If $\lambda = 1$ and $\alpha = 1$, (15) reduces to $1 - w[q] > w(1 - q)$. This inequality holds if and only if $\beta < 1$.