

Preference Anomalies, Preference Elicitation and the Discovered Preference Hypothesis

JACINTO BRAGA and CHRIS STARMER*

*School of Economics, University of Nottingham, Nottingham, NG7 2RD, UK; *Author for correspondence (e-mail: chris.starmer@nottingham.ac.uk)*

Accepted 11 March 2005

Abstract. There is wide-ranging evidence, much of it deriving from economics experiments, of ‘anomalies’ in behaviour that challenge standard preference theories. This paper explores the implications of these anomalies for preference elicitation methods. Because methods that are used to inform public policy, such as contingent valuation, are based on standard preference theories, their validity may be called into question by the anomaly data. However, on a new interpretation, these anomalies do not contradict standard theory but are errors in stated preference that can be expected to disappear as people become more experienced in relevant decision environments. We explore the evidence for this interpretation and what implications follow for preference elicitation methodology.

Key words: contingent valuation, experiments, learning, preference elicitation, preference theory

JEL classifications: C91, D01, D60, D81, D83, H40

1. Introduction

This paper explores the implications for preference elicitation of observed “anomalies” in behaviour relative to standard economic theories of preference. Theories such as the Hicksian theory of consumer preferences, or the expected utility model of risk preference, lie at the heart of modern economic theory. They are fundamental ingredients in models used to predict and to normatively evaluate economic outcomes and they have also played important roles in various aspects of public policy formation. In particular, models like this provide the theoretical underpinnings of various tools of preference elicitation such as those used in cost benefit analysis. There is, however, a substantial body of evidence – much of it deriving from experimental tests of preference theory, but some from field data – which suggests that standard preference models are inadequate as descriptive representations of the preferences of ordinary people.¹ Taken at face value, such evidence raises questions about the validity of conclusions drawn from policy appraisal techniques grounded on models challenged by the existing data.

The evidence has prompted some theorists to develop new models of preference to explain the anomaly evidence but, although some work is

underway, the significance of these efforts at the foundations of economics has yet to fully percolate through to practical questions of policy formation. There is, however, emerging evidence that at least some anomalies apparently decay in certain kinds of repeated choice environment. One interpretation of this evidence is that anomalies are errors in stated preference that can be expected to disappear in environments that foster certain kinds of learning. If this is the right way to interpret anomaly evidence, it suggests a quite different agenda for those seeking to develop methods of preference elicitation with a view to providing reliable data for input to public policy. On this view, we should not be seeking to alter methods to reflect changes in the theoretical foundations of economics. Rather, we should work within the standard preference theoretic framework and develop elicitation methods that incorporate suitable controls for the types of errors that may arise in stated preference.

In this paper, we seek to evaluate this new interpretation of the anomaly evidence and, in light of that, consider what implications follow for preference elicitation methodology. In Section 2, we explain why preference anomalies might be taken to constitute significant problems for preference elicitation and welfare economics more generally. In Section 3, we consider alternative ways of responding to anomaly evidence and introduce a specific hypothesis leading to an interpretation of anomalies as errors relative to standard preference theory: that is the Discovered Preference Hypothesis proposed by Plott (1996). Section 4 examines the extent of support for this hypothesis and Section 5 explores some possible mechanisms that might explain its workings. Section 6 attempts to extract some implications for elicitation methodology. Some tentative conclusions are drawn in Section 7.

2. Some Preference Anomalies and Why They Matter for Public Policy

To motivate the case for thinking that preference anomalies might matter for public policy, we present some widely discussed cases: the Allais paradoxes, the WTA/WTP disparity, the closely related endowment effect and the preference reversal phenomenon. In much of the discussion that follows, we will draw on evidence relating to these phenomena but it is important to keep in mind that these specific examples are part of a broader literature documenting many other preference anomalies.

The Allais paradoxes (Allais 1953) are widely observed choice patterns that appear to violate the independence axiom of expected utility theory (EUT). The independence axiom states that for any three prospects² Q , R , and S , and any probability $p \in (0,1)$, the compound prospect $S' = (S, p; Q, 1 - p)$ – which offers S with probability p , otherwise Q – is preferred to $R' = (R, p; Q, 1 - p)$ if and only if S is preferred to R . The normative rationale for this axiom is that when choosing between S' and R' one should

attend only to what is different between them. While many people find this rather compelling as a normative principle, for appropriately chosen Q , R , S , and p , many people also tend to systematically violate it.

The common ratio effect is one such violation. It involves two pairs of binary prospects. One pair consists of a safer prospect, S , and a riskier one, R . To illustrate with an example from Kahneman and Tversky (1979), $S = (3000, 1)$ and $R = (4000, 0.8; 0, 0.2)$ (monetary values are Israeli pounds). The other pair, S' and R' , can be obtained from the first one by scaling down by a common factor the probabilities of the best outcomes of S and R . Again in Kahneman and Tversky's study, $S' = (3000, 0.25; 0, 0.75)$, and $R' = (4000, 0.2; 0, 0.8)$. It can be seen that $S' = (S, 0.25; 0, 0.75)$ and $R' = (R, 0.25; 0, 0.75)$. Thus, the independence axiom implies that S' should be preferred to R' if and only if S is preferred to R . However, people tend to choose S over R and R' over S' . In Kahneman and Tversky's study, 80% of subjects chose S , but only 35% chose S' .

The other famous Allais example is the common consequence effect. This also involves two pairs of prospects. In one pair, two prospects S and R share a common consequence. For instance, in Allais' (1953) example, letting m denote a million, and substituting pounds for the original French francs, $S = (£1\text{ m}, 1)$ and $R = (£5\text{ m}, 0.1; £1\text{ m}, 0.89; £0, 0.01)$. These prospects share a common consequence of £1 m with probability 0.89 (this becomes more apparent if we represent $S = (£1\text{ m}, 0.1; £1\text{ m}, 0.89; £1\text{ m}, 0.01)$). The second pair is obtained from the first by replacing the common consequence with a worse one, in Allais' example, £0. Thus, $S' = (£1\text{ m}, 0.11; £0, 0.89)$ and $R' = (£5\text{ m}, 0.1; £0, 0.9)$. The independence axiom implies that S' should be preferred to R' if and only if S is preferred to R . Yet, with pairs appropriately crafted, as in the present example, people tend to choose S and R' .

The independence axiom is a crucial building block of expected utility theory. So, its failure calls into question that theory and a rather wide range of prescriptive decision techniques based on that theory. For instance, failure of the independence axiom raises questions about the validity of any technique that evaluates risks by simply weighting some measure of the value of outcomes by probabilities. This includes, any approach that uses standard decision tree analysis identifying optimal paths through trees as those that maximise expected (i.e. probability weighted) payoffs. To see why, consider a choice between two risky medical interventions I_1 and I_2 . Suppose that I_1 results in H_1 with probability p , and with probability $1 - p$ results in any probability distribution over health states denoted Q . I_2 , by comparison, generates H_2 with probability p , otherwise Q . Now consider the position of an analyst who has some (let us assume valid) utility measure over individual health states and is assessing the relative merits of I_1 and I_2 . Notice that, if EUT is used to form this assessment, the ranking of the interventions will necessarily be independent of the value of p . But for an agent whose

preferences exhibit a common ratio effect, their own ranking may vary systematically with the value of p . Similarly, while an expected utility analysis will conclude that the specification of Q is irrelevant to the choice of intervention, for an agent whose preferences exhibit a common consequence effect, their preferred intervention may vary as the common component Q changes. The more general point is that observed violations of EUT provide the basis for concern about whether standard decision-analytic techniques built upon that theory will accurately reflect the desires of agents who may have non-expected utility preferences.

Our second illustrative anomaly is the well-known WTA/WTP disparity observed in applications of contingent valuation. Contingent valuations are typically elicited as willingness to pay (WTP for short) or willingness to accept (WTA for short). The standard (Hicksian) model of consumer preferences implies that these two measures when applied in any specific context should generate similar valuations, yet it has been widely observed in both field data and laboratory experiments that measured WTA is often several times larger than measured WTP: this observed discrepancy is the WTA/WTP disparity (see Bateman et al. 1997 for a discussion of the empirical findings). While there have been intense debates about how large a WTA/WTP disparity can be explained consistently with Hicksian theory,³ experimental tests embodying controls to eliminate accounts of the WTA/WTP disparity consistent with standard theory have reproduced the disparity in one-off decisions (see for example, Knetsch and Sinden 1984; Knetsch 1989; Bateman et al. 1997). This is not, it seems, a phenomenon that can be explained consistently with standard preference theory.

The WTA/WTP disparity has immediate significance in relation to public policy partly because it relates to a generic problem of needing to compare the “worth” of qualitatively different goods. The problem naturally arises from the fact that most public policies tend to generate both upside gains and downside losses across qualitatively different outcome domains. For example, reducing speed limits on motorways may save lives, but cost time. So, any policy determining speed limits (including, implicitly, the *status quo*) implies some trade-off between these dimensions. But what is the appropriate trade-off between timesaving and lifesaving? In principle, contingent valuation methodology provides a tool for measuring – and hence ultimately comparing – upside and downside effects on a common scale. The WTA/WTP disparity, however, raises awkward questions about this otherwise extremely useful technique such as: why does the disparity occur; which of WTP and WTA is the best measure of value; do WTP and WTA measure different things; is either of them capturing a meaningful measure of anything at all?

The WTA/WTP disparity can be thought of as an instance of the so-called endowment effect, which has been observed in simple choice experiments and field data. Consider two goods, A and B , and three different decision prob-

lems: in one you are endowed with A and given the choice of exchanging it for B ; in another you are endowed with B and given the choice of exchanging it for A ; finally, suppose you are endowed with neither good but can choose to have one of them. Given a standard preference model, these problems are equivalent ways of asking which good, A or B , you prefer and the preference ordering elicited should not depend on the initial endowment. Contrary to this prediction, there is a tendency for agents to prefer the good they are endowed with – this is the endowment effect (see Knetsch 1989). Some have interpreted this as compelling evidence of *loss aversion* in preferences. Models incorporating loss aversion (e.g., Tversky and Kahneman 1991) assume that agents evaluate changes in their circumstances as gains and losses relative to some reference point (typically the *status quo*). Loss aversion can then be captured by assuming that agents are relatively more sensitive to losses than to corresponding magnitude gains. It follows that agents will be less willing to swap between a pair of goods if endowed with one of them (rather than endowed with none) because giving up the endowed good is evaluated as a loss relative to the status quo. If loss aversion is a characteristic feature of human preferences this has immediate practical significance for contingent valuation methodology. Jack Knetsch in a paper in this Special Issue argues with conviction that – contrary to conventional wisdom – WTA is the correct measure for welfare changes which occur in dimensions of loss and that applying WTP will lead to potentially very significant underestimates of true values. Knetsch's argument rests, in good measure, on the claim that loss aversion is a feature of human preferences and the endowment effect is a key part of the evidence supporting that presumption.

Our last illustrative anomaly is the preference reversal phenomenon. This was first reported in experiments that required individuals to carry out two types of tasks. In one type of task people choose between two simple gambles, one of which (often called the \$-bet) offers a small chance of winning a "good" prize; the other (the "P-bet") offers a larger chance of winning a smaller prize. The second type of task asks subjects to state money values – usually elicited as minimum selling prices – separately for each of the two bets. The preference reversal phenomenon is a tendency for individuals to pick the P-bet in the choice task but place a higher value on the \$-bet (see Seidl 2002 for a review of evidence). It presents a puzzle for economics because, viewed from the standard theoretical perspective, both tasks constitute ways of asking essentially the same question, that is, "which of these two bets do you prefer?" In these experiments, however, the preference ordering revealed appears to depend upon the elicitation procedure. The phenomenon appears not to be caused by random errors, since the pattern of reversal is clearly asymmetrical: many people choose the P-bet and value the \$-bet more highly, while very few do the opposite (that is, choose \$ and value P more highly).

Preference reversal raises a very awkward question: if choices and valuations reveal different preference orderings, which, if either, reflect true preference? Without an answer to this question we do not know on which elicitation methods, if any, we can rely for obtaining sound preference data. Of course we should not presume that puzzles observed in simple experiments with gambles necessarily extend to other, policy relevant contexts, but there are some indications that they may. For instance, Slovic (1995) gives an example of preference reversal in an experiment involving choices between and valuations of improved air quality and a computer upgrade. 41% of subjects chose the improved air quality, but placed a higher monetary value on the computer upgrade. A phenomenon akin to preference reversal, *the choice-matching discrepancy*, was observed by Tversky et al. (1988) in a policy relevant context. One study involved two beach clean-up programmes: a “comprehensive” scheme costing \$750,000 and a “partial” one costing \$250,000. Subjects expressed their preferences between schemes in either a choice task or a matching task. In the matching task, subjects were presented with the two schemes but the cost of one was omitted, and they were required to identify a value for the missing cost which would make the two schemes equally attractive. The comprehensive programme was revealed preferred in 48% of the choices but in only 12% of the rankings inferred from the matching values. These findings indicate that inconsistencies between preferences elicited via choice and different forms of valuation may be fairly pervasive and extend to policy relevant contexts.

The potential significance of anomalies runs far beyond their implications for specific techniques of preference elicitation. For example, a feature shared by preference reversal, the WTA/WTP disparity and indeed other preference anomalies is that, on face value interpretation, they appear to reflect failures of *procedure invariance*. The principle of procedure invariance is the widely held methodological presumption that, when a given theoretical system postulates the existence of some entity (e.g. gravitational force, certainty equivalent) which is in principle measurable, then alternative ways of measuring it, which are equivalent under the theory concerned, should not produce systematically different answers. When theoretically equivalent techniques produce inconsistent measures, that raises questions about the validity of an underlying theoretical scheme which implies their equivalence. In the case of the WTA/WTP disparity, two techniques for attaching money values to a good, which should theoretically produce similar answers, often generate very different measured values. In the case of preference reversal, two procedures for assessing the preference ordering of a pair of gambles generate inconsistent rankings. In this sense, preference anomalies raise very basic concerns about the preference theoretic foundations of economics warranting serious thought and investigation. These concerns feed directly through to policy analysis built upon standard theoretical foundations. As

Shogren (2003) puts it in a recent survey article considering the implications of anomalies for environmental policy:

“These anomalies, and many others, matter for the rational valuation of environmental protection because they undercut the internal validity of the cost–benefit estimates economists use to evaluate alternate protection strategies... Challenges to rational choice theory are threats to rational nonmarket valuation”

3. Responding to Anomalies – the Discovered Preference Hypothesis

One line of response to anomalies taken by some theorists has been to develop new preference models to explain them. Work in this genre includes the development of non-standard models of risk preference (see Starmer 2000) and adaptations of Hicksian preference theory such as those incorporating reference points and loss aversion. Amongst these new theories are models that can potentially accommodate the WTA/WTP disparity (Tversky and Kahneman 1991) or preference reversal (see Loomes and Sugden 1983) and some models can do both (e.g. Sugden 2003, Schmidt et al. 2004). The displacement of standard models in favour of non-standard alternatives holds the implication that some techniques of policy analysis – i.e. those based on standard models – may be unreliable or inappropriate. It may, nevertheless, be feasible to develop counterpart techniques that rely on improved (non-standard) theoretical foundations. One example of work in this spirit is Bleichrodt et al. (2001), who developed utility elicitation procedures that are robust to specific violations of EUT (i.e., non-linear probability weighting). Robert Sugden’s contribution to this Special Issue takes a more radical line and responds to anomalies by proposing new foundations for welfare economics that require very little of the usual preference theoretic architecture.

But not all economists have been convinced that experimentally observed preference anomalies make a strong enough case for departing from standard theoretical models. For example, Ken Binmore (1994) strikes a sceptical note arguing that anomalous behaviour is not interesting unless it survives some key tests of robustness:

“But how much attention should we pay to experiments that tell us how inexperienced people behave when placed in situations with which they are unfamiliar, and in which the incentives for thinking things through carefully are negligible or absent altogether?... In brief, two questions about experiments with human subjects always need to be asked. Does the behavior survive when the incentives are increased? Does it survive after the subjects have had a long time to familiarize themselves with all the wrinkles of the unusual situation in which the experimenter has

placed them? If not, then the experimenter has probably done no more than inadvertently to trigger a response in the subjects that is adapted to some real-life situation, but which bears only a superficial resemblance to the problem the subjects are really facing in the laboratory” (Binmore 1994, pp. 184–185).

This is a head-on challenge to the significance of large swathes of anomaly evidence generated in one-shot decision settings. And the challenge is all the more pertinent given the emergence of recent evidence showing that some anomalies apparently decay in some repeated decision settings with real incentives. For example, in some studies that allowed subjects to acquire experience through repetition with feedback and incentives, WTP and WTA converged (Coursey et al. 1987; Shogren et al. 1994, 2001; List and Shogren 1999) and preference reversal subsided (Cox and Grether 1996; Braga and Starmer 2003). We shall discuss this evidence in more detail below.

An interesting interpretation of decaying anomalies has been proposed by Plott (1996): the *Discovered Preference Hypothesis (DPH)*. Roughly stated, this assumes that when people face new decisions, in unfamiliar environments, they may be unclear about what it is in their interest to do. But, if agents make repeated decisions, and if the decision environment provides feedback on the consequences of actions and if those consequences are significant enough for agents to attend to, then we may expect behaviour to evolve. Plott characterises this as a three-stage process with a particular trajectory: decisions progressively exhibit less randomness and greater rationality. To quote Plott at some length:

“The theory is that rational choices evolve through three stages reflecting experience and practice. Stage one occurs when experience is absent. Untutored choices reflect a type of myopia. The individual is purposeful and optimizing, but exhibits little awareness about the immediate environment or the possible longer run consequences of any act that might be taken. Responses are ‘instantaneous’ or ‘impulsive’, reflecting whatever may have been perceived as in self-interest at the instant. To an ‘outsider’, such behavior could appear to have a substantial random component because inconsistencies among choice may be present. Systematic aspects of choices might exist, reflecting attention and perceptions, but they might not make sense when viewed from the perspective of a preference based model. Stage two is approached as repeated choices, practice, incentives (feedback), provide sobering and refocusing experiences. Problems of the type found in the first class of example are no longer present in the data. Choices begin to reflect and incorporate an awareness of the environment, and can be recognised by an ‘outsider’ as a stable form of ‘strategy’ or ‘decision’. The full constancy of the rational model begins to find support in the data. However, problems of the sort con-

tained in the second class of examples can still be detected. Stage three, the final stage, is one in which choices begin to anticipate the rationality reflected in the choices of others. The fact that others might be acting rationally, and the consequences of that rationality, as it works through the interdependent fabric of social institutions, become reflected in the choices of each agent.” (Plott 1996, pp. 226–227).

This offers a new and interesting interpretation of anomalies and one which, we shall argue, has some interesting implications in relation to methods of preference elicitation. As a prelude to this discussion, it will be useful to highlight what we take to be three key aspects of DPH.⁴

1. *Stable and context free preferences*: DPH assumes that agents have stable and context free preferences that exist independently of the discovery process. This is a key feature demarcating Plott’s DPH from the *constructed preference* view, of which he is an open critic. The latter perspective rejects the usual presumption of economic theory that behaviour can be understood in terms of a unique set of preferences which are stable and pre-exist choice. On one reading of this view, preferences are not prior to, but arise as a consequence of, decision tasks.⁵ The constructed preference view has been promoted primarily by psychologists (see for example Slovic, 1995; Kahneman, 1996) and, thus far, it has found little favour amongst economists, no doubt partly because adopting it is antithetical to conventional economic methodology that builds from preference assumptions. Correspondingly, the constructed preference view also has radical implications at the more practical level of public policy. Much of the thrust of conventional welfare economics is directed towards enhancing preference satisfaction, and specific methods such as CV play a functionary role in this objective by seeking to quantify preferences. This general strategy, however, looks problematic if preferences lack stability or, worst still, incoherent if they don’t exist at all. Plott (op cit., p. 227) makes the point crisply:

“Of course, if no preferences exist, then there is no foundation for a theory of optimization and no foundation for a theory of strategic behaviour and game theory. The idea of constructed preferences would seem to leave little room for economics...”

2. *Standard preferences*: It will be part of our interpretation of DPH that underlying preferences are of a standard (i.e., anomaly free) kind. We think this is clearly in the spirit of Plott’s hypothesis. It is hinted at in phrases such as “...the final product of the process may be a preference-like object that is very familiar to economic theory” p. 227. And, it is characteristic of the evidence he reviews that behaviour becomes more consistent with standard theory.

3. *The Nature of Discovery*: Given 1 and 2, it is natural to interpret anomalies as (possibly systematic) *errors* relative to underlying preferences.

On this reading, discovery is a process of learning to avoid errors and preferences will be more accurately revealed in environments that foster error reduction by allowing subjects to learn from experiencing the consequences of their decisions. This requires task repetition, feedback on the consequences of decisions, and that those consequences be real.⁶ For the subsequent discussion, it will be useful to distinguish between two types of learning postulated in DPH. One type is learning about how best to achieve well-defined goals in a new decision environment: we will call this “*institutional learning*”. For example, in the context of a repeated auction (such as a second price auction) institutional learning would include learning about the rules of the auction, learning about how other people bid in the auction and, learning how to bid optimally, given your own preferences, the auction rules and the behaviour of others. Much of Plott’s discussion explores different types of institutional learning. But beyond this Plott explicitly suggests that discovery may involve learning at a more basic level: that is, *discovering features of one’s own preferences*, such as one’s attitude to risk.⁷ We will call such processes of learning about one’s own desires “*value learning*”.

The DPH provides a novel interpretation of preference anomalies with some rather distinctive implications. At the theoretical level, in contrast to those who interpret anomalies as reasons for the development of descriptively more accurate preference models, DPH provides a possible defence for standard theory. If apparently systematic deviations from standard theory turn out, on closer inspection, to be merely errors that disappear with experience, the case for non-standard preference modelling may evaporate. So, at the applied level, does the case for alternative elicitation methods based on such non-standard models.

The practical implications of the DPH for preference elicitation methodologies such as CV are, however, less clear. Assume that, as DPH suggests, valuation data will be more accurate when elicited in the context of some repeated market institution where agents make real decisions with real consequences. It is no accident that the goods traded in experimental markets are usually low value *private goods*, such as money vouchers, mugs, chocolate bars and the like. The rights to these goods can be easily transferred from one experimental subject to another, and the goods themselves can be physically passed to the agents for consumption. In contrast, many of the goods which we have interest in valuing for policy purposes are *public goods* such as changes in air quality, noise or risk. The rights to these goods, because of their very public nature, cannot typically be assigned to individuals and, therefore, cannot be traded. So repeated markets cannot be used to elicit anomaly-free data when it comes to public goods.

A way around the non-tradability problem can possibly be found in cases where it is possible to simulate the effects of a public good on an individual. For instance, one could in principle expose experimental subjects to differing

private levels of, say, air quality, noise or risk in a laboratory with a view to observing their private valuations. However, this approach, even when feasible, suffers from severe limitations. One is that a subject's valuation of the effect of a public good, say, some air quality improvement, during the course of an experiment may not be a reliable indicator of her valuation of the same effect over a year or her lifetime. Other limitations are ethical. For instance, while public policies will often impact upon the probabilities of life or death associated with particular activities (e.g. public transport), ethical considerations will preclude scope for placing experimental volunteers in life threatening situations.

Even though public goods cannot typically be traded in markets, given the truth of DPH, there may still be some useful lessons to draw from the experimental literature relating to it. In the worst case scenario we could conceivably be driven to conclude that, in some circumstances of interest, there simply is no way of eliciting reliable preference data when direct observation of experienced decisions is not possible. There are, however, some much more positive possibilities. For example, to the extent that the anomalies present in inexperienced decisions are "similar" comparing across private and public goods, the evidence from repeated experimental markets for private goods may provide lessons for the elicitation of preferences on public goods. If inexperienced decisions are typically prone to predictable biases and if experience in repeated markets makes biases disappear, observation of behaviour in those markets may reveal the nature and magnitude of the biases. This could allow us to calibrate one-shot valuations for goods that cannot by their nature be valued in a repeated experimental market. There might also be opportunities to use the methodology of repeated market elicitation to compare the reliability of competing elicitation methods that can be applied elsewhere. Such results could be useful in obtaining more reliable preference data in contexts where repetition of real decisions with incentives is not practical, as in the case of public goods.

Of course, before thinking there are lessons to draw from DPH, one should evaluate empirical support for it. With this in mind, in the next section, we explore some of the evidence relating to DPH. We suggest that there is some, albeit qualified, support for the DPH. We then move on to ask what mechanisms might lie behind the erosion of anomalies when and where this does appear to occur.

4. Is There Evidence for the DPH?

We begin by examining evidence relating to the impacts of experience on conformity with EUT in lottery choice experiments. There is some evidence that experience in the form of just repeating simple choice tasks may sometimes produce greater conformity between observed behaviour and the

predictions of EUT. For example, several studies show a tendency for behaviour to converge towards expected value maximisation (i.e. a special case of EUT with linear utility) when subjects repeatedly face Allais type binary choice problems (see for example, Keren and Wagenaar 1987; Barron and Erev 2003). To read this as evidence for DPH, we must assume that underlying preferences are risk neutral. There may of course be other interpretations, for instance, that subjects adopt expected value maximisation not because this reflects their true underlying preference but because it is a simple rule to follow in a repeated choice context. We discuss this possibility further in the next section.

While currently this is harder to find, there is at least some evidence of experience promoting behaviour consistent with non-risk neutral EUT. The clearest example that we know of comes from an experiment run by Loomes and Sugden (1998) in which subjects made a single repetition of 45 pairwise lottery choices. Loomes and Sugden (1998) and Loomes et al. (2002) analyse that data set using different methods but reach this common conclusion: the choice patterns violate EUT, even after allowing for stochastic variation in decision making, but the degree of violation subsides with repetition. Loomes et al. (2002) econometrically fitted several models to the choice data and the best fit was provided by a version of rank-dependent expected utility theory which relaxes the independence axiom of EUT. Their econometric model featured two sources of stochastic variation: one source of this variation was intended to capture preference imprecision; the other was designed to capture decision errors akin to the idea of *trembles* in game theory. Their data set allowed them to examine the impact of experience and they report two distinct effects. One effect of experience was that the tremble component of stochastic variation became insignificant; the second was a tendency for preferences to become *more* consistent with EUT.

On one interpretation, this evidence from binary choice studies provides some support for DPH: experience in the form of task repetition reduces noise and fosters greater conformity with the standard model of risk preference. The evidence for that claim, however, is so far rather modest and not all of it points in the same direction. Some studies dating back as far as the 1970s show that opportunities for reflection about choices can lead to *increases* in Allais type behaviour (See Slovic and Tversky 1974, and more recently Bone et al. 1999).

There is perhaps stronger support for DPH in evidence relating to the WTA/WTP disparity. Coursey et al. (1987) provided the earliest evidence of the subsidence of the WTA/WTP disparity. Following the first experimental observation of the WTA/WTP disparity by Knetsch and Sinden (1984), Coursey et al. investigated its robustness in a repeated market setting. In their experiment subjects repeatedly stated WTP to avoid, or WTA to endure tasting an unpleasant substance. Valuations were elicited using a Vickrey

auction mechanism. In an n th price Vickrey auction, with $1 < n \leq$ number of bidders, the $n-1$ agents submitting the highest (lowest) bids buy (sell) at a price equal to the n th highest (lowest) bid. A key feature of the Vickrey mechanism is that “sincere bidding” – i.e., bidding one’s true WTP or WTA – is a weakly dominant strategy. Coursey et al. employed the, now widely used, second-price variant of the Vickrey auction,⁸ in which $n = 2$, and their main finding was that mean values of WTP and WTA converged with repetition of the auction. While this suggests that the WTA/WTP disparity is eroded by market experience, a subsequent study by Kahneman et al. (1990) reports a persistent WTA/WTP disparity in a repeated quasi-market environment. In their experiment, valuations were elicited using the Becker-DeGroot-Marschak (BDM) incentive mechanism. The contrast between these two sets of results might seem surprising given that the BDM mechanism is also a variant of the Vickrey auction mechanism,⁹ rendering truthful revelation an optimal strategy. But subsequent research has replicated the finding that while a WTA/WTP disparity may be robust to repeated experience of the BDM mechanism, it tends to be eroded as a consequence of exposure to some market mechanisms, including the second-price variant of the Vickrey auction (see for example, List and Shogren 1999; Shogren et al. 1994,¹⁰ 2001¹¹).

Strikingly similar findings have been reported by Cox and Grether (1996) in relation to preference reversal. They elicited WTA valuations for simple gambles using several different mechanisms including the BDM procedure and the second price auction. We focus on their results across these two institutions. Subjects in their experiment were exposed to just one of the mechanisms. Values for P and \$ bets were elicited five times each, followed by a straight choice between them. Subjects also had immediate feedback about the consequences of their actions: if they sold a bet, they were paid the appropriate price; if they held a bet at the end of a round, they played it out. The data from the group exposed to the BDM mechanism revealed the typical asymmetric pattern of reversals when choices were compared with initial or with final valuations. The data from their second price auction showed the standard asymmetric pattern of preference reversal – many people chose P and valued \$ more highly, whereas very few did the opposite – present when comparing initial valuations with choice. But, by round five, although reversals were still quite common, the asymmetric pattern was no longer present. These symmetrical inconsistencies might be interpreted as random deviations from (discovered) preferences of a conventional kind. This same tendency towards erosion of the systematic component of preference reversal following repeated exposure to a second-price auction has been observed in a subsequent experiment of our own (Braga and Starmer 2003).

The above evidence provides support for the standard preference model, if we assume, in the spirit of DPH, that behaviour after repetition better reflects

subjects' preferences. This should not be taken for granted, though. Indeed, in the next section we will explore some reasons for thinking that some forms of experience could possibly be sources of bias. Nevertheless, there is at least some evidence consistent with the DPH view that experience is de-biasing.

More evidence pertaining to DPH is provided by experimental markets using *induced-value* techniques. In an induced-value market, the experimenter seeks to "control" a subject's true valuation of a virtual good by guaranteeing to redeem it at a fixed price, known in advance to the subject. One of the major reasons for using the technique is that such control over preferences facilitates prediction of equilibrium market outcomes, and one can then examine whether particular variables tend towards theoretical expectations. The technique can also be used for another purpose of special interest to us here. Because, in an induced-value experiment, we know what each subject's true value ought to be, we can see whether and how values stated within different elicitation mechanisms, such as Vickrey auctions and the BDM mechanism, deviate from true values. Several studies have used this approach to investigate the value revealing properties of Vickrey auctions and the BDM Mechanism, and the evidence is somewhat mixed.

One of the most recent studies (Noussair et al. 2004) provides the clearest support for DPH. The authors elicit WTP with the BDM mechanism and the second-price Vickrey mechanism. They report that both auctions have a substantial bias towards underbidding in early periods, but with repetition bids converge to true values. As in the studies reviewed above, the Vickrey mechanism emerges as a significantly more effective elicitation device: initial bids are closer to values and convergence between bids and values is faster and goes further in the Vickrey auction. In the last round of the Vickrey auction, nearly 90% of subjects bid within 10% of their induced values and 77% of bids were within 2% of value, whereas in the BDM treatment the corresponding figures were about 65% and 41% only. In other induced-value studies the most obvious common finding was that deviations from true values were small from the first round. Some of these studies observed some convergence towards true values (Coppinger et al. 1980, Cox et al. 1982; Irwin et al. 1998) whereas others did not (Kagel et al. 1987; Kagel and Levin 1993). These two last studies mean that the induced-value literature does not provide unequivocal support for DPH, but it does not plainly refute it either, and as we will see in the next section the persistence of small deviations between stated and true values is compatible with a qualified version of DPH.

To conclude, not all evidence surveyed in this section is fully consistent with DPH, but some support can be found for it: in some circumstances behaviour initially deviates from, but with experience converges to, the predictions of standard theory. Interestingly, some elicitation mechanisms are more effective than others in promoting that convergence. That leads to interesting questions about what drives such changes in behaviour and why

some mechanisms seem to be more effective than others? We suspect that multiple factors may be involved in the evolution of stated preferences with experience. In the next section, we explore some possible contributory factors that may be significant elements in a full understanding of these processes.

5. Some Possible Mechanisms Promoting the Decay of Anomalies

In this section, we present three types of mechanisms that might plausibly promote the decay of anomalies and for which support can be found in the existing literature: institutional learning, value learning and the adoption of heuristics. Learning mechanisms are in the spirit of DPH; heuristics are not, and while they may in some circumstances coincidentally promote convergence to the predictions of standard theory, in other circumstances they may do the opposite. These mechanisms are not mutually exclusive and in some situations more than one may be at work simultaneously. From a practical point of view, the three types of mechanism have different implications for preference elicitation and we will tentatively explore this issue in Section 6.

5.1. INSTITUTIONAL LEARNING

At one level the notion of institutional learning is easy to grasp. We will all surely recognise from daily life that unfamiliarity with an environment makes good decision making more difficult. For example, a resident of the UK may find that buying a house in France is a more difficult task than buying one in the UK at least partly because the rules of the “French game” will be less well known to the newcomer. So we might reasonably expect both a tendency for agents to make mistakes in unfamiliar decision environments and a tendency for decision competence to grow – and the incidence of decision errors to fall – as familiarity with the decision environment grows. Since some of the institutions used in the lab for preference elicitation, such as the Vickrey auction and BDM procedure, are fairly complex, it would be no great surprise if subjects placed in them imperfectly understand those institutions – at least initially – and consequently mis-predict the consequences of their decisions within them. Thus, their decisions could deviate from standard theoretical predictions, even if their preferences conform to standard theory.

But whilst unfamiliarity provides a basis for expecting “noise” to be present in complex institutions, the anomaly literature we have reviewed reveals more structured phenomena, involving first the presence of *systematic* anomalies in behaviour, and second the erosion of those anomalies apparently as a consequence of experience within particular institutions. If DPH is to ultimately provide a satisfactory account of such patterns in terms of institutional learning it will need to specify both how unfamiliarity gives rise to *systematic* biases and the *learning processes* involved in their erosion. This

strikes us as a rich and important, yet relatively underdeveloped, research topic. Nevertheless, we can begin to see how these requirements could be met and in what follows we offer some illustrations relating to the phenomena of preference reversal and the WTA/WTP disparity. Specifically, following a suggestion voiced in Bateman et al. (2004) we argue that these particular anomalies might arise (at least in part) as a consequence of decision errors caused by the operation of a decision heuristic, which gives rise to sub-optimal behaviour in some contexts (the “caution heuristic”). We then discuss a possible mechanism, which might be involved in promoting the erosion of sub-optimal decision rules (the “market discipline hypothesis”).

A recent experiment reported by Bateman et al. (2005) revealed the familiar WTA/WTP disparity. We saw in Section 2 that the WTA/WTP disparity could be caused by loss aversion. While Bateman et al. argue that loss aversion may be a contributory element their analysis suggests that other factors are probably involved too. This conclusion is prompted by the fact that they use different measures for loss aversion – aside from comparing WTA and WTP – which suggest the presence of only weak loss aversion. In light of this, the authors suggest another possible cause for the familiar WTA/WTP disparity: the operation of a “*cautious heuristic*”. This heuristic is an assumed tendency of individuals to overstate any true *incoming valuation* (where the subject states the minimum amount of a good he is to receive, e.g. WTA) and to understate any true *outgoing valuation* (where a subject states the largest amount of a good he is willing to forgo, e.g. WTP). This heuristic may be well adapted to various real bargaining environments that ordinary people will be familiar with (consider for instance the case of bargaining over buying and selling of property where some degree of haggling is common). However, in institutions that provide incentives for true value revelation, as was the case in Bateman et al., this heuristic is disadvantageous. But if subjects fail to understand the incentives, they may quite understandably fall back on the heuristics they are familiar with.

The cautious heuristic is consistent with the data observed by Bateman et al. (1997, 2004). It is also consistent with much of the evidence reviewed in Section 4 for the presence and subsequent erosion of the WTA/WTP disparity in repeated markets. In those studies, the erosion of the WTA/WTP disparity usually involves a reduction of stated WTA and, sometimes, an increase in stated WTP. If we assume, in the spirit of DPH, that behaviour in the last repetitions of auctions better reflect true values, then those changes in stated values would mean that in the early repetitions of the auctions WTA was being overstated, and WTP was being understated. Such initial misstatements of and subsequent change in values are consistent with the idea that subjects initially misunderstand the incentives provided by the auction and rely on the cautious heuristic, but, with repetition and feedback, learn to adopt a more suitable or refined bidding strategy.

This same conjecture may also account for the changes observed in preference reversal patterns. In two studies cited above where preference reversal was seen to decline in a repeated market (Cox and Grether 1996; Braga and Starmer 2003), valuations were elicited as WTA. Preference reversal declined in these studies because WTA for the \$-bet fell markedly while stated WTA for the P-bet remained broadly unchanged. So, the systematic tendency for subjects to value \$ above P while choosing P over \$ was driven away. If the cautious heuristic and its gradual replacement by a more refined valuation strategy is what drives the behaviour of \$ valuations, one must ask why the same does not happen with the P-bet? A plausible explanation is that the P-bets used in these experiments were nearly certain amounts of money (binary gambles with probabilities of 97% or 81% to win otherwise nothing or a very small loss). There is, therefore, little room for a “sensible” P-valuation to vary, and most observed P-valuations, in any round, are just under the amount to win and very close to the expected value of the lottery. In contrast, \$-bets have much lower winning probabilities (19% and 31% in those studies). There is, therefore, more scope for variation in “reasonable” valuations.

Some induced-value experiments also provide support for the idea that subjects initially relied on the cautious heuristic and later adopted a better valuation strategy. That is the case of Noussair et al. (2003), the only study that observed large deviations between stated and true values, and also of Coppinger et al. (1980) and Cox et al. (1982). All of these studies found initial systematic underbidding that subsided with repetition in second price auctions. Deviations between stated and true values observed in other studies are not consistent with the cautious heuristic but are also relatively mild: Kagel et al. (1987) and Kagel and Levin (1993) found overbidding in the second-price auction, while Irwin et al. (1998) found deviations from true values in both directions in buying and selling tasks conducted with the BDM mechanism.

These induced-value experiments provide an interesting insight into the nature of any learning that might be taking place. We showed that misunderstanding of the auctions and institutional learning is consistent with empirical findings, but did not show that this is the only possible interpretation of the evidence. While we find it plausible that several factors may be behind the observed behaviour, if any learning is taking place in the induced-value experiments, it must be institutional rather than value learning: when we observe agents' valuations for induced value vouchers gravitate towards their money face values, it seems most reasonable that, whatever agents have learned, it reflects some better grasp of the strategic environment and what it is optimal for them to do (given a desire for money) rather than any change in their attitudes to money. Additional support for this idea comes from the comparison of the results from induced-value second-price and English auctions. In an English auction the price keeps increasing until the last bidder

but one withdraws from the auction; at that stage the price stops and the remaining bidder buys the good at that price. As in the second-price auction, the optimal bid in the English auction – that is, the price at which it is optimal for the bidder to withdraw from the auction – is the agent's true value.¹² Yet, in induced-value English auctions (Coppinger et al. 1980; Kagel et al. 1987; McCabe et al. 1990; van Huyck et al. 1993) convergence of bids to true values is much faster than in second-price auctions.¹³ It is unlikely that the English and second-price auctions would differ in their ability to foster value learning, even if there were anything to learn about the monetary value of a money voucher. A more plausible explanation for the difference in performance between the two institutions is that the consequences of one's actions are far more transparent in the English auction, and that makes it easier to find the optimal bidding strategy. For instance, the unsuitability of the cautious heuristic may pass unnoticed in a second-price auction, but it will be obvious in an English auction, as underbidding would require the subject to choose not to buy when the price is still below his private value.

In those environments where there does appear to be a tendency for agents to move away from reliance on (biased) decision rules and towards more optimal decisions, a pertinent question that arises is what drives such learning? One might argue that this question need not be answered; so long as we know that learning moves in the direction of the predictions of standard theory that will suffice. That answer, however, is not entirely satisfactory because as we have seen above, institutions appear to differ in their propensities to promote learning (recall the significantly better performance of the second price auction, relative to the BDM mechanism in revealing induced values). So if we are to understand such differences it will be necessary to explore the learning process. We believe that some progress is being made in this direction.

One example is provided by Loomes et al. (2003), who suggest a simple learning hypothesis that is consistent with much of the evidence we have discussed. The authors propose a model of bid adjustment – they call it the “market discipline hypothesis” – in which subjects respond, in a systematic way, to *costly mistakes*. On this hypothesis, a subject has made a “costly mistake” if they have deviated from their optimal strategy *and* received a lower payoff (ex post) than they would have, had they played their (ex ante) optimal strategy. Institutional learning would then result from subjects learning, through some combination of reflection and trial and error, to avoid costly mistakes. In the context of Vickrey auctions, even though bidding true value is a weakly dominant strategy, not all deviations of bid from value will be costly mistakes. In fact, a misstatement of value is costly, in the above sense, only if the market price (including a randomly-generated BDM price) falls between a bidder's true and stated values. To illustrate this, take the case of stated WTA in an auction.¹⁴ Suppose an agent overasks (i.e. $\text{ask} > \text{true value}$). Whether or not this deviation from true

bidding is a costly mistake will depend on the resulting market price. If the price falls between the ask and value (i.e. $\text{ask} > \text{price} > \text{value}$) then the agent has made a costly mistake because they will fail to sell at an attractive price. On the other hand if price exceeds the ask, the deviation is not costly – the agent sells, they would wish to sell at this price and bidding true value would not have affected their payoff. Now suppose the agent overasks and the price is less than the value. Again this is not a costly deviation from sincere bidding. The agent does not sell, but they would not want to sell at this price and truthful bidding would not have changed their payoff.

Loomes et al. (2003) attempted to directly test the market discipline hypothesis exploiting two implications of it. One is that, in Vickrey auctions, there can be agents stating bids which deviate by an arbitrary amount from their true values, but with no incentive to adjust their bids because their deviations are not costly to them (e.g. agents who over-ask in selling auctions where the price turns out to be less than their true value). On the other hand, the market discipline hypothesis puts quite tight bounds on the bids of marginal traders: in equilibrium, marginal traders must be bidding “close” to their true values. So, to the extent that the WTA/WTP disparity is caused by misstatements of true valuations, there should be a stronger tendency for convergence of WTP and WTA in bids of marginal traders than there is in average bids. Loomes et al. test this prediction in the context of repeated median price auctions and find it supported.

The market discipline hypothesis offers a plausible explanation of why anomalies subsided more slowly in the BDM mechanism than in the second price auction. Although both institutions make it a weakly dominant strategy to reveal one’s true values, deviations from true values have probably been more costly in the second price auction than in BDM procedure. In typical implementations of BDM, prices have been drawn from uniform distributions with rather high upper supports, with the consequence that BDM prices have likely exceeded asks most of the time. For instance, Shogren et al. (2001) generated prices for coffee mugs from a uniform distribution ranging from US\$0 to US\$9.50. Given that the mean ask was never higher than US\$3.68, the BDM price would have been higher than most asks most of the time. When that happens, subjects bear no cost for overstating true WTA because they sell at a price attractive to them. In contrast, in Vickrey auctions, a generalised tendency to substantially overask will ensure that the market price – one of the asks – will exceed most private values unless they are highly dispersed. A similar analysis suggests that understating WTP in typical implementations of BDM is less likely to result in costly errors than in Vickrey auctions. This idea is supported by Noussair et al. (2004). They explicitly calculate the expected cost of deviating from true revelation in their experiment – as a percentage of expected earnings at the optimum – and find that the cost is always higher in the Vickrey auction than in the BDM mechanism.

The market discipline hypothesis also provides a possible explanation for the persistence of mild deviations between stated and true values observed in some induced-value second-price auctions. Take the case of generalised underbidding in a second-price buying auction. Typically, induced private values vary considerably across subjects. So, if underbidding is mild, it is possible that the second highest bid exceeds all but the two highest private values. So, potentially, only the second highest bidder will bear a cost. If there is generalised mild overasking in a selling auction, it is possible that only in the second lowest asker will bear a cost.

Additional support for the market discipline hypothesis comes from our work on preference reversal (Braga and Starmer 2003). Like Cox and Grether, we observed that \$-asks decline during repetitions of the second-price auction, and the asymmetrical, standard preference reversal appeared in early auction rounds but had disappeared by the fifth. We extended the Cox and Grether study by also eliciting WTA in a second-to-last-price auction, where all but the two highest askers sell their lotteries. In such an auction, the market price exceeds most subjects' asks as long as there are more than three participants (there were seven or eight in our case). Thus, overstatement of one's true WTA will be costly less often than in a second-price auction. In line with the implications of the market discipline hypothesis, \$-asks declined only gently in the second-to-last price auction, and preference reversal persisted.

The market discipline hypothesis is just one very simple theory providing a possible explanatory mechanism for the operation of the DPH and much more research is warranted to properly understand the operation of such processes. Nevertheless, we think there is at least a tentative basis for the claim that mechanisms involving responses to costly errors may form part of the explanation for some of the behaviours that appear consistent with DPH.

5.2. VALUE LEARNING

Having made the case for institutional learning, we turn to some experimental evidence that suggests agents' *consumption experience* may play a role in the evolution of stated preference. Such evidence, we tentatively suggest, signals the possibility of something closer to value learning than to institutional learning, and that possibility has some distinctive implications in relation to preference elicitation methodology.

One example of what we have in mind as evidence of an effect attributable to "consumption experience" is provided in a recent study by Kuilen and Wakker (2004). In their experiment, two groups of subjects made repeated decisions in Allais choice problems. Both groups knew that their payoff from the experiment would be the outcome of their choices in one of

the decision tasks and which task would be for real was to be selected randomly at the end of the experiment. The difference between the treatments was that one group saw the chosen lotteries played out following each decision (ongoing outcome feedback) while the other group saw only the outcome of the lottery that turned out to be for real at the end (terminal outcome feedback). The researchers report Allais type deviations from EUT for both groups in early decisions but the effects of task repetition differed between the groups. Allais deviations were eroded in the group that had ongoing outcome feedback, but persisted in the group that had only terminal outcome feedback. The authors conjecture that Allais deviations may reflect biases in the way that subjects respond to stated probability information and that observing resolutions of the chance events allows subjects to get a better feel for the meaning of stated probabilities thereby eroding the initial bias.

One aspect of the results from Braga and Starmer (2003) confirms that feedback on the resolution of risks matters, but also suggests that the mechanisms involved may go beyond subjects developing a feel for the meaning of probability information. In common with other studies, we observed a tendency for average WTA of \$-type bets to fall significantly across the rounds of repeated Vickrey auctions. But a striking characteristic of our data is that individual WTA tends to fall most immediately after a round in which a given subject has played a lottery and *lost*. A feature of our design was that those who had sold their lotteries also observed what would have happened had they kept and played them. Hence, in this case, we can infer that the impact of losing upon valuations is not just driven by probability learning because seeing other people losing does not produce a similar effect: the *personal* experience of actually losing appears to play a key role. A possible interpretation is that it reflects subjects learning about particular affective states: specifically, what it *feels like* to experience losses and through that process they reflect on, and adjust, their auction bids.

The idea that consumption experience might change or refine an agent's perception of their own preferences may seem more plausible when set against the backdrop of emerging evidence indicating that individuals may be prone to systematic biases in anticipating their tastes (see for example, Kahneman and Snell 1990; Loewenstein and Adler 1995; Loewenstein et al. 2003). Some of the evidence relates to riskless choice. For example in Loewenstein and Adler's "Experiment 1" some subjects were asked to predict their selling prices for a mug, prior to being endowed with it, and then asked to reveal actual selling prices after the endowment had occurred. Other subjects were simply endowed with the mug and required to reveal selling prices. The predicted selling prices were significantly lower than actual selling prices and Loewenstein and Adler interpret this as evidence that there is an endowment effect but that individuals fail to anticipate it. There is parallel

evidence suggesting that individuals may systematically mis-predict their risk attitudes in dynamic decision contexts. Specifically, several studies of individual behaviour in decision trees show a tendency for agents to display less risk aversion when they are required to precommit to an action for a future node in the tree than when they decide only after having reached that node (see the review of the evidence in Cubitt et al. 2004). One interpretation of this is that, in precommitting to decisions, agents systematically over-estimate their own willingness to take risks further along the tree. Cubitt and Sugden (2001) find that the same effect is not present when subjects have experience of practicing the decision before making their real choices. A tentative interpretation of this is that, through experience, subjects learn how to better predict their willingness to take future risks.

We should note that, in principle, some of the above findings which we have used to motivate the case for value learning might also be explained as forms of institutional learning. For instance, the effects of observing risk resolutions in Kuilen and Wakker's study could be a form of institutional learning if, for instance, agents initially distrust stated probability information provided by the experimenter but adjust their priors towards stated probabilities as they observe evidence consistent with them. Similarly, the impact of loss experience observed by Braga and Starmer would be consistent with subjects learning about the elicitation mechanism given the added assumption – in the spirit of the market discipline hypothesis – that it takes a loss to prompt subjects to reflect on the workings of the auction and their bidding strategies. A subject who sold their lottery at an attractive price, might feel no urge to revise her asks, but subjects who failed to sell and lost while playing the lottery might feel inclined to ponder the merits of lower asks in the following round.

In Section 6, we suggest that value and institutional learning may have some quite different implications for elicitation methodology. As yet, the existing evidence is in many cases not clear enough to allow discrimination between them as accounts of the evolution of stated preference. Given this, we confine ourselves here to the more modest objective of attempting to identify alternative hypotheses, suggested by the data, that may be worthy of further research and scrutiny.

5.3. REPEATED EXPERIENCE AND CHOICE HEURISTICS

It is possible that some systematic changes of stated preference resulting from experience may reflect neither institutional nor value learning and, in this section, we consider another possible factor: heuristics. There is a large literature, much of it emanating from psychology, discussing the role of simple choice heuristics in various types of decision making. One branch of this literature has demonstrated that a range of decision biases might be caused

by agents using various simplifying decision heuristics or choice rules (see for example, the collection of articles in Gilovich et al. 2002). Another branch has emphasised how simple decision rules can be well adapted for effective decision making, even in rather complex environments (Gigerenzer et al. 1999).

Choice heuristics provide possible explanations for the presence of some anomalies in one-shot decisions, including at least some of those discussed in Section 2. For example, to the extent that preference reversal reflects over-valuation of \$-bets, it might be explained by a heuristic such as that posited under the “anchoring and adjustment hypothesis” (Slovic et al., 1982). Consider an agent evaluating a bet which gives a p chance of winning prize x . On this hypothesis, agents initially “anchor” their valuation on the winning amount x , they adjust downwards to take account of the risk, but over-valuation follows from the assumption that adjustment away from the anchor is in some sense “insufficient”.

In this section, we are mostly interested in heuristics that may coincidentally be promoting behaviour that appears consistent with DPH. Knetsch et al. (2001) and Loomes et al. (2003) have suggested such a heuristic and found empirical support for it. As part of the motivation for their studies, both sets of authors observe that the elimination of the WTA/WTP disparity reported by Shogren et al. (1994) could be explained on the assumption that there was some tendency for subjects to follow a simple rule of *adjusting their bids towards observed market prices*. In the Shogren et al. data there is a tendency for WTP to fall and WTA to rise in repeated second price auctions. But notice that, when subjects are buying in a second price auction, the market price is the second *highest* WTP; when they are selling, the market price is the second *lowest* WTA. So, when there are more than three bidders, price following would tend to eliminate the WTA/WTP disparity. As Loomes et al. explain, a similar argument also provides a possible account of the disappearance of preference reversal in Cox and Grether’s study. To the extent that price following drives changes in stated preferences in repeated markets, those changes do not necessarily imply movements towards true preference and the apparent convergence of behaviour towards the predictions of standard theory may be merely a coincidental product of agents adopting a naive rule. If price following is at work in the erosion of anomalies, then DPH will appear less compelling. The studies of Knetsch et al. (2001) and Loomes et al. (2003) provide evidence that price following may actually be at work.

To discriminate between price following and competing hypotheses, Knetsch et al. elicited WTP and WTA for mugs in six rounds of both a second price and a second-to-last price auctions. In the second-to-last price auction with more than three participants (there were ten in their experiment), most asks will be below the market price and most bids will be above

it. Therefore, in contrast to what happens in the second price auction, price following will tend to amplify the WTA/WTP disparity. With the second price auction the WTA/WTP disparity was small even in the first round, and, as in previous experiments, disappeared with repetition. In the second-to-last price auction, the WTA/WTP disparity was very large in the first round, and *increased* with repetition – the ratio of the mean WTA to the mean WTP jumped from 2.35 in the first round to 7.45 in the sixth round. This provides some support for the price following hypothesis. Note, however, that price following does not explain the gentle decline in asks for the \$-bet in our own second-to-last price auction (Braga and Starmer 2003).

Loomes et al. (2003) elicited WTA and WTP in median price auctions for binary lotteries. These auctions are variants of the Vickrey mechanism where the price is determined by the bid (or ask) at the middle of the bid/ask distribution. Within each auction, two subgroups of the participants were bidding for (or offering) two different lotteries. Both lotteries had the same money prize (12 UK pounds), but one gave a low chance of winning the prize while the other gave a high chance winning the prize. In some groups, a majority of participants would bid for (or offer) the attractive lottery, and the remaining subjects would bid for (or offer) the unattractive lottery; in other groups it was the reverse. Typically the market price would be set by a member of the majority and so the market price tended to be higher in the groups where the majority bid for (or held) the more attractive lottery. Loomes et al. report a significant price following effect: agents buying (or selling) the low probability bet set higher bids (asks) when they had previously observed higher market prices. The same effect did not appear to operate on the high probability bet, and there may be a parallel here with the earlier discussion of overvaluation: it may be that price following is confined to situations where there is scope for an agent to be uncertain about what their bid or ask should be.

The evidence for price following in these experiments is not necessarily incompatible with the idea that people learn to make better decisions in response to costly errors. While there is apparently no rational basis for price following in these designs, from the point of view of individual bidders, it may be harmless too. For instance, if it reflects a systematic drift in the bids of non-marginal traders it may involve few if any costly errors. Even so, price following may impart a bias to average valuations in the market. So, if our interest is in how to elicit accurate preference data, we should be alert to the possibility of this type of effect.

The possibility of price following raises issues which run well beyond those of preference elicitation techniques. Evidence of price following points in the direction of the constructed preference perspective and, as such, raises some potentially serious concerns relative to the standard perspective of welfare economics. The latter framework depends, in important respects, on the

background assumption that agents have stable preferences that are exogenous to the market. If revealed preferences themselves turn out to be endogenous to market institutions, the very notions of efficiency which have been the cornerstone of economic policy analysis may require re-examination.

6. Implications for Preference Elicitation

The research we have reviewed shows that there is indeed considerable evidence of stated preferences changing systematically in environments that provide specific kinds of experience. At least some of the evidence points to experience promoting greater consistency between behaviour and the predictions of standard preference theories. We have suggested three mechanisms that might be among the processes involved in these changes: institutional learning (in particular responding to costly errors), value learning resulting from consumption experience and the adoption of simple heuristics. In this section we seek to take stock of the arguments that have been set out with the particular aim of teasing out possible implications for preference elicitation methodology.

When it comes to private goods, one may be tempted to draw a simple lesson from the literature reviewed above: to elicit more reliable preference data, repeat decision tasks with incentives and feedback. But things are not so straightforward, as that literature also suggests that repetition and feedback may fail to reduce biases if they are not costly. As the evidence for price following shows, repetition with feedback may even induce further biases. This question was recently raised by Harrison et al. (2004). They argue that when institutions such as Vickrey auctions are used to elicit valuations, participants should take part in only one auction round to avoid biases that might be induced through affiliation of values. While we accept the possibility of bias from affiliation our analysis suggests the presence of other biases arising in mechanisms such as Vickrey auctions (e.g. the use of a cautious heuristic), which may only be eroded by exposure to repeated experience. This suggests a need to weigh these considerations in the design of elicitation procedures, or the calibration of research findings.

In this paper, however, we are mainly interested in elicitation of preferences over public goods, in which case repetition with feedback and incentives is typically not feasible. Here, to the extent that repetition with feedback and incentives improves the reliability of preference elicitation, we have mainly two positive suggestions to offer: calibration of valuation measures by correcting for what appears to be a bias in inexperienced decisions, and use of choice-based elicitation devices. Before we elaborate on these suggestions, however, there is an important issue that we will raise but not fully address here: that of incentives.

6.1. THE INCENTIVE QUESTION

A potential problem of elicitation of preferences over public goods is that, typically, it cannot be made incentive compatible. There is controversy as to the effects of incentives (see Camerer and Hogarth 1999; Hertwig and Ortmann 2001). Many economists regard data from hypothetical tasks with some reservations. On the other hand, in some cases, results seem to be unaffected by incentives. For instance, similar patterns of preference reversal have been observed in one-shot experiments both with and without incentives (see for instance Lichtenstein and Slovic 1971; Grether and Plott 1979).

Incentives are crucial in DPH and play an important role in the learning mechanisms we suggested above: institutional and value learning are enhanced if decisions have real consequences and may not occur at all if there are no real consequences. Importantly, all the experiments reviewed above were incentive compatible. So, any generalisation of findings of those experiments to public goods must be qualified by the fact that typical methods of elicitation of preferences over public goods, such as contingent valuation, cannot be made incentive compatible. Specifically, we cannot be sure that procedures that improve data quality in incentive-compatible exercises will have the same effect if there are no incentives.

We are aware that several collective decision mechanisms exist that, in theory, overcome the incentive problem: for instance, the Clarke (1971) mechanism makes it optimal for individuals to reveal their true values; other mechanisms, such as those proposed by Groves and Ledyard (1977) and Falkinger (1996), in their Nash equilibria, reveal aggregate preferences. These mechanisms, one could think, could advantageously replace contingent valuation in many situations. However, they are fairly complex and their incentive structure is far from transparent, so it is by no means guaranteed that they will work in practice. But as they have been insufficiently tested in the laboratory, it is too early to judge their effectiveness (see Chen 2005, for a review of the theoretical and experimental literature). Even if these mechanisms do work, there may be obstacles to their implementation in many situations. One is that people may have political objections to them. For instance, in a field study with college fraternities in the US (Tideman 1983), some members objected to replacing majority voting with the more efficient (in the Pareto sense) Clarke mechanism, on the grounds that it departs from the one-person-one-vote rule. We can think of similar objections to some of the other incentive-compatible mechanisms. Another obstacle to the implementation of these mechanisms is that in many situations one is not interested in immediately arriving at a collective decision, but simply in eliciting preferences over different public goods to explore different policy options. If the mechanism is applied with no real decisions being arrived at, that is, if people participate in the mechanism on a just-imagine basis then, given the

complexity of the incentive structure of those mechanisms, it is doubtful that any data obtained from the exercise will be of any meaning at all.

For these reasons, although we think that these incentive-compatible mechanisms offer interesting new possibilities and deserve further empirical testing, we feel that they are not currently viable replacements for contingent valuation. It is also noteworthy that, even though some of the mechanisms have been known for over two decades, they have not, to our knowledge, been used for actual policy purposes.¹⁵ Therefore, it makes sense to explore the experimental literature to find ways of improving the reliability of contingent valuation.

6.2. CALIBRATING VALUATION MEASURES

We suggest that DPH may provide some new perspective on the reliability of standard contingent valuation measures by consideration of what typically happens to stated WTP and WTA for goods that are elicited in situations that can be expected to foster learning. As discussed above in Section 4, there is some evidence of a tendency for the WTA/WTP disparity to close in repeated market settings with incentives and feedback. This involves both a tendency for WTA to fall and for WTP to rise, although the latter is less pronounced (see for example, List and Shogren, 1999). To the extent that post-experience measures of value in such environments are better estimates of true values, it suggests that standard approaches for eliciting valuations may underestimate WTP and overestimate WTA. That evidence suggests that a figure between elicited WTP and WTA may be a better measure of value than either WTP or WTA themselves. Alternatively, one might consider eliciting measures, such as *equivalent gain*, which tend to lie *between* WTP and WTA. This needs some qualification, however, as some induced value studies observed overstatement, although mild, of WTP.

6.3. A CASE FOR INSTITUTIONAL SIMPLICITY

In Section 5, we suggested that some biases may be caused by subjects' failure to understand complex elicitation institutions and we presented evidence consistent with that view. On this interpretation of the data, the erosion of anomalies would result from subjects learning through repetition with feedback and incentives to better implement their preferences in the institution. In the case of public goods it is impractical, if not impossible, to implement elicitation devices embodying opportunities for institutional learning. This provides an argument favouring the use of choice-based elicitation mechanisms, as they are simple and transparent and hence require less learning than other alternative elicitation devices. In fact, given that all a subject has to do in a choice task is to pick which of two alternatives she prefers, there is very

little, if anything at all, for the agent to learn about the institution itself. Choices should therefore be less prone to biases than alternative elicitation devices. This idea is consistent with evidence we have presented above. In the studies of Cox and Grether (1996) and Braga and Starmer (2003) subjects' preference rankings of goods were elicited in two ways: through choice and through pairwise comparisons of WTA valuations. The WTA valuations were elicited in repeated markets and both studies revealed a tendency for the rankings revealed in WTA to become more consistent with choices over time. On the DPH view, this convergence of valuations on choice would favour a presumption that, in the absence of experience, choices are more reliable indicators of preference than valuations.

However, there is evidence showing that institutional transparency is not a panacea for generating anomaly free data. For instance, Bateman et al. (2004) observed the WTA/WTP disparity despite eliciting values from binary choices. In their design values were obtained via a series of *multiple dichotomous choices*. For instance, to elicit WTA for 10 chocolates, a subject would be given 10 chocolates and then would make a series of choices where one option was always "You keep your 10 chocolates" and the other was "You give us your 10 chocolates and take X in exchange" with X varying from £0.30 to £7.50 in £0.30 increments. A possible explanation for the observation of the disparity in this design is that subjects interpreted the series of choices as a device to elicit a valuation and then adopted a customary heuristic for valuation tasks (e.g. the caution heuristic discussed above) leading to overstatement of true WTA and understatement of WTP.

This conjecture is consistent with the findings of Fischer et al. (1999). They looked for the choice-matching discrepancy (see section 2) across several types of matching tasks. In one treatment, the matching value was obtained via an open-ended question (straight matching); in another, via a series of dichotomous choices (choice-based matching). In the latter case, the series was not pre-set; instead, in each choice, the value was adjusted according to the subject's previous answer. For instance, in the example above, if a subject chose to keep the chocolates, X would be increased; otherwise it would be decreased. The procedure was continued until the matching value was known to be within an interval of pre-defined width. Fischer et al. found that the choice-matching discrepancy persisted when matching was choice-based, and preferences elicited with choice-based matching were much closer to those elicited with straight matching than those elicited with a simple choice. As in Bateman et al. (2005), it looks as though subjects responded to choice-based matching as if it were a valuation task. But in another treatment Fischer et al. concealed the aim of the choice-based matching tasks by interweaving several choice series and filler questions. In this condition, no choice-matching discrepancy was observed. Here

then is an instance where complete transparency of an institution is not a virtue. The good news, however, is that it appears to be possible to elicit relatively unbiased valuations via a series of simple and transparent dichotomous choices, although the aim of the whole exercise must be opaque to the subjects.

There is of course another reason why institutional transparency may not guarantee anomaly free responses: some alleged anomalies may reflect features of true (but non-standard) preferences. There is some support for this conclusion in at least some cases. For example, Cubitt et al. (2001) conduct a meta-analysis examining whether Allais violations of EUT occur in very simple choice environments. The “purest” of these is a design in which subjects make a *single* binary choice between a pair of prospects knowing that it is the *only task* and that *it is for real*. They find that Allais-type violations are present even in the purest single choice design. It seems implausible to suppose that those violations of EUT could be due to errors arising from any kind of misunderstanding of the choice institution since all the agent does is to make a choice between two options. This suggests the conclusion that non-expected utility models are indeed needed to properly model risk preferences, at least for explaining one shot decisions.

6.4. IMPLICATIONS OF VALUE LEARNING

The Cubitt et al. paper just mentioned does not set out to examine whether Allais-type anomalies are robust to repetition, particularly in the presence of outcome feedback. This is significant given the evidence we presented above suggesting that Allais anomalies might be eroded by exposure to outcome feedback. While the evidence for this is so far rather modest it does fuel interest in the hypothesis that value learning may be playing some role in the dynamics of stated preference. To the extent that value learning occurs, it suggests that agents’ reactions to researchers’ descriptions of choice alternatives may differ systematically from the evaluations they would arrive at having directly experienced those alternatives. Moreover, if there is scope for subjects to learn from consumption experience about their own tastes in relation to very simple experimental goods like binary gambles, then the same should be true in the case of much more complex entities such as environmental public goods or changes in risk to life and so forth. An implication of this is that good preference elicitation techniques may require opportunities for subjects to directly experience goods being assessed in order to facilitate appropriate learning.

Yet, for reasons discussed at the end of Section 3, there are clearly severe constraints upon an experimenter’s ability to “place” in the laboratory analogues of many, if not most, of the real public goods that practitioners typically wish to value. Nevertheless, there may be methodologies that could

help take subjects “closer” to the particular consumption experiences that may be required for value learning to take place. For instance, there may be some scope for *simulating* goods to be valued in an experimental laboratory and as an example of this strategy, a group of researchers at the University of East Anglia who use experimental methods to investigate environmental issues have recently invested in virtual reality technology with a view to creating virtual environments in a laboratory context. So, for instance, subjects asked to think about what value they place upon some beautiful landscape may be given the opportunity to experience it vicariously in cyberspace. This strikes us as an exciting new development that may allow exploration of the connections between real and simulated experiences.

6.5. DOES EXPERIENCE NECESSARILY IMPROVE PREFERENCE DATA?

In various parts of the discussion we have interpreted experimental data adopting the presumption that, when experience promotes changes in behaviour, the post-experience decision tends to be a better indicator of individual welfare. However, we have also discussed at least two types of reason why this might not always be a legitimate presumption. The first is suggested by the discussion of heuristics and, in particular, the possibility that the emergence of behaviour consistent with standard theory might be a coincidental consequence of agents following some simple rule of thumb. The second reason follows from allowing that certain kinds of influences on choice, which are biases, might nevertheless form part of an appropriate welfare benchmark if they affect agents’ perceptions of their own well being.

Consider the evidence surrounding the endowment effect. Although this effect seems highly replicable in well-defined settings, there is some evidence that it may not be robust if decisions are repeated (see Plott and Zeiler 2005). The work of List (2003) reveals parallel findings in naturally occurring market settings: in the markets that he investigates, traders are, in the aggregate, prone to an endowment effect but the size of the effect across traders depends on their individual level of market experience. And, for the most experienced traders, the endowment effect is negligible. Suppose that the reason for its decay is that experienced decision makers somehow come to recognise loss aversion as a bias in their decisions. Does this necessarily mean that we should seek to correct measures of value when they can be expected to contain a bias created by loss aversion (e.g. should we reduce, on these grounds, the WTA of an inexperienced chooser)? We believe that in some cases the answer may be “no”. Even if choosers would recognise their valuation as biased were they to become more experienced traders of the good in question, the bias may be relevant to an assessment of changes in their welfare if the bias concerned affects the agents’ own perceived sense of well-being. To ground the thought in more concrete circumstances: suppose a

householder facing a compulsory purchase order on their property is moved by feelings of loss aversion. Even if we know that an experienced real estate agent valuing the same property would not, it may be appropriate for a policy to be informed by the welfare as experienced by real agents even when their experiences contain recognised ‘bias’.

7. Conclusion

The existence of various anomalies in stated preference raises some serious questions about techniques of preference elicitation based on conventional models of preference. This paper was prompted by evidence that some preference anomalies appear to decay in some repeated markets and Plott’s Discovered Preference Hypothesis interpretation of that fact. We believe that this hypothesis sheds some interesting light on issues of preference elicitation which are relevant to environmental economics and public policy more generally.

One message of the paper is to reject a naïve reading of DPH: the evidence we have discussed does not support the claim that all anomalies of preference can be interpreted as errors relative to standard theory. Consequently, in our judgement DPH does not provide a blanket defence of standard preference theory. New models of preference are needed to organise the observed data and we believe progress is being made in this regard.

That said, we believe that the DPH interpretation offers important insights into the nature and dynamics of stated preference particularly through its emphasis on the roles of error and learning. We believe that, given the state of current research, it is too early to attempt to distil very crisp and simple messages for preference elicitation methodology. But, we hope that our discussion has successfully highlighted issues worthy of further consideration and investigation.

Acknowledgements

We are grateful for financial support from the Leverhulme Trust (award no. F/00204/K) and for discussions with participants at the workshop “Coping with preference anomalies in environmental decision-making”, University of East Anglia, September 2002. We are grateful for the assistance of the Danish Home Store. We would also like to thank Ian Bateman and Robert Sugden for very helpful comments on an earlier draft.

Notes

1. See Starmer (2000) for a review of the evidence relating to decision under risk or Camerer (1995) for a more general discussion of individual choice.
2. Where a prospect is to be understood as a probability distribution over outcomes.

3. A “large” WTA/WTP disparity could be consistent with standard theory when the good being valued has no close substitutes and/or there are large income effects. See Hanemann (1991).
4. While our discussion of DPH here is intended to distil what we take to be key elements of Plott’s DPH, our account simplifies a more complex and subtle discussion in Plott. As such, we do not attribute this interpretation of DPH directly to Plott: instead, we present it as a working hypothesis motivated by Plott’s discussion.
5. An alternative version might allow prior existence, but deny uniqueness of preferences.
6. “repeated choices, practice, incentives (feedback), provide sobering and refocusing experiences.” (Plott 1996, p. 227)
7. “The hypothesis suggests that attitudes like expectations, beliefs, risk-aversion and the like, are *discovered*, as are other elements of the environment. People acquire an understanding of what they want through a process of reflection and practice.” (emphasis in original, p. 227).
8. They used a Vickrey auction but modified with a veto rule so that, strictly speaking, their mechanism was not incentive compatible with sincere bidding.
9. When the BDM mechanism is used to elicit WTP, a single agent submits a valuation v knowing that a random device will generate a bid r from a given distribution. If $r \leq v$, the agent buys at price $= r$. If $r > v$, there is no trading. The procedure works in an analogous way for elicitation of WTA. Hence, from the bidder’s point of view, the BDM mechanism is equivalent to a Vickrey auction with one other participant who bids at random.
10. Shogren et al. (1994), examined the WTA/WTP disparity in repeated second-price Vickrey auctions comparing results for two types of goods: some with close substitutes (mugs or candy); and another with no close substitutes (reduced health risks in sandwiches). They found convergence of WTP and WTA for goods with close substitutes but not for valuations of reducing the health risk. Their interpretation is that the results observed by the end of repeated trials can be interpreted as consistent with standard theory which allows a relatively large WTA/WTP disparity for goods with no close substitutes.
11. Shogren et al. (2001) examine the robustness of the WTA/WTP disparity comparing the effects of three different elicitation mechanisms: the BDM mechanism, the second price Vickrey auction; and a “random nth price” auction. The latter is a variant of the Vickrey auction in which bidders know that n is selected at random in each auction round, after bids have been submitted. They report a WTA/WTP disparity in the initial rounds for each mechanism. It persists with BDM but goes away within both second price and random nth price auctions.
12. As the price in the English auction typically increases by fixed increments each time, there may be a difference between the optimal bid and true value, but that difference is smaller than the price increment.
13. Coppinger et al., Kagel et al. and McCabe et al. report the market price only, that is, the second highest bid. But as private values were randomly attributed to subjects in each auction round, it is highly unlikely that price convergence would result from the bidding behaviour of a few subjects only. Van Huyck et al. do not publish any data at all, as their induced-value English auction was not the main focus of their study.
14. Analogous reasoning applies to stated WTP.
15. The Clarke mechanism was actually used by some college fraternities in the US, but that was part of the aforementioned field study (Tideman 1983) and the practice was later discontinued.

References

- Allais, M. (1953), 'Le Comportement de l'Homme Rationel devant le Risque, Critique des Postulats et Axiomes de l'Ecole Américaine', *Econometrica* **21**, 503–546.
- Barron, G. and I. Erev (2003), 'Small Feedback-based Decisions and Their Limited Correspondence to Description-based Decisions', *Journal of Behavioral Decision Making* **16**, 215–233.
- Bateman, I., A. Munro, B. Rhodes, C. Starmer and R. Sugden (1997), 'A Test of the Theory of Reference Dependent Preferences', *Quarterly Journal of Economics* **112**, 479–505.
- Bateman, I., D. Kahneman, A. Munro, C. Starmer and R. Sugden (2005), 'Testing Alternative Models of Loss Aversion: An Adversarial Collaboration', *Journal of Public Economics*.
- Binmore, K. (1994), *Playing Fair*, Cambridge, MA: MIT Press.
- Bleichrodt, H., J. L. Pinto and P. P. Wakker (2001), 'Making Descriptive Use of Prospect Theory to Improve the Prescriptive Use of Expected Utility', *Management Science* **47**, 1498–1514.
- Braga, J. and C. Starmer (2003), *Does Market Experience Eliminate Preference Reversal?*, Mimeo: University of Nottingham.
- Bone, J., J. Hey and J. Suckling (1999), 'Are Groups More or Less Consistent than Individuals?', *Journal of Risk and Uncertainty* **8**, 63–81.
- Camerer, C. F. (1995), 'Individual Decision Making', in J. Kagel and A. E. Roth, eds., *Handbook of Experimental Economics*. Princeton, NJ: Princeton University Press.
- Camerer, C. F. and R. Hogarth (1999), 'The Effects of Financial Incentives in Economics Experiments: A Review and Capital-Labor-Production Framework', *Journal of Risk and Uncertainty* **19**, 7–42.
- Chen, Y. (2005), 'Incentive-Compatible Mechanisms for Pure Public Goods: A Survey of Experimental Literature', in C. Plott and V. Smith, eds., *The Handbook of Experimental Economics Results*, Amsterdam: North-Holland.
- Clarke, E. (1971), 'Multipart Pricing of Public Goods', *Public Choice* **11**, 17–33.
- Coppinger, V. M., V. L. Smith and J. A. Titus (1980), 'Incentives and Behavior in English, Dutch and Sealed-Bid Auctions', *Economic Inquiry* **18**(1), 1–22.
- Cox, J. C. and D. M. Grether (1996), 'The Preference Reversal Phenomenon: Response Mode, Markets and Incentives', *Economic Theory* **7**, 381–405.
- Cox J., B. Roberson and V. Smith (1982), 'Theory and Behavior of Single Object Auctions', in V.L. Smith, ed., *Research in Experimental Economics*, Vol. 2, Greenwich, CT: JAI Press.
- Coursey, D., J. Hovis and W. Schultz (1987), 'The Disparity Between Willingness To Accept and Willingness To Pay measures of Value', *Quarterly Journal of Economics* **102**, 679–690.
- Cubitt, R. P., C. Starmer and R. Sugden (2001), 'Discovered Preferences and the Experimental Evidence of Violations of Expected Utility Theory', *Journal of Economic Methodology* **8**, 385–414.
- Cubitt, R. P., C. Starmer and R. Sugden, (2004), 'Dynamic Decisions: Some Recent Evidence from Economics and Psychology', in I. Brocas and J. D. Carrillo, eds., *The Psychology of Economic Decisions*, Vol. 2: Reasons and Choices, Oxford: Oxford University Press.
- Cubitt, R. and R. Sugden (2001), 'Dynamic Decision Making Under Uncertainty: An Experimental Investigation of Choices Among Accumulator Gambles', *Journal of Risk and Uncertainty* **22**, 103–28.
- Falkinger, J. (1996), 'Efficient Private Provision of Public Goods by Rewarding Deviations from Average', *Journal of Public Economics* **62**, 413–422.
- Fischer, G., Z. Carmon, D. Ariely and G. Zauberman (1999), 'Goal-based Construction of Preferences: Task Goals and the Prominence Effect', *Management Science* **45**, 1057–1075.
- Gigerenzer, G. and P. M. Todd, The ABC Research Group (1999), *Simple Heuristics that Make Us Smart*, New York: Oxford University Press.

- Gilovich, T., D. Griffin and D. Kahneman, eds. (2002), *Heuristics & Biases: The Psychology of Intuitive Judgment*, New York: Cambridge University Press.
- Grether, D. and C. Plott (1979), 'Economic Theory of Choice and the Preference Reversal Phenomenon', *American Economic Review* **69**, 623–638.
- Groves, T. and J. Ledyard (1977), 'Optimal Allocation of Public Goods: A Solution to the 'Free-Rider' Problem', *Econometrica* **45**, 783–809.
- Hanemann, W. M. (1991), 'Willingness to Pay and Willingness to Accept: How Much Can They Differ?', *American Economic Review* **81**, 635–647.
- Harrison, G. W., R. M. Harstad and E. E. Rutström (2004), 'Experimental Methods and Elicitation of Values', *Experimental Economics* **7**, 123–140.
- Hertwig, R. and A. Ortmann (2001), 'Experimental Practices in Economics: A methodological Challenge for Psychologists?', *Behavioral and Brain Sciences* **24**, 383–403.
- Irwin, J. R., G. H. McClelland, M. McKee, W. Schulze and N. Norden (1998), 'Payoff Dominance vs Cognitive Transparency in Decision Making', *Economic Inquiry* **36**, 272–285.
- Kagel, J., R. Harstad and D. Levin (1987), 'Information Impact and Allocation Rules in Auctions with Affiliated Private Values: A Laboratory Study', *Econometrica* **55**, 1275–1304.
- Kagel, J. and D. Levin (1993), 'Independent Private Value Auctions: Bidder Behaviour in First-, Second- and Third-Price Auctions with Varying Numbers of Bidders', *The Economic Journal* **103**, 868–879.
- Kahneman, D. (1996), Comment on Plott's Rational Individual Behavior in Markets and Social Choice Processes: The Discovered Preference Hypothesis, in K. Arrow, E. Colombatto, M. Perleman and C. Schmidt, eds., *Rational Foundations of Economic Behavior* (pp. 251–254). London: Macmillan and St. Martin's.
- Kahneman, D., J. Knetsch and R. Thaler (1990), 'Experimental Tests of the Endowment Effect and the Coase Theorem', *Journal of Political Economy* **98**, 1325–1348.
- Kahneman D. and Snell J. (1990), 'Predicting Utility', in R.M. Hogarth, eds., *Insights in Decision Making*, Chicago, IL: University of Chicago Press.
- Kahneman, D. and A. Tversky (1979), 'Prospect Theory: An Analysis of Decision under Risk', *Econometrica* **47**, 263–291.
- Keren, G. and W. Wagenaar (1987), 'Violations of Utility Theory in Unique and Repeated Games', *Journal of Experimental Psychology* **13**, 387–391.
- Knetsch, J. (1989), 'The Endowment Effect and Evidence of Non-reversible Indifference Curves', *American Economic Review* **79**, 1277–1284.
- Knetsch, J. and J. A. Sinden (1984), 'Willingness to Pay and Compensation Demanded: Experimental Evidence of an Unexpected Disparity in Measures of Values', *Quarterly Journal of Economics* **99**, 507–521.
- Knetsch, J., F. -F. Tang and R. Thaler (2001), 'The Endowment Effect and Repeated Market Trials: Is the Vickrey Auction Demand Revealing?', *Experimental Economics* **4**, 257–269.
- Kuilen, G. van de and P. P. Wakker (2004), 'Learning in the Allais Paradox', paper presented at FUR XI, Paris, June 2004, Mimeo, University of Amsterdam.
- Lichtenstein, S. and P. Slovic (1971), 'Reversals of Preference Between Bids and Choices in Gambling Decisions', *Journal of Experimental Psychology* **89**, 46–55.
- List, J. (2003), 'Does Market Experience Eliminate Market Anomalies', *Quarterly Journal of Economics* **118**, 41–71.
- List, J. and J. Shogren (1999), 'Price Information and Bidding Behavior in Repeated Second-Price Auctions', *American Journal of Agricultural Economics* **81**, 942–949.
- Loewenstein, G. and D. Adler (1995), 'A Bias in the Prediction of Tastes', *Economic Journal* **105**, 929–937.
- Loewenstein, G., T. O'Donoghue and M. Rabin (2003), 'Projection Bias in Predicting Future Utility', *Quarterly Journal of Economics* **118**, 1209–1248.

- Loomes, G., P. Moffatt and R. Sugden (2002), 'A Microeconomic Test of Alternative Stochastic Theories of Risky Choice', *Journal of Risk and Uncertainty* **24**, 103–130.
- Loomes, G., C. Starmer and R. Sugden (2003), 'Do Anomalies Disappear in Repeated Markets?', *Economic Journal* **113**, C153–C166.
- Loomes, G. and R. Sugden (1983), 'A Rationale for Preference Reversal', *American Economic Review* **73**, 428–432.
- Loomes, G. and R. Sugden (1998), 'Testing Different Stochastic Specifications of Risky Choice', *Economica* **65**, 581–598.
- McCabe, K.A., S.J. Rassenti and V.L. Smith (1990), 'Auction Institutional Design: Theory and Behaviour of Simultaneous Multiple-Unit Generalisations of the Dutch and English Auctions', *American Economic Review* **80**, 1276–1283.
- Noussair, C., S. Robin and B. Ruffieux (2004), 'Revealing Consumers' Willingness-To-Pay: A Comparison of the BDM Mechanism and the Vickrey Auction', *Journal of Economic Psychology* **25**, 725–741.
- Plott, C. R. (1996), 'Rational Individual Behavior in Markets and Social Choice Processes: The Discovered Preference Hypothesis', in K. Arrow, E. Colombatto, M. Perleman and C. Schmidt, eds., *Rational Foundations of Economic Behavior* (pp. 225–50). London: Macmillan and St. Martin's.
- Plott, C. R. and K. Zeiler (2005), 'The Willingness to Pay/Willingness to Accept Gap, the "Endowment Effect" and Experimental Procedures for Eliciting Valuations', *American Economic Review*, forthcoming.
- Seidl, C. (2002), 'Preference Reversal', *Journal of Economic Surveys* **6**, 621–655.
- Shogren, J., S. Cho, C. Koo, J. List, C. Park, P. Polo and R. Wilhelm (2001), 'Auction Mechanisms and the Measurement of WTP and WTA', *Resource and Energy Economics* **23**, 97–109.
- Shogren, J., S. Shin, D. Hayes and J. Kliebenstein (1994), 'Resolving Differences in Willingness To Pay and Willingness To Accept', *American Economic Review* **84**, 255–270.
- Shogren J. S. (2003), 'The X-Chapter. Experimental Methods and Valuation', in K. G. Maler and J. Vincent, eds., *Handbook of Environmental Economics*, Vol. II, Amsterdam: North-Holland.
- Schmidt, U., C. Starmer and R. Sugden (2004), 'Explaining Preference Reversal', paper presented at FUR XI, Paris June 2004, Mimeo.
- Slovic, P. (1995), 'The Construction of Preferences', *American Psychologist* **50**, 364–371.
- Slovic, P., B. Fischhoff and S. Lichtenstein (1982), Response Mode, Framing and Information-processing Effects in Risk Assessment, in R. M. Hogarth, eds., *New Directions for Methodology of Social and Behavioral Science: The Framing of Questions and the Consistency of Responses* (pp. 21–36). San Francisco: Jossey-Bass.
- Slovic, P. and A. Tversky (1974), 'Who Accepts Savage's Axiom?', *Behavioral Science* **19**, 368–373.
- Starmer, C. (2000), 'Developments in Non-expected Utility Theory: The Hunt for A Descriptive Theory of Choice under Risk', *Journal of Economic Literature* **38**, 332–382.
- Sugden, R. (2003), 'Reference-dependent Subjective Expected Utility Theory', *Journal of Economic Theory* **111**, 172–191.
- Tideman, T. (1983), 'An Experiment in the Demand Revealing Process', *Public Choice* **41**, 387–401.
- Tversky, A. and D. Kahneman (1991), 'Loss Aversion in Riskless Choice: A Reference-dependent Model', *Quarterly Journal of Economics* **106**, 1039–1061.
- Tversky, A., S. Sattath and P. Slovic (1988), 'Contingent Weighting in Judgement and Choice', *Psychological Review* **95**, 371–384.
- Van Huyck, J.B., R.C. Battalio and R.O. Beil (1993), 'Coordination Failure, Game Form Auctions, and Tacit Communication', *Games and Economic Behaviour* **5**, 485–504.