[16 May 2001]

Is there loss aversion in buying?
An adversarial collaboration

Ian Bateman[1]

Daniel Kahneman[2]

Alistair Munro[3]

Chris Starmer[4]

Robert Sugden[3]

1. School of Environmental Sciences, University of East Anglia
2. Woodrow Wilson School for Public and International Affairs, Princeton University
3. School of Economic and Social Studies, University of East Anglia
4. School of Economics, University of Nottingham

Address of corresponding author:

Robert Sugden, School of Economic and Social Studies, University of East Anglia, Norwich NR4 7TJ, England (e-mail address: <r.sugden@uea.ac.uk>)

[Acknowledgement to : Leverhulme Trust]

Abstract

(Provisonal)

This paper reports an adversarial collaboration between Daniel Kahneman and a group of British economists. The collaboration grew out of our attempts to reconcile an apparent difference between the results of two experiments, both of which have been presented as evidence that preferences are reference-dependent. We design and run a joint experiment to try to settle the question of whether, when consumers consider giving up money in exchange for goods, they feel loss aversion with respect to potential money outlays. The paper discusses both the substantive findings of the experiment and the methodological issues involved in adversarial collaboration.

This paper reports an exercise in a research methodology that we believe is new to experimental economics: *adversarial collaboration*. An adversarial collaboration is an experiment (or other empirical investigation) carried out jointly by two individuals or research groups who, having proposed conflicting hypotheses, seek to resolve the issue in dispute. The work we describe arises from our attempts to reconcile an apparent difference between the results of two experiments, both of which have been presented as evidence that preferences vary with reference states, and thus as evidence against the received theory of consumer choice. We attempt to settle the question of whether, when consumers consider giving up money in exchange for goods, they feel loss aversion with respect to potential money outlays. This issue in the theory of reference-dependent preferences has direct practical significance, since it bears on the explanation of the widely observed disparity between willingness-to-pay and willingness-to-accept valuations of costs and benefits, which has proved so problematic for contingent valuation studies.

This paper can be read on two different levels. At the level of methodology, it pioneers a method of resolving scientific disputes which has wide applicability across experimental economics. At the substantive level, it reports one of the largest experimental investigations to date into the determinants of loss aversion. Although the experiments we report are designed to discriminate between two alternative formulations of the theory of reference-dependent preferences, the null hypotheses for our tests are given by standard consumer theory. Hence, our experiments also constitute tightly-controlled tests of whether, as standard theory predicts, preferences over consumption bundles are invariant with respect to changes in reference points.

1. Adversarial collaboration

In an adversarial collaboration, the two parties agree on the design of a experiment which they will conduct jointly. Before knowing what the experiment will find, they accept its validity as a test – not necessarily a conclusive test – of their respective hypotheses. Each party anticipates its interpretation of possible outcomes of the experiment, particularly those that it does *not* predict. The two parties agree that particular outcomes of the experiment would support one hypothesis, and particular other outcomes would support the other. Both parties commit to publishing the results, whatever they may be.[1] We believe that this

methodology has advantages over the more conventional form of scientific debate, in which each research group designs and runs experiments independently, chooses which of its results to publish, and can challenge the validity of other groups' experimental designs after knowing the results those designs have produced. Adversarial collaboration encourages a more constructive approach to the resolution of disagreements.

Adversarial collaboration, as compared with conventional scientific debate, requires different attitudes on the part of researchers – in particular, more attention to understanding the other side's arguments, and less to rhetorical strategies for defeating them. But it also requires different expectations on the part of the scientific community as a whole. We are all used to reading journal articles which report apparently clear-cut experimental results and which draw strong conclusions from them. But as readers, we learn to apply some discount to such claims. We have to allow for confirmation bias in the design of experiments – the tendency for researchers to look for 'tests' which seem likely to confirm their prior hypotheses. We also have to keep in mind that the experiments that are reported in the journals are not necessarily representative of the larger set of experiments that have been run: we have to allow for the possibility that research groups publish only their most 'successful' experiments, and use conformity with their prior hypotheses as one of their criteria of success. And we expect the authors of journal articles to talk up their conclusions, drawing wide-ranging implications from them and down-playing doubts and ambiguities. Adversarial collaboration must be expected to lead to a different kind of publication.

Because the experimental designs used in adversarial collaboration have to be agreed by both parties, each party has to subject its hypothesis to a genuinely stringent test. The guarantee of stringency is this: each party's hypothesis is being subjected to a test that the other party expects it to fail. Thus, one of the mechanisms which tends to generate *apparently* decisive experimental results in the existing literature, positive confirmation bias, is neutralized. The commitment to publication, backed up by the two parties' common knowledge of the outcomes of the experiment, neutralizes another such mechanism: selection bias at the publication stage.

Adversarial collaboration will not always bring the parties into full agreement about the issue in dispute: they may have different interpretations of what their jointly-conducted experiment has found. Scientific debate is better served if such differences are reported frankly than if they are concealed by bland generalities. From the reader's point of view, a

report of this kind may be more useful than the superficially more definite conclusions that are customarily expected of non-adversarial research papers. Ultimately, however, the value of an adversarial collaboration is to be found in the validity and power of the experimental design it has adopted, and in the quality of the data this has generated. Whether, having seen the results, the parties to the collaboration agree on how they should be interpreted is a secondary matter: it is the reader who must draw the conclusions.

2. Theoretical background

The hypothesis that there is an asymmetry between individuals' attitudes to gains and to losses was first brought to the attention of economists by Kahneman and Amos Tversky (1979) and by Richard Thaler (1980). Since then, there has been an accumulation of evidence – from experiments, from survey data, and from the field – which suggests that individuals' choices are more responsive to anticipated losses than to equal and opposite anticipated gains.[2] These findings are inconsistent with the standard (or *Hicksian*) theory of consumer choice, in which preferences over final consumption states are independent of individuals' current endowments.

Probably the most fully-developed theoretical explanation of this asymmetry is the theory of *reference-dependent preferences*, proposed by Tversky and Kahneman (1991). In this theory, individuals have preference orderings over bundles of goods, as in Hicksian consumer theory, but these preferences are defined relative to *reference states*. A reference state is a point in goods space which the individual treats as the status quo or normal expectation; gains and losses in the various dimensions of goods space are defined in terms of displacements from the reference state. In notation, reference states are represented by subscripting the preference relation; thus 'x is weakly preferred to y, viewed from the reference state r' is written as $x \succsim_r y$.

Tversky and Kahneman propose a hypothesis of *loss aversion* which links changes in (reference-dependent) preferences with changes in the reference state. Let $x = (x_1, ..., x_n)$ and $y = (y_1, ..., y_n)$ be two bundles of the same n goods, such that for some good i, $y_i > x_i$. Let $r = (r_1, ..., r_n)$ and $s = (s_1, ..., s_n)$ be potential reference states, such that $r_i = x_i$, $s_i = y_i$, and $r_j = s_j$ for all $j \neq i$. For any given i, there is *loss aversion in good i* if, for all such x, y, r, s: $y \succsim_r x \Rightarrow y \succ_s x$. Tversky and Kahneman's hypothesis is that there is loss aversion, so defined, in all

goods.[3] It is an implication of this hypothesis that, for all x, y: $y \succsim_x x \Rightarrow y \succ_y x$. This latter property of the theory may be interpreted as a tendency, other things being equal, to prefer the reference state to other alternatives.

In order to apply Tversky and Kahneman's theory to a concrete choice problem, it is necessary to specify the chooser's reference state. The theory itself does not tell us how reference states are determined; in this respect, the formal model is left uninterpreted (Tversky and Kahneman, 1991, pp. 1046-1047). Thus, how the theory should be applied to specific decision-making environments can be a matter of judgement. One such difference of judgement led to our adversarial collaboration.

Bateman, Munro, Starmer and Sugden have worked together since the early 1990s on the effects of loss aversion in economic environments. Since all of them were based at the University of East Anglia when the work we report was carried out, we shall call them 'the UEA group'. In their experimental tests of reference-dependent preference theory, the UEA group have taken an individual's reference state to be whatever bundle of goods she currently owns. This *current endowment hypothesis* provides a simple method of deriving testable hypotheses from reference-dependent preference theory. In the light of their previous experimental findings, the UEA group have treated the current endowment hypothesis as descriptively adequate, at least in relation to simple laboratory tasks of buying, selling and choosing.[4] Munro and Sugden (1998) have used the hypothesis more generally in an analysis of how the workings of markets are affected by loss aversion.

In contrast, Tversky and Kahneman (1991) do not fully endorse the current endowment hypothesis, even with respect to simple laboratory tasks. Among the cases in which they think this hypothesis may be inappropriate is an experimental design in which subjects buy and sell coffee mugs for money, indicating their willingness to trade by reporting their valuations of mugs as 'buyers', as 'sellers' and as '
two experiments using this design, choosers' valuations were much closer to buyer's valuations than to sellers' (Kahneman, Knetsch and Thaler, 1990), Tversky and Kahneman conclude: 'The buyers in these markets do not appear to value the money they give up in a transaction as a loss. These observations are consistent with the standard theory of consumer choice, in which the decision of whether or not to purchase a good is treated as a choice between it and other goods that could be purchased instead' (1991, p. 1055). Tversky and Kahneman's hypothesis is that a subject who is considering buying a coffee mug construes

her reference state as including *neither* the mug *nor* the money she would have to spend to buy it. If she chooses to buy, she gains the mug; if not, she gains whatever she would buy with the money instead. We shall call this hypothesis *no loss aversion in buying* (NLAIB).

Explaining the intuition behind NLAIB, Tversky and Kahneman give two further examples. The first is one in which a person exchanges one $5 bill for five $1 bills. In this case, they argue, the person does not feel any pain of loss. The transaction is not construed as the combination of a loss and a separate gain (which would bring loss aversion into play); instead, the two sides of the transaction are integrated before any valuation takes place, so that the whole transaction is represented mentally as a net gain of zero on a single dimension ('cash'). The second example is a routine commercial transaction in which a trader (say, a shoe merchant) sells goods from stock. The merchant does not feel a pain of loss in parting with shoes, because to him the shoes already 'have the status of tokens for money'. The general intuition is that if a person gives up A to get B, there is loss aversion if and only if A is valuable to her in some way that B is not.

Tversky and Kahneman do not generalize beyond these examples. In the course of the adversarial collaboration, the UEA group discussed with Kahneman how to firm up the NLAIB hypothesis by specifying the class of buying tasks to which it applies. In response, he proposed a theory of *separability*. In this context, 'separability' means the separate evaluation of the gains and losses associated with a given transaction; its opposite is the *integration* of those gains and losses prior to evaluation. When a loss and an equal and opposite gain are integrated in this way, loss aversion does not arise. Kahneman proposes that consumers normally have *budget reserves*, i.e. reserves of money that are available for unanticipated spending. When an individual faces an unanticipated opportunity to buy a good, and is able to finance this spending from her budget reserve, gains and losses are integrated: the money that has to be spent to buy the good is already seen as a token for unspecified goods. Thus, Kahneman argues, NLAIB holds in such situations. In contrast, if the individual faces an unanticipated buying opportunity which she can finance only by forgoing some *specific* consumption plan, the act of buying involves a definite loss, separable from the gain; and so NLAIB does not hold. In the converse case of selling, gains and losses are integrated if the proceeds of the sale are earmarked for the purchase of a replacement good; but separability holds if those proceeds will be added to the budget reserve. When experimental subjects are dealing in low-value but non-staple goods such as coffee mugs, it

seems reasonable to assume that purchases can be financed from budget reserves, and that the proceeds from sales are not earmarked for replacements (Kahneman and Novemsky, 2000). The development of this analysis of separability illustrates one of the positive features of adversarial collaboration: theoretical progress may be stimulated through the process of the parties attempting to understand one another's positions.

3. Measures of individuals' valuations

Before going further, it will be useful to define some measures of individuals' valuations of changes in consumption. The current endowment hypothesis and the NLAIB hypothesis lead to different implications about the nature of disparities between alternative valuation measures. These conflicting implications are of practical significance, since they concern measures that are frequently used in contingent valuation and cost-benefit studies. They are also central to the logic of the experiment we report in this paper.

Consider a model in which there are only two goods; quantities of these goods are represented by x_i, x_j. We shall often interpret this model so that one good is some particular private consumption good (for short, we shall call this 'the good') and the other is an index of general purchasing power, or 'money'. When we use this interpretation, the good will be denoted by 1 and money by 2. For any given individual, consider how one unit of good i can be valued in units of good j. (Notice that there is no loss of generality in speaking of the value of 'one unit' of good i: as modellers, we are free to choose the units in which good i is measured so that any given real quantity counts as one unit.) More specifically, for any given quantities x_i', x_j' of the two goods, consider how we might express in units of good j the value of consuming the bundle $(x_i' + 1, x_j')$ rather than (x_i', x_j').

We define six such measures of this value. The first four of these measures are well-known; the significance of the other two will emerge later:[5]

1. *Willingness to pay (WTP).* Suppose the individual's current endowment is (x_i', x_j'). WTP_{ji} is the largest amount of good j that the individual would be willing to give up in return for a gain of one unit of good i.

2. *Willingness to accept (WTA).* Suppose the individual's current endowment is $(x_i' + 1, x_j')$. WTA_{ji} is the smallest amount of good j that the individual would be willing to accept in return for accepting a loss of one unit of good i.

3. *Equivalent loss (EL).* Suppose the individual's current endowment is $(x_i' + 1, x_j')$. EL_{ji} is the largest amount of good j that the individual would be willing to give up in place of a loss of one unit of good i.

4. *Equivalent gain (EG).* Suppose the individual's current endowment is (x_i', x_j'). EG_{ji} is the smallest amount of good j that the individual would be willing to accept in place of a gain of one unit of good i.

5. *Risky willingness to pay (RWTP).* Suppose the individual's current endowment is (x_i', x_j'). Consider a gamble with two mutually exclusive outcomes, each with probability 0.5. One outcome is that the individual gains one unit of good i, with no change in good j. The other is that she loses some amount of good j, with no change in good i. $RWTP_{ji}$ is the largest such loss of good j consistent with her being willing to accept the gamble.

6. *Risky willingness to accept (RWTA).* Suppose the individual's current endowment is $(x_i' + 1, x_j')$. Consider a gamble with two mutually exclusive outcomes, each with probability 0.5. One outcome is that the individual loses one unit of good i, with no change in good j. The other is that she gains some amount of good j, with no change in good i. $RWTA_{ji}$ is the smallest such gain of good j consistent with her being willing to accept the gamble.

It is important to notice that these definitions do not presuppose any theory of preferences; they are to be interpreted as observable magnitudes, as revealed in the individual's behaviour in some given setting, and not as constructs within a particular theory (in the sense that the Hicksian concepts of compensating and equivalent variation are).

4. Theoretical predictions: valuation under certainty

In this Section, we consider WTP, WTA and EG measures of the relative value of a consumption good (good 1) and of money (good 2).

Hicksian theory predicts that, if wealth effects are weakly positive:

$$WTP_{21} \leq EG_{21} = WTA_{21}, \text{ and} \qquad (1a)$$

$$WTP_{12} \leq EG_{12} = WTA_{12}. \qquad (1b)$$

If (as in most experimental environments), WTA_{21} is small relative to the individual's total wealth, credible values of the rate of change of WTP_{21} with respect to wealth imply that the

value of WTA_{21}/WTP_{21} is close to 1.[6] For symmetrical reasons, the value of WTA_{12}/WTP_{12} can also be expected to be close to 1 in experiments using low-value goods.

The theory of reference-dependent preferences, *in conjunction with the current endowment hypothesis*, predicts that if wealth effects are weakly positive:

$$WTP_{21} < EG_{21} < WTA_{21}, \text{ and} \quad (2a)$$

$$WTP_{12} < EG_{12} < WTA_{12}. \quad (2b)$$

The prediction $EG_{21} < WTA_{21}$ is an implication of loss aversion in the good. Intuitively, WTA_{21} is based on a comparison between gains of money and *losses* of the good, while EG_{21} is based on a comparison between gains of money and *gains* of the good. If preferences are more responsive to losses than to gains, as the hypothesis of loss aversion implies, WTA_{21} will be greater than EG_{21}. Symmetrically, $EG_{21} < WTAG_{21}$ is an implication of loss aversion in money. The other two inequalities result from the conjunction of loss aversion and wealth effects; loss aversion in money contributes to $WTP_{21} < EG_{21}$, while loss aversion in the good contributes to $WTP_{12} < EG_{12}$.[7]

However, if the theory of reference-dependent preference is combined with NLAIB, the implications are different. Notice that WTP_{21}, EG_{21} and WTA_{21} correspond with the valuations of 'buyers', 'choosers' and 'sellers' respectively, as discussed by Tversky and Kahneman in the context of the coffee-mug experiment. If NLAIB applies (which, according to Kahneman's theory, will be the case if the relevant money outlays can be covered from budget reserves), the money outlays involved in WTP_{21} valuations are not construed as losses. Thus, differences between EG_{21} and WTP_{21} can be caused only by wealth effects. Recall that WTA_{12} measures an individual's willingness to accept the good in return for giving up money: this measure, just like WTP_{21}, refers to situations in which the individual spends money to buy the good. So if NLAIB applies, the money outlays involved in WTA_{12} valuations are not construed as losses; since wealth effects cannot contribute to differences between WTA_{12} and EG_{12}, the implication is that $WTA_{12} = EG_{12}$. Thus, again on the assumption that wealth effects are weakly positive, the predictions of reference-dependent theory *in conjunction with NLAIB* are:

$$WTP_{21} \leq EG_{21} < WTA_{21}, \text{ and} \quad (3a)$$

$$WTP_{12} < EG_{12} = WTA_{12}. \quad (3b)$$

The differences between these three sets of predictions allow the two variants of reference-dependent preference theory to be tested, both against conventional Hicksian consumer theory and against one another.

These conflicting predictions also have significance for contingent valuation studies. The large disparities between WTP_{21} and WTA_{21} that are frequently found in survey data pose a serious problem for contingent valuation methodology.[8] Many practitioners recommend using WTP_{21} rather than WTA_{21} valuations on the grounds that the former better reflect individuals' true preferences (e.g. Arrow et al, 1993). *If* loss aversion can be interpreted as a bias,[9] NLAIB (if true) provides support for that recommendation: if NLAIB holds, WTA_{21} valuations pick up the effects of loss aversion in the good, while WTP_{21} valuations do not pick up *any* loss aversion effects. But if instead the current endowment hypothesis is true, there is a relation of symmetry between WTP_{21} and WTA_{21}: the former picks up loss aversion in money, while the latter picks up loss aversion in the good. Then, if one wishes to screen out the effects of loss aversion in contingent valuation studies, EG_{21} valuations – which do not require respondents to consider losses, and which treat money and the good symmetrically – may be more suitable than either WTP_{21} or WTA_{21} (Bateman et al, 1997).

5. Theoretical predictions: risky willingness to pay and risky willingness to accept

In the discussions which preceded the design of our experiment, Kahneman proposed a restricted form of prospect theory as a means of generating sharp predictions, particularly about RWTP and RWTA valuations. The reasoning which leads to this model combines elements of Tversky and Kahneman's (1991) theory of reference-dependent preferences and of Kahneman and Tversky's (1979) prospect theory, along with additional simplifying assumptions.

Assume, as in prospect theory, that reference-dependent preferences over consumption bundles depend only on the displacement of each bundle from the reference state.[10] Assume that these preferences can be represented by an additively separable *value function*. Thus, in the case of two goods, the value function can be written as $v(\Delta x) = v_1(\Delta x_1) + v_2(\Delta x_2)$, where $\Delta x = (\Delta x_1, \Delta x_2)$ is a *displacement vector* of changes in consumption relative to the reference state. Let Δx^G and Δx^L be displacement vectors, such that Δx^G is an

unambiguous improvement relative to the reference state and Δx^L is an umambiguous worsening. Consider a *balanced lottery* which gives each of Δx^G and Δx^L with probability 0.5. Applying prospect theory,[11] whether the individual prefers this lottery to the reference state depends on the sign of $\pi(0.5)v(\Delta x^G) + \pi(0.5)v(\Delta x^L)$, where $\pi(.)$ is the *probability weighting function*. Because the probabilities of gain and loss are equal, the $\pi(.)$ terms can be cancelled out: a balanced lottery is preferred to, indifferent to, or less preferred than the reference state according to whether $v(\Delta x^G) + v(\Delta x^L)$ is greater than, equal to, or less than zero.

As a simplification, assume the following functional form for the value function: for each good i, $v_i(\Delta x_i) = a_i(\Delta x_i)^\beta$ if $\Delta x_i \geq 0$ and $v_i(\Delta x_i) = b_i(\Delta x_i)^\beta$ if $\Delta x_i \leq 0$, where a_i, b_i and β are constants satisfying $a_i > 0$, $b_i > 0$, $1 \geq \beta > 0$. Given this functional form, the value of b_i/a_i is a natural measure of loss aversion in good i. This model implies the following relationships:

$$EG_{ji} = [a_i/a_j]^{1/\beta} \Delta x_i \quad (4)$$

$$EL_{ji} = [b_i/b_j]^{1/\beta} \Delta x_i \quad (5)$$

$$RWTP_{ji} = [a_i/b_j]^{1/\beta} \Delta x_i \quad (6)$$

$$RWTA_{ji} = [b_i/a_j]^{1/\beta} \Delta x_i. \quad (7)$$

Hence:

$$RWTA_{ji} / EG_{ji} = [b_i/a_i]^{1/\beta} \quad (8)$$

$$EG_{ji} / RWTP_{ji} = [b_j/a_j]^{1/\beta}. \quad (9)$$

Recall that b_i/a_i and b_j/a_j are indices of loss aversion in good i and good j respectively.

If separability holds, the model also implies:

$$WTP_{ji} = RWTP_{ji} = [a_i/b_j]^{1/\beta} \Delta x_i \quad (10)$$

$$WTA_{ji} = RWTA_{ji} = [b_i/a_j]^{1/\beta} \Delta x_i. \quad (11)$$

The equalities $WTP_{ji} = RWTP_{ji}$ and $WTA_{ji} = RWTA_{ji}$ constitute a surprising and distinctive prediction of prospect theory: that there is no risk aversion with respect to balanced lotteries.[12] On the hypothesis that lies behind NLAIB, separability holds when an individual considers giving up the good to gain money. Thus, (10) applies in the case of WTP_{12} and

(11) in the case of WTA_{21}. But, according to the same hypothesis, separability does *not* hold when money is being given up to gain the good, as in the tasks that elicit WTP_{21} and WTA_{12}. Such tasks are equivalent to choices among gains. Thus:

$$WTP_{21} = EG_{21} = [a_1/a_2]^{1/\beta} \Delta x_1 \quad (12)$$

$$WTA_{12} = EG_{12} = [a_2/a_1]^{1/\beta} \Delta x_2. \quad (13)$$

6. Empirical background

The starting point for our adversarial collaboration was an apparent conflict between the findings of two previous experimental investigations of loss aversion.

In the experiments to which we have already referred, Kahneman et al (1990, experiments 6 and 7) elicited money valuations of coffee mugs using three different measures: WTP_{21}, EG_{21} and WTA_{21}. In one experiment, the median values of WTP_{21}, EG_{21} and WTA_{21} were respectively \$2.87, \$3.12, and \$7.12 (in Canadian dollars). In a second experiment, the median valuations were \$2.00, \$3.50, and \$7.00. Kahneman et al interpret these results as evidence that, when individuals buy low-value goods, there is little loss aversion in money. (Notice that in interpreting these results as supporting their hypothesis about buying tasks, Kahneman et al are implicitly weakening that hypothesis from '*no* loss aversion in buying' to some proposition of the form '*attenuated* loss aversion in buying'.)

Bateman et al (1997) elicited WTP_{21}, WTA_{21}, EG_{21}, EL_{21}, WTP_{12}, WTA_{12}, EG_{12} and EL_{12} measures of individuals' relative valuations of money and particular consumption goods (in one condition, luxury chocolates, in another, cans of Coke). The experimental design tested for loss aversion in goods and in money by comparing the preferences implied by these different measures; comparisons were made in a way that screened out income effects. Overall, the results gave strong support to the theory of reference-dependent preferences, interpreted in terms of the current endowment hypothesis. The design allowed four independent tests for loss aversion in money in relation to tasks in which subjects gave up money to buy goods (the comparisons between WTP_{21} and EG_{21}, and between WTA_{12} and EG_{12}, for the chocolate and Coke conditions). In each of these comparisons, the null hypothesis that the two measures elicit the same preferences is an implication of NLAIB as well as of Hicksian theory. In three out of the four tests, the null hypothesis is rejected (the exception is the comparison between WTA_{12} and EG_{12} in the case of Coke). The data do not

support even the weaker and less formal hypothesis that loss aversion in buying is a weak effect, relative to loss aversion in selling: the median values of WTP_{21}, EG_{21} and WTA_{21} were £0.50, £1.50 and £2.00 for chocolate, and £0.60, £1.00 and £1.50 for Coke.[13]

Our adversarial collaboration developed out of our attempts to reconcile these apparently conflicting conclusions about the truth or falsity of NLAIB. In informal discussions with the UEA group, Kahneman pointed to two features of the Bateman et al experiment which, he suggested, might account for its failure to confirm NLAIB. First, unlike the Kahneman et al experiment, it had used a *random lottery* design. Each subject faced several different tasks, knowing that just one of these would be for real; which task was for real was determined by a random process at the end of the experiment. The subject's endowment varied between tasks. Thus, when a subject faced any given task, her 'current endowment' (as defined by the experimenters) was conditional on that task being the one selected by the random lottery mechanism. Kahneman argued that this feature of the design would tend to reduce all forms of loss aversion.[14]

Second, Kahneman noted that Bateman et al elicited valuations using *open-ended valuation* tasks (e.g. 'What is the largest amount of money you would be willing to pay for ... ?'). He thought this design feature less satisfactory than the *multiple dichotomous choice* tasks (e.g. 'Would you be willing to pay $1 for ... ?', 'Would you be willing to pay $2 for ...?', and so on) used by Kahneman et al. To explain the point of this criticism, we need to make a distinction between two forms of valuation. An *outgoing* valuation records the largest amount of some good that an individual is willing to transfer *to* someone else; an *incoming* valuation records the smallest amount of some good that an individual is willing to accept as a transfer *from* someone else. WTP_{21}, WTP_{12}, EL_{21} and EL_{12} are outgoing valuations, while WTA_{21}, WTA_{12}, EG_{21} and EG_{12} are incoming valuations. Subjects may follow *tactical* or *cautious* heuristics which lead them to understate their true outgoing valuations and to overstate their true incoming valuations. Although such heuristics do not in fact serve a subject's interests in incentive-compatible experiments such as those we discuss in this paper, they may be well-adapted to many real-world situations in which terms of trade are determined through bargaining. More generally, the principle of erring on the side of caution in appraising other people's trading proposals may be useful in a world in which not all economic interactions are positive-sum games. Kahneman argued that dichotomous choice tasks are less likely than open-ended valuation tasks to evoke tactical heuristics,

because the format of a dichotomous choice task prompts subjects to respond as though facing an exogenously determined price.

Arguing from these two features of the experiment, Kahneman suggested that the Bateman et al results might be primarily the product of tactical heuristics, rather than of loss aversion. In the course of our discussions, it became clear that the most effective way of resolving our disagreement would be to carry our a further experiment. Hence our adversarial collaboration.

7. Designing a new experiment: principles

In trying to develop a design for a new experiment to test NLAIB, we quickly agreed that it would have the following four features. First, it would work by eliciting subjects' relative valuations of money and some low-value, non-staple consumption good. Second, valuations would be elicited by multiple dichotomous choices. Third, each subject would face just one valuation task, which would be for real. Finally, there would be no tasks (like the EL_{21} and EL_{12} tasks in the Bateman et al experiment) in which subjects have to choose between losing money and losing a specific good.

The first feature was necessary to satisfy the preconditions of the NLAIB hypothesis, as firmed up by Kahneman. We were able to agree on the second feature because, in the light of their experience after designing the Bateman et al experiment, the UEA group had come to agree with Kahneman that open-ended valuations are more likely to evoke tactical heuristics than are dichotomous choices. The third and fourth features ensure that, if a subject is given any initial endowment by the experimenters, she genuinely owns it: it is not conditional on any random event, and (since she has the option of refusing to engage in any exchanges) she is free to leave the experiment with it. We agreed that these features would tend to make anticipations of loss salient to subjects.

We further agreed that the experiment should not only test NLAIB in its strict form, but should also generate evidence relevant for the appraisal of a weaker and less precise hypothesis about loss aversion in buying tasks. Kahneman argued that his theory of separability, which implies NLAIB in its strict form, should be interpreted as a simplified model of a more complex reality; the economically significant core of his claim about buying tasks was that, in these tasks, loss aversion is *either* absent *or* greatly attenuated. As this

weaker hypothesis is not the implication of any fully-specified model, it can be firmed up in different ways. In the spirit of adversarial collaboration, we sought to agree in advance what evidence would count in support of that hypothesis.

Beyond these points, however, agreement on an experimental design proved more difficult to achieve.

Kahneman proposed a design based on comparisons between WTP_{21}, WTA_{21}, EG_{21}, $RWTP_{21}$ and $RWTA_{21}$. On the assumption that separability holds when the good is given up in exchange for money, but not when money is given up in exchange for the good, the model set out in Section 5 generates the prediction:

$$RWTP_{21} < WTP_{21} = EG_{21} < WTA_{21} = RWTA_{21}. \quad (14)$$

The first inequality in (11) reflects loss aversion in money, while the second reflects loss aversion in the good. (In the model, $WTP_{21} / RWTP_{21} = [b_2/a_2]^{1/\beta}$, and $WTA_{21} / EG_{21} = [b_1/a_1]^{1/\beta}$.) Kahneman's expectation, based on previous experimental results, was that the ratios $WTP_{21} / RWTP_{21}$ and WTA_{21} / EG_{21} would take values close to 2.[15] He argued that, considered as a whole, (14) was a bold prediction. As far as he knew, it was made by no other theory. Were this particular pattern to be found among the five valuations, that would provide confirmation for a theory which included NLAIB as one of its elements. The onus would then be on the proponents of rival theories to find alternative explanations for the regularity. However, if observed valuations showed the pattern $WTP_{21} < EG_{21} < WTA_{21}$, that would be evidence against NLAIB, and in favour of the current endowment hypothesis.

As a way of assessing whether loss aversion is attenuated in buying tasks, Kahneman proposed to compare the ratios EG_{21}/WTP_{21} and WTA_{21}/EG_{21}. If the value of EG_{21}/WTP_{21} was only slightly greater than 1 and markedly less than WTA_{21}/EG_{21}, that would count as evidence in favour of attenuation. If, instead, EG_{21}/WTP_{21} and WTA_{21}/EG_{21} were both greater than 1 and were similar in magnitude to one another, that would count as evidence against.

The UEA group proposed a design based on comparisons between WTA_{21} and EG_{21}, and between WTA_{12} and EG_{12}. As we showed in Section 4, Hicksian theory predicts $WTA_{21} = EG_{21}$ and $WTA_{12} = EG_{12}$. These would be the null hypotheses. The conjunction of reference-dependent preference theory and the current endowment hypothesis predicts $EG_{21} < WTA_{21}$ (i.e. loss aversion in the good) and $EG_{12} < WTA_{12}$ (i.e. loss aversion in money). The

conjunction of reference-dependent theory and NLAIB predicts $EG_{21} < WTA_{21}$ and $EG_{12} = WTA_{12}$. Thus, the two versions of reference-dependent theory could be tested against Hicksian theory and against each other. If the observed values of WTA_{12}/EG_{12} and WTA_{21}/EG_{21} were both greater than 1, but the former ratio was much closer to 1 than was the latter, that would count as evidence of the attenuation of loss aversion.

The UEA group favoured this design as a means of testing for loss aversion while controlling for a wide range of potentially confounding factors. For example, consider a comparison between EG_{12} and WTA_{12} as a test for loss aversion in money. (The following arguments about experimental control apply with equal force to comparisons between EG_{21} and WTA_{21} as a means of testing for loss aversion in the good.) Suppose that individuals' valuations of £1.00, in units of chocolate, are elicited by multiple dichotomous choices. A representative choice question for eliciting EG_{12} is: 'Choose one of the following options: *Either* we give you £1.00, *or* we give you x chocolates'. The corresponding question for eliciting WTA_{12}, asked of a subject who has previously been given £1.00, is: 'Choose one of the following options: *Either* you give us your £1.00 and take x chocolates in exchange, *or* you keep your 10 chocolates'. Notice that the two questions offer *exactly the same* pair of alternative final states: either the subject leaves the experiment with £1.00 more than she came with, or she leaves it with x chocolates more. Thus, these choice problems are identical from the viewpoint of *any* theory that postulates that an individual has a single set of preferences over final states. So systematic differences between EG_{12} and WTA_{12} cannot be explained by any such theory. In particular, such differences cannot be attributed to Hicksian income or substitution effects. Nor can such differences plausibly be attributed to tactical heuristics. If, instead of treating each dichotomous choice as a distinct task, a subject construes the whole set of choices as a single valuation task, tactical heuristics might come into play; but since both EG_{12} and WTA_{12} are incoming valuations, the two types of question have the same tendency to prompt the over-statement of true valuations.

As a way of testing Kahneman's hypotheses about the valuations of balanced lotteries, the UEA group suggested an extension to this design, in which $RWTA_{21}$ and $RWTA_{12}$ would also be elicited. The restricted form of prospect theory presented in Section 5 generates the predictions $EG_{12} = WTA_{12} < RWTA_{12}$ and $EG_{21} < WTA_{21} = RWTA_{21}$. Since all the relevant valuations are incoming, tests of these predictions are not confounded by the effects of tactical heuristics.

Each of the parties to the adversarial collaboration recognised the validity of both designs. But even after prolonged discussion, each continued to prefer its own proposal. Kahneman had reservations about the UEA group's proposal to use tasks which elicit valuations of fixed amounts of money in units of a consumption good such as chocolate. He accepted the theoretical arguments that generated the UEA group's null and alternative hypotheses; but he was concerned that subjects would have difficulty in understanding tasks of this kind, or not construe them as buying tasks, even when the tasks were presented in the form of dichotomous choices. The UEA group accepted that, if an experiment based on Kahneman's proposal found a clear pattern corresponding with (14), that result would be strong evidence in support of NLAIB. But if the experiment failed to find that pattern, the results would be open to many alternative interpretations. Thus, the design did not have the capacity to deliver sharp results in confirmation of the current endowment hypothesis, were that hypothesis true.

In order to continue the adversarial collaboration, we combined the two proposals. We agreed on a design that would elicit all of the valuations EG_{21}, WTP_{21}, WTA_{21}, $RWTP_{21}$, $RWTA_{21}$, EG_{12}, WTA_{12} and $RWTA_{12}$. We agreed to carry out both the tests proposed by Kahneman and those proposed by the UEA group. Without agreeing on the relative importance to be attached to the two sets of tests, we agreed that each of them was valid as a means of discriminating between NLAIB and the current endowment hypothesis.

8. The experiment

The experiment was carried out at the University of East Anglia. Subjects were recruited from the undergraduate population by means of e-mailed invitations; they were broadly representative of that population in terms of age, gender, and subject of study. Subjects were required to bring cash to the experiment, but were assured that any opportunities to spend money would be optional. Initially, we recruited 320 subjects. Each subject was allocated at random to one of eight treatments, each of which was designed to elicit one of the valuations EG_{21}, WTP_{21}, WTA_{21}, $RWTP_{21}$, $RWTA_{21}$, EG_{12}, WTA_{12} and $RWTA_{12}$. The specific good took the form of luxury chocolates sold by a specialist shop located in the centre of Norwich, easily accessible from the university campus. These chocolates are sold by weight, at an average price of about £0.30 each. To allow exchanges in units of single chocolates to be carried out conveniently, transactions within the experimental sessions were carried out in

vouchers. A voucher entitled its holder to a specified number of chocolates, free of charge, when presented at the shop.

On arrival at a session, subjects were told that the experiment had two separate parts. They were told nothing about Part 2 until they had completed Part 1, except that any payoffs they might receive in Part 2 would be additional to their payoffs from Part 1. In fact, Part 2 was a choice, for real, between two lotteries with money prizes. The main purpose of this part of the experiment was to supplement subjects' final earnings, particularly in two treatments (WTP_{21} and $RWTP_{21}$) in which those earnings would otherwise have been rather low.[16] In this paper, we are concerned only with Part 1 of the experiment.

Depending on which task they had been assigned, subjects were given 'endowments' (which in some cases were 'nothing'). Subjects in the WTA_{12} and $RWTA_{12}$ groups were given £1.00. Those in the WTA_{21} and $RWTA_{21}$ groups were given 10 chocolates (in the form of vouchers). All other subjects were given nothing. Endowments (money or vouchers) were physically handed over to subjects. It was explained that subjects' endowments were theirs to keep if they so chose. The conditions for the use of the vouchers were explained, and samples of the chocolates were shown; no information was given about the price of the chocolates.

The concept of a 'lottery' was also explained. Lotteries would be resolved by the subject's drawing a disc from a bag containing 100 discs, numbered from 1 to 100. The outcomes of lotteries would be described in terms of what the subject would gain or lose, conditional on the number of the disc drawn.

Subjects were given booklets in which their tasks were set out. The rest of the instructions were printed in these booklets. As far as possible, the instructions were common to all treatments; the common elements were also read out by an experimenter, who fielded any questions. Subjects were told that they had to make twenty-five 'choices'. For each subject, one of these choices was for real. On arrival at the experiment, subjects had been shown a box containing twenty-five sealed envelopes, each containing a ticket with one of the numbers 1-25. Each subject picked one envelope from the box, to be opened at the end of the experiment. The number in the envelope was the number of the choice problem that was for real for that subject. This device was used to dramatize the fact that one and only one problem was for real, and that the identity of this problem was independent of the subject's

responses. At the end of the experiment, each subject carried out whatever transaction (if any) he had chosen in the problem that was for real.

Each choice problem required the subject to choose one of two 'options', one displayed on the left-hand side of the page and one on the right, by ticking the appropriate box. For any given subject, the right-hand option was the same in all twenty-five problems. The left-hand options differed only in respect of one parameter. In the EG_{21}, WTP_{21}, WTA_{21}, $RWTP_{21}$ and $RWTA_{21}$ treatments, this parameter was an amount of money from the set {£0.30, £0.60, £0.90, ..., £7.50}; we shall say that in these treatments the *response mode* was money. In the EG_{12}, WTA_{12} and $RWTA_{12}$ treatments, it was a number of chocolates from the set {1, ..., 25}; in these cases, the response mode was chocolate. For each treatment, half of the booklets presented the problems in ascending order (i.e. in Choice 1, the parameter was £0.30 or 1, in Choice 2 it was £0.60 or 2, and so on), while the other booklets presented them in descending order. Subjects were allocated randomly between these two presentations, so as to control for order effects. We required each subject's choices to be mutually consistent in the sense of respecting dominance.[17]

The EG_{21}, WTP_{21}, WTA_{21}, $RWTP_{21}$ and $RWTA_{21}$ treatments elicit, to within £0.30 bands, money valuations of 10 chocolates. (In the notation of Section 3, with m denoting a subject's money wealth before coming to the experiment and on the assumption that no one then owned chocolates: 10 physical chocolates constitute one 'unit' of chocolate, $x_1' = 0$, and $x_2' = m$.) The EG_{12}, WTA_{12} and $RWTA_{12}$ treatments elicit chocolate valuations of £1.00 (i.e. £1.00 constitutes one 'unit' of money, $x_1' = 0$, and $x_2' = m$). Figure 1 shows how a typical problem for each treatment (in each case, the twelfth problem in ascending order) was displayed. The entries in square brackets were not seen by the subjects; these have been added for the benefit of the reader, to identify the relevant treatment.

Two additional treatments were run in a follow-up experiment. The responses to the treatments described above turned out to indicate surprisingly low levels of loss aversion for chocolate, as measured by the ratio WTA_{21} / EG_{21}. Kahneman conjectured that this was the result of our having used vouchers rather than actual chocolates in the experimental sessions. The use of vouchers, he suggested, might attenuate loss aversion by mentally distancing subjects from the consumption experiences associated with the chocolates, and thus weakening the sense of ownership associated with chocolate endowments. In addition, since vouchers have some of the properties of money, the psychological mechanisms which (on his

account) give rise to NLAIB might also affect tasks in which vouchers are given up in trade. To test this conjecture,[18] we ran a follow-up experiment which repeated the EG_{21} and WTA_{21} treatments, exactly as before except for one detail: the '10 chocolates' took the form of a pre-packed box of 10 chocolates (the same kind as we had used before) rather than vouchers to be redeemed at the supplier's shop. At the start of the follow-up experiment, each subject in the WTA_{21} treatment was handed such a box as his endowment; subjects who retained or gained chocolates in the course of the experiment took a box of chocolates away with them. Using the same procedures as before, we recruited an additional 107 subjects and divided them at random between the two additional treatments. These treatments will be denoted by EG_{21}* and WTA_{21}*; we shall say that they involved *immediate chocolate* as contrasted with chocolate vouchers.

9. Results

The responses to the ten treatments are summarized in Tables 1 and 2. In presenting the data, we use the following conventions. Recall that for any given subject in any given treatment, there are 25 choice problems. Since subjects' choices are required to respect dominance, there are 26 alternative permissible ways of answering any such set of problems: *either* the left-hand option is chosen in every problem, *or* there is a switch from left to right (or right to left, depending on the task) after exactly one of the choices 1, ..., 24, *or* the right-hand option is chosen in every problem. Each response indicates a different *valuation* of the right-hand option, expressed in terms of one of the valuation measures defined in Section 3.

We assign these responses the values 1, ..., 26, in ascending order of the valuation of 10 chocolates (for tasks in which the response mode is money) or in ascending order of the valuation of £1.00 (for tasks in which the response mode is chocolate). Thus, for tasks in which the response mode is money, the valuation 1 corresponds with the range of money values of 10 chocolates from zero to £0.30; the valuation 2 corresponds with values from £0.30 to £0.60, and so on up to the valuation 26 which corresponds with values from £7.50 upwards. For tasks in which the response mode is chocolate, the valuation 1 corresponds with the range of chocolate values of £1.00 from 0 to 1 chocolate; the valuation 2 corresponds with values from 1 to 2 chocolates, and so on up to the valuation 26 which corresponds with values from 25 chocolates upwards.

We also report subjects' *implicit preferences* between 10 chocolates and £1.00. In treatments in which the right-hand option is 10 chocolates, a subject whose valuation is 1, 2 or 3 has chosen to have £0.90 rather than 10 chocolates, and so can be presumed to prefer £1.00 to 10 chocolates. Conversely, a subject whose valuation is 5 or more has chosen to have 10 chocolates rather than £1.20, and so can be presumed to prefer 10 chocolates to £1.00. (A valuation of 4 does not reveal the subject's preference either way.) In treatments in which the right-hand option is £1.00, the valuations 1, ..., 10 reveal an implicit preference for £1.00 over 10 chocolates, while the valuations 11, ... , 26 reveal the opposite preference. Implicit preferences are of interest because they are comparable across all treatments, irrespective of whether the response mode is money or chocolate.

Table 1 refers to the seven treatments for which the response mode was money. The upper part of the table reports, for each treatment, the geometric mean, arithmetic mean, median and standard deviation of the distribution of subjects' implicit valuations. We shall give particular attention to the geometric mean. This is because we are concerned with the values of *ratios* and, as summary statistics for describing ratios, geometric means are more satisfactory than arithmetic means. More specifically, we are concerned with two kinds of ratios. First, it is fundamental to our experimental design that the money and chocolate response modes are symmetrical with one another. Viewed in this perspective, valuations should be interpreted as ratios between quantities of money and (equally preferred) quantities of chocolate; whether these ratios are expressed as '£ per chocolate' of 'chocolates per £' is arbitrary. Suppose we have valuations v_1, ..., v_n in units of £/chocolate from n subjects. Since the geometric mean of $1/v_i$ is equal to the inverse of the geometric mean of v_i, the information content of the geometric mean of those valuations is independent of the units in which they are expressed. Second, in order to appraise the hypothesis of attenuated loss aversion in buying, we need to look at ratios of valuations (such as the ratio WTA_{21} / EG_{21}, used as an indicator of the extent of loss aversion in chocolate). Suppose we have valuations v_1, ..., v_n from one treatment and valuations w_1, ..., w_n from another treatment, with different subjects in the two treatments. However we index these subjects, the ratio between the geometric mean of v_i and the geometric mean of w_i is equal to the geometric mean of (v_i / w_i).[19]

The lower part of Table 1 reports, for each treatment, the distribution of responses classified by implicit preferences. Table 2 presents the corresponding data for the three treatments for which the response mode was chocolate.

Table 3 reports summary statistics for those comparisons between treatments that are relevant for our design. The first entry in each row identifies a ratio of two valuations. The next five entries indicate whether particular causal factors, if operating, have a tendency to increase the value of that ratio above unity. (Each of the five factors, if present, works in the same direction.) The second entry states whether loss aversion in chocolate and/or money would tend to increase the value of the ratio, and in the case of loss aversion in money, whether this tendency is conditional on the truth or falsity of NLAIB. The third entry states whether risk aversion with respect to balanced lotteries would have the same tendency. (Recall that prospect theory predicts the absence of such risk aversion.) The fourth entry states whether subjects' use of tactical heuristics would have that tendency. In this column, 'yes' signifies that the numerator of the ratio is an incoming valuation and that the denominator is an outgoing valuation; 'no' signifies that both valuations are of the same type. The fifth entry states whether Hicksian income and substitution effects would have that tendency. Here, 'yes' signifies that it is a prediction of Hicksian theory that, if income effects are normal, the ratio of valuations is greater than unity. (Notice, however, that the *size* of such predicted effects is likely to be tiny: see Section 3.) 'No' signifies that Hicksian theory predicts that the two valuations are exactly equal; 'n.a.' signifies that Hicksian theory makes no firm predictions (because of the presence of risk). The sixth entry indicates whether a difference between subjects' attitudes to immediate chocolate and to chocolate vouchers would impact on the value of the ratio; we assume that such an effect, if it existed, would imply higher valuations for immediate chocolates than for chocolate vouchers.

The seventh entry in each row is the ratio between the geometric means of the relevant valuations. We use this as our main indicator of the similarity or divergence between responses to different treatments. (If the reader wishes to compare medians or arithmetic means, the relevant information is given in Tables 1 and 2.) The final entry reports the z-statistic for a Mann-Whitney test for differences between the distributions of valuations in the two treatments; an asterisk denotes significance at the 5 per cent level in a one-tail test; two asterisks denote significance at the 1 per cent level. A positive sign indicates that the

median value of the 'numerator' valuation is greater than that of the 'denominator' valuation.[20]

To provide a benchmark for other comparisons, we begin by comparing the WTA_{21} and WTP_{21} treatments. Recall that subjects in the WTA_{21} treatment are endowed with chocolate and report their willingness to accept money in exchange for giving up their endowment; subjects in the WTP_{21} treatment are given no endowments and report their willingness to spend money to buy chocolate. Many experiments and surveys have found willingness to accept to be greater than willingness to pay in comparisons of this kind: we shall call this the *classic WTA/WTP comparison*. Our experiment replicates the familiar result: the ratio of geometric means is 2.13 and the difference between the distributions of valuations is overwhelmingly significant. This result is not surprising, but it gives some assurance that our experiment is picking up whatever causal factors lie behind commonly-observed differences between willingness to accept and willingness to pay.

As the entries in the first row of Table 3 highlight, the comparison between WTA_{21} and WTP_{21} is deficient in experimental control. A WTA_{21}/WTP_{21} ratio greater than unity could be evidence of loss aversion in chocolate, of loss aversion in money (combined with the falsity of NLAIB), of tactical heuristics, of Hicksian income and substitution effects, or of any combination of these factors. Our experiment was designed to allow the two kinds of loss aversion to be disentangled from each other and from these other effects.

We now turn to the relevant tests. In this section, and in the spirit of adversarial collaboration, we confine ourselves strictly to the data generated by our experiment and to the tests we planned in advance of seeing those data. Discussion, interpretation, and comparisons with the results of other experiments are postponed to Section 10.

Recall that the conjunction of prospect theory, reference-dependent theory and NLAIB, as modelled in Sections 4 and 5, implies $WTP_{21}/RWTP_{21} > 1$, $RWTA_{12}/WTA_{12} > 1$, and $RWTA_{21}/WTA_{21} = 1$; the two inequalities are manifestations of loss aversion in money. In fact, the ratio of geometric means for each of these comparisons (0.87, 0.97 and 1.02 respectively) is close to 1, and in each case the difference between the two distributions is not significant. These results give some support to the hypothesis that there is no risk aversion with respect to balanced lotteries – a distinctive prediction of prospect theory. But they give no support to the conjunction of reference-dependent theory (which predicts loss aversion in money) and NLAIB.

In addition, NLAIB implies $EG_{21}/WTP_{21} = 1$. In fact, the ratio of geometric means for this comparison is 1.75, and the two distributions of valuations are significantly different. Thus, NLAIB in its strict form is rejected. Further, the evidence from this comparison does not suggest that loss aversion is attenuated in buying tasks. However, it is possible that the EG_{21}/WTP_{21} comparison has picked up the effects of tactical heuristics and/or Hicksian effects.

We now consider the tests for loss aversion which control for attitudes to risk, tactical heuristics and Hicksian effects. There are two such tests for loss aversion in chocolate: the original WTA_{21}/EG_{21} comparison using data from treatments in which chocolates were represented by vouchers, and the WTA_{21}^*/EG_{21}^* comparison using data from treatments in which subjects were given boxes of chocolates. The first of these comparisons yields a ratio of geometric means of 1.22 and no significant difference between distributions. In the second comparison (with a larger sample size), the ratio of geometric means is 1.30 and the difference between the distributions *is* significant. In relation to NLAIB, the crucial comparison is between WTA_{12} and EG_{12}. NLAIB implies $WTA_{12}/EG_{12} = 1$. In contrast, the conjunction of reference-dependent theory and the current endowment hypothesis implies $WTA_{12}/EG_{12} > 1$, as the manifestation of loss aversion in money. In fact, the ratio of geometric means is 1.28; the difference between distributions is not significant.

The one firm conclusion we can draw from the tests described in the previous paragraph is that, after Hicksian effects and tactical heuristics have been controlled for, there is loss aversion in chocolate. We cannot reject NLAIB; but there is no evidence that positively supports either that hypothesis or the weaker hypothesis of ALAIB. The observed divergence between WTA_{12} and EG_{12}, which the current endowment hypothesis would allow us to treat as a measure of loss aversion in money, is similar in magnitude to the divergences between WTA_{21} and EG_{21} and between WTA_{21}^* and EG_{21}^*, which (according to either hypothesis) are measures of loss aversion in chocolate.

Finally, we compare the two immediate chocolate treatments with the corresponding chocolate voucher treatments.[21] The distributions of valuations in the EG_{21}^* and EG_{21} treatments are remarkably similar, while WTA_{21}^* valuations are rather greater than WTA_{21} ones (the ratio of geometric means is 1.14, and the difference between distributions is significant). This finding is consistent with the hypothesis that loss aversion is attenuated if what individuals stand to lose is a voucher rather than an immediately consumable good.

Notice, however, that the truth value of this hypothesis has no implications for our tests of NLAIB: those tests are concerned only with loss aversion *in money*.

10. Discussion

One of the most striking features of our results is the relative weakness of loss aversion effects in *all* those comparisons in which either or both of the parties to the collaboration expected such effects to show up. (The unexpected weakness of these effects is the main reason why our tests of NLAIB were less decisive than we had hoped.)

Recall that both parties agreed about the tests to be used to detect loss aversion in chocolate: the comparisons between WTA_{21} and EG_{21} (for treatments using vouchers) and between WTA_{21}* and EG_{21}* (for treatments using immediate chocolates). Even in the latter case, WTA_{21}* valuations were only 30 per cent greater than EG_{21}* ones. We cannot be as categorical about the degree of loss aversion in money, since the form in which such loss aversion shows up depends on the truth value of NLAIB. However, if NLAIB is false, loss aversion in money is picked up in the comparison between WTA_{12} and EG_{12}; and here we found a divergence between the two valuations of a little less than 30 per cent. If NLAIB is true, and given the special assumptions of the model presented in Section 5, loss aversion in money is picked up in the comparisons between WTP_{21} and $RWTP_{21}$ and between WTA_{12} and $RWTA_{12}$. No evidence of loss aversion was found in either of these comparisons.

Nevertheless, in the classic WTA/WTP comparison, we found a divergence of over 100 per cent between WTA_{21} and WTP_{21}. This suggests that our experiment has picked up *some* relatively strong causal factor – or combination of factors – which contributes to the classic WTA/WTP discrepancy that has been found in so many experiments and surveys. So what is that factor? In trying to answer this question, we go beyond the issues on which the parties reached agreement before running the experiment. We can offer two suggestions, one of which is favoured by Kahneman, the other by the UEA group.

The starting point for the first suggestion is the fact that, in our experiment, there is a marked disparity between EG_{21} and WTP_{21} valuations. In Kahneman's model, $EG_{21}/WTP_{21} = 1$ if separability does not hold (i.e. if NLAIB is true). However, if separability *does* hold, it follows from (4) and (10) that $EG_{21}/WTP_{21} = [b_2/a_2]^{1/\beta}$; under this assumption, the EG_{21}/WTP_{21} comparison picks up loss aversion in money, while the $RWTP_{21}/WTP_{21}$

comparison does not. Kahneman's tentative interpretation of the results is that, contrary to his prior expectation, subjects treated money given up in return for chocolates as a loss. On this interpretation, our experiment fails to confirm NLAIB, but it does find some evidence of strong loss aversion in money. This account would be compatible with Kahneman's *general* theory of separability if the UEA subjects were so financially constrained that they did not perceive themselves as having budget reserves. Kahneman suggests that this may have been the case, and that this may amount to an unanticipated difference between the UEA subject pool and the subject pools that he has used previously.

The members of the UEA group favour a different interpretation. Like Kahneman, they see the results as failing to confirm NLAIB. But they believe that the most controlled test for loss aversion in money is the WTA_{12}/EG_{12} comparison. Thus, they interpret the results as showing that loss aversion in chocolate and loss aversion in money are both relatively weak effects. The implication is that although each of these forms of loss aversion is a contributory cause of the classic WTA/WTP discrepancy, they cannot be the *only* causes. The UEA group suggests that a third factor is at work: subjects are using tactical or cautious heuristics which, irrespective of the response mode, generate relatively high incoming valuations and relatively low outgoing ones. On a hypothesis of this kind, we should expect relatively chocolate-loving preferences to be revealed in those money-response tasks that elicit incoming valuations (i.e. EG_{21}, WTA_{21}, $RWTA_{21}$, EG_{21}* and WTA_{21}*). Conversely, we should expect relatively money-loving preferences to be revealed both in the chocolate-response tasks (i.e. $RWTA_{12}$, WTA_{12} and EG_{12}), since those tasks elicit incoming valuations in units of chocolate, and in those money-response tasks that elicit outgoing valuations (i.e. $RWTP_{21}$ and WTP_{21}). Implicit preferences do in fact show this general pattern (see Tables 1 and 2); the classic discrepancy between WTA_{21} and WTP_{21} is part of that pattern.

So far, we have discussed the results of our experiment in isolation. However, our adversarial collaboration began from an apparent conflict between the findings of two earlier experimental investigations, those of Kahneman et al (1990) and Bateman et al (1997). We now reconsider those (and other related) findings in the light of the experiment reported in this paper.

The main results of our joint experiment are remarkably similar to those of the Bateman et al experiment, despite the differences between the two designs. Specifically: NLAIB is not supported; loss aversion in money (as measured by WTA_{12}/EG_{12}) is similar in

magnitude to loss aversion in the good (as measured by WTA_{21}/EG_{21}); both effects are quite weak; incoming valuations are markedly greater than outgoing valuations; the classic disparity between WTA_{21} and WTP_{21} appears as a strong effect. If the only data to be organized were those produced by these two experiments, it might seem natural to reject the NLAIB hypothesis, and to conclude that loss aversion in money and loss aversion in consumption goods are real but relatively weak effects. One might also conclude that there is a pervasive disparity between incoming and outgoing valuations, presumably due to tactical or cautious heuristics, which is not eliminated when valuations are elicited by multiple dichotomous choices.

It is much more difficult to reconcile the results of our joint experiment with those of Kahneman et al. The methods used to elicit valuations in the two experiments are similar in all respects except for apparently minor matters of presentation. The goods used have similar characteristics – in both cases they are low-value, non-staple private consumption goods. Yet the Kahneman et al experiments produced evidence in support of NLAIB (or, at least, in support of the hypothesis that loss aversion is much attenuated in buying tasks); and they found much stronger loss aversion in the good than was the case in our joint experiment.

We now know of further relevant experimental results. Between 1986 and 1991, Kahneman ran a series of experiments which elicited various combinations of $RWTP_{21}$, WTP_{21}, EG_{21}, WTA_{21} and $RWTA_{21}$ valuations. These experiments used essentially the same methods as were used for the corresponding treatments in our joint experiment; the 1125 subjects used in the various treatments were students at north American universities. At the time the joint experiment was designed, Kahneman believed that the results of these earlier experiments had been lost in a fire which destroyed his home in 1991. However, they were found intact in 1999, after the first stage of our joint experiment had been run. These results (and those of a further experiment which Kahneman and Nathan Novemsky ran after our joint experiment) are reported by Kahneman and Novemsky (2000). Averaging across all these experiments, Kahneman and Novemsky's computations of the ratios of median valuations are: $RWTA_{21}/WTA_{21} = 0.91$, $WTA_{21}/EG_{21} = 1.80$, $EG_{21}/WTP_{21} = 1.08$, $WTP_{21}/RWTP_{21} = 2.23$. These ratios are consistent with the results of the Kahneman et al experiments; they are consistent with Kahneman's predictions; they support NLAIB; they indicate relatively strong loss aversion in both money and consumption goods; and they show no evidence of a systematic difference between incoming and outgoing valuations. However, these data also

show a considerable degree of variation across experiments – variation that cannot be explained merely as the result of random factors, relative to a background hypothesis of no differences between subject pools. Thus, while unexplained cross-experiment differences clearly exist within the whole body of data bearing on NLAIB, it is debatable whether or not there is a particular discontinuity between, on the one hand, the two experiments conducted at UEA and, on the other, the set of experiments conducted by Kahneman in north America. Kahneman judges that there is such a discontinuity, and speculates that it is due to some difference between British and north American student populations – possibly that British students live under tighter budget constraints. The UEA group remains agnostic about whether there is a transatlantic discontinuity to be explained.

11. Conclusions

The question of whether individuals are subject to loss aversion when they spend money to buy goods is an important issue in the developing theory of reference-dependent preferences. In the context of cost-benefit analysis and contingent valuation, the corresponding question is whether willingness-to-pay (WTP) valuations of benefits do or do not pick up respondents' loss aversion with respect to money. If they do not, there is an asymmetry between WTP and willingness-to-accept (WTA) valuations – only WTA picks up any kind of loss aversion – and, if loss aversion is interpreted as a bias, that asymmetry might justify contingent valuation practitioners in using WTP as the standard of valuation. But if WTP *does* pick up loss aversion in money, WTP and WTA are symmetrical with one another; if we want a measure which abstracts from loss aversion, there is a case for using valuations of 'equivalent

The experiment reported in this paper was conducted as an adversarial collaboration between Kahneman, who predicted that loss aversion would not occur in buying tasks, and the other authors, who predicted that it would. Both parties agree that the evidence *from this experiment* favours the latter prediction, although not decisively so. However, in the light of conflicting evidence from other related experiments, the role of loss aversion in buying tasks must be considered as still an open question, on which more research is needed.

One unusual feature of our experiment was the parallel elicitation of money values of a fixed amount of a private consumption good (chocolate) *and* chocolate valuations of a fixed amount of money. This allows us isolate the effects of loss aversion from the effects of the

difference between ‘outgoing’ valuations (valuations of one thing in terms of willingness to give up another) and ‘incoming’ valuations (valuations of one thing in terms of willingness to accept another). Our results suggest that there is a tendency, additional to those induced by loss aversion, for individuals to understate their true valuations when responding to outgoing valuation tasks (such as WTP) and to overstate them in incoming tasks (such as WTA). We have suggested that this tendency may be due to subjects’ use of tactical or cautious heuristics that are well-adapted to everyday trading situations. The fact that such effects can occur in what appears to be a transparently incentive-compatible design, using multiple dichotomous choices to elicit valuations of private consumption goods, points to the difficulty of eliminating the corresponding effects in contingent valuation studies. The implication is that loss aversion in goods, loss aversion in money, and tactical and cautious heuristics may all be implicated in the familiar disparity between WTA and WTP, as found in contingent valuation.

In terms of scientific method, we believe that our work has demonstrated the value of adversarial collaboration in experimental economics. While we do not fully agree about how best to interpret our findings, we have gone a long way in narrowing down the areas of disagreement. We recommend this method to other experimental researchers as a constructive way of resolving conflicts between rival hypotheses.

Table 1: Responses to tasks with money as the response mode

	task (in = incoming, out = outgoing)						
	$RWTP_{21}$ (out)	WTP_{21} (out)	EG_{21} (in)	WTA_{21} (in)	$RWTA_{21}$ (in)	EG_{21}* (in)	WTA_{21}* (in)
valuations of 10 chocolates (units of £0.30):							
geometric mean	5.38	4.66	8.17	9.95	10.17	8.69	11.30
arithmetic mean	6.75	5.55	10.00	10.80	12.70	10.24	12.46
median	5	6	10	10	12	10	14
standard deviation	4.99	2.82	5.05	4.58	7.35	4.94	5.19
implicit preferences:							
no of subjects who:							
prefer £1	9	9	5	0	3	3	3
not clear[a]	7	4	1	1	1	8	1
prefer 10 chocolates	24	27	34	39	36	44	48
(% who prefer chocolates)	(60.0)	(67.5)	(85.0)	(97.5)	(92.3)	(80.0)	(92.3)
total	40	40	40	40	40	55	52

a. Subjects whose responses indicated that the valuation of 10 chocolates was at least £0.90 but no more than £1.20.

Table 2: Responses to tasks with chocolate as the response mode

	task (in = incoming, out = outgoing)		
	$RWTA_{12}$ (in)	WTA_{12} (in)	EG_{12} (in)
valuations of £1.00 (units of 1 chocolate):			
geometric mean	9.42	9.62	7.52
arithmetic mean	12.70	10.95	8.85
median	12.5	10	8
standard deviation	7.75	5.70	4.84
implicit preferences:			
no of subjects who:			
prefer £1	24	17	12
prefer 10 chocolates	16	23	28
(% who prefer chocolates)	(40.0)	(57.5)	(70.0)
total	40	40	40

Table 3: Comparisons of valuations

comparison	picks up loss aversion in:	picks up risk aversion in balanced lotteries	picks up tactical heuristics	picks up Hicksian effects	picks up voucher effects	ratio of geometric means	Mann-Whitney test: z-statistic
WTA_{21}/WTP_{21}	money (if NLAIB false) and chocolate	no	yes	yes	no	2.13	5.46**
WTA_{21}/EG_{21}	chocolate	no	no	no	no	1.22	0.37
$RWTA_{21}/WTA_{21}$	neither	yes	no	n.a.	no	1.02	0.93
EG_{21}/WTP_{21}	money (if NLAIB false)	no	yes	yes	no	1.75	4.23**
$WTP_{21}/RWTP_{21}$	money (if NLAIB true)	yes	no	n.a.	no	0.87	– 0.57
WTA_{12}/EG_{12}	money (if NLAIB false)	no	no	no	no	1.28	1.49
$RWTA_{12}/WTA_{12}$	money (if NLAIB true)	yes	no	n.a.	no	0.97	1.01
WTA_{21}^*/EG_{21}^*	chocolate	no	no	no	no	1.30	2.03*
EG_{21}^*/EG_{21}	neither	no	no	no	yes	1.06	0.31
WTA_{21}^*/WTA_{21}	neither	no	no	no	yes	1.14	1.82*

In final column, * denotes significance at 5 per cent level in a one-tail test; ** denotes significance at 1 per cent level.

Figure 1: Examples of tasks used in the experiment

Choice 12 [EG_{21}]	We give you £3.60	9	**or**	We give you 10 chocolates	9
Choice 12 [WTA_{21}]	You give us your 10 chocolates and take £3.60 in exchange	9	**or**	You keep your 10 chocolates	9
Choice 12 [$RWTA_{21}$]	You enter the lottery shown below: discs 1-50 / discs 51-100 You lose your 10 chocolates / You keep your 10 chocolates and win £3.60 in addition	9	**or**	You keep your 10 chocolates	9
Choice 12 [WTP_{21}]	You give us £3.60 and take take 10 chocolates in exchange	9	**or**	You do not trade	9
Choice 12 [$RWTP_{21}$]	You enter the lottery shown below: discs 1-50 / discs 51-100 You lose £3.60 / You win 10 chocolates	9	**or**	You do not enter it	9
Choice 12 [WTP_{21}]	You give us £3.60 and take take 10 chocolates in exchange	9	**or**	You do not trade	9
Choice 12 [EG_{12}]	We give you 10 chocolates	9	**or**	We give you £1.00	9
Choice 12 [WTA_{12}]	You give us your £1.00 and take 10 chocolates in exchange	9	**or**	You keep your £1.00	9
Choice 12 [$RWTA_{12}$]	You enter the lottery shown below: discs 1-50 / discs 51-100 You lose your £1.00 / You keep your £1.00 and win 10 chocolates in addition	9	**or**	You keep your £1.00	9

Notes

1. A proposed protocol for adversarial collaboration is included in Mellers, Hertwig and Kahneman (2001).

2. Bateman et al (1997), Sugden (1999) and Starmer (2000) give references to relevant experimental and survey research. See also Myagkov and Plott (1997), who find loss aversion in an experimental market setting, and Samuelson and Zeckhauser (1988) and Benartzi and Thaler (1995), who find evidence of loss aversion in 'real' markets.

3. In fact, Tversky and Kahneman propose a strictly stronger hypothesis, namely that the implication $y \succsim_r x \Rightarrow y \succ_s x$ holds for all x, y, r, s such that $y_i > x_i$, $r_i = x_i$, $y_i \geq s_i > x_i$, and $r_j = s_j$ for all $j \neq i$. However, the hypothesis we have stated is sufficient for the purposes of this paper, and the intuition behind it is easier to explain.

4. The UEA group does not claim that the current endowment hypothesis applies to *all* decision situations. For example, that hypothesis might not apply if decision-makers perceive current endowments as unfair or morally wrong.

5. The first four of these measures are discussed in more detail in Bateman et al (1997). EL_{21} and EL_{12} measures are not used in the experiment reported in this paper, but are included for completeness.

6. For example, consider a student subject with money wealth of $10,000 whose WTA for a coffee mug is $10. If her WTP for a mug increases in proportion to her wealth, the Hicksian prediction is that, given her current wealth, her WTP for the mug is $9.995. For more discussion of this point, see Sugden (1999).

7. The reasoning that leads to the predictions stated in this and the preceding paragraph is presented in more detail in Bateman et al (1997).

8. Mitchell and Carson (1989) provide an overview of the contingent valuation method. The problems caused by disparities between WTP and WTA are discussed (from different perspectives) by Hanemann (1999) and Sugden (1999).

9. Whether or not loss aversion should be interpreted as a bias in the context of valuation is an interesting issue. We view this as an open question which we do not attempt to address here.

10. In the framework of reference-dependent preference theory, this is a simplifying assumption. In that theory, preferences depend both on the absolute levels of consumption and on the reference state.

11. This analysis uses prospect theory as presented by Kahneman and Tversky (1979), or (equivalently in this case) the rank-dependent formulation proposed by Starmer and Sugden (1989). In Tversky and Kahneman's (1992) rank-dependent formulation, there are separate probability weighting functions for gains and losses.

12. In expected utility theory, in which preferences depend on final consumption states, concavity of the utility function implies risk-aversion with respect to balanced lotteries. In rank-dependent utility theory, the standard assumption that the probability weighting function $\pi(.)$ has the property $\pi(0.5) < 0.5$ is a further factor inducing risk-aversion with respect to such lotteries (Quiggin, 1993).

13. Bateman et al (1997) do not report these results in this form; instead, responses are analysed in terms of the implicit preferences that they reveal. We use median valuations here to allow comparisons with the Kahneman et al experiment.

14. Loewenstein and Adler (1995) offer some evidence that loss aversion is less when endowments are contingent on random events than when they are certain.

15. Using data from an experiment in which subjects reported certainty equivalents for lotteries with money consequences, Tversky and Kahneman (1992) fit a model of the form presented in Section 5, separately for each subject. The median value of b_2/a_2 (i.e. the index of loss aversion in money) is 2.25. The median value of β is 0.88. Thus, Kahneman expected typical values of $[b_2/a_2]^{1/\beta}$ to be in the region of $2.25^{0.88} = 2.04$. On the basis of the results reported by Kahneman et al (1990), he expected typical values of WTA_{21}/EG_{21} to be close to 2 also.

16. The second part of the experiment is reported by Cubitt, Starmer and Sugden (2000).

17. For example, if an EG valuation was being elicited in ascending order, a subject who chose the left-hand option in any given choice problem was not allowed to choose the right-hand option in a subsequent problem. If a subject's responses were mutually inconsistent, the nature of the inconsistency was explained to him, and he was asked to revise those responses. In fact, all but three of the 427 subjects responded consistently at the first attempt.

18. Kahneman's proposed protocol for adversarial collaboration includes the principle that, after the initially-agreed experiment has been run, and before the principle that binds the parties to publication comes into play, each party may propose one additional experiment 'to exploit the fount of hindsight wisdom which commonly becomes available when disliked results are obtained' (Mellers, Hertwig and Kahneman, 2001).

19. A possible objection to the use of mean valuations (whether arithmetic or geometric) is that they are sensitive to extreme values. However, our subjects rarely used the extremes of the response scales. For example, of the 307 subjects using the money response mode, only 13 recorded the lowest valuation 1 and only 9 recorded the highest valuation 26. For the 120 subjects using the chocolate response mode, the corresponding numbers were 6 and 5.

20. We use a non-parametric test because, for several treatments, the hypothesis that the distribution of valuations is normal can be rejected. The alternative Kolmogorov-Smirnoff test for differences between distributions gives very similar results to the Mann-Whitney test; the only case of disagreement is the comparison between EG_{21}* and WTA_{21}* distributions, where the Kolmogorov-Smirnoff test does not find a significant difference ($z = 0.96$).

21. In making these comparisons, it should be remembered that the subjects in the WTA_{21}* and EG_{21}* treatments were recruited separately from those in the WTA_{21} and EG_{21} treatments. We have no reason to expect any systematic differences between the two subject pools, but the possibility that such differences exist cannot be ruled out completely.

References

Arrow, Kenneth J., Robert Solow, Paul R. Portney, Edward E. Leamer, Roy Radner and E. Howard Schuman (1993). Report of the NOAA Panel on Contingent Valuation. *Federal Register* 58(10): 4601-4614.

Bateman, Ian, Alistair Munro, Bruce Rhodes, Chris Starmer and Robert Sugden (1997). A test of the theory of reference-dependent preferences. *Quarterly Journal of Economics* 112: 479-505.

Benartzi, Shlomo and Richard H. Thaler (1995). Myopic loss aversion and the equity premium puzzle. *Quarterly Journal of Economics*, 110: 73-92.

Cubitt, Robin, Chris Starmer and Robert Sugden (2000). Discovered preferences and the experimental evidence of violations of expected utility theory. Forthcoming in *Journal of Economic Methodology*.

Hanemann, W. Michael (1999). The economic theory of WTP and WTA. In Ian J. Bateman and Kenneth G. Willis (eds), *Valuing Environmental Preferences: Theory and Practice of the Contingent Valuation Method in the US, EU, and Developing Countries* (Oxford: Oxford University Press).

Kahneman, Daniel, Jack L. Knetsch and Richard H. Thaler (1990). Experimental tests of the endowment effect and the Coase theorem. *Journal of Political Economy* 98: 1325-1348.

Kahneman, Daniel and Nathan Novemsky (2000). Loss aversion in risky and riskless transactions. Mimeo, Princeton University.

Kahneman, Daniel and Amos Tversky (1979). Prospect theory: an analysis of decision under risk. *Econometrica* 47: 263-291.

Loewenstein, George and Daniel Adler (1995). A bias in the prediction of tastes. *Economic Journal* 105: 929-937.

Mellers, Barbara, Ralph Hertwig and Daniel Kahneman (2001). Do frequency representations eliminate conjunction effects? An exercise in adversarial collaboration. Forthcoming in *Psychological Science*, July 2001.

Mitchell, Robert C. and Richard T. Carson (1989). *Using Surveys to Value Public Goods: the Contingent Valuation Method*. Washington, D.C.: Resources for the Future.

Munro, Alistair and Robert Sugden (1998). A theory of general equilibrium with reference-dependent preferences. Economics Research Centre Discussion Paper no. 9803, University of East Anglia.

Myagkov, Mikhail and Charles R. Plott (1997). Exchange economies and loss exposure: experiments exploring prospect theory and competitive equilibria in market environments. *American Economic Review* 87: 801-828.

Quiggin, John (1993). *Generalised Expected Utility Theory: the Rank-Dependent Model*. Boston, Mass.: Kluwer.

Samuelson, William and Richard Zeckhauser (1988). Status quo bias in decision making. *Journal of Risk and Uncertainty* 1: 7-59.

Starmer, Chris (2000). Developments in non-expected utility theory: the hunt for a descriptive theory of choice under risk. *Journal of Economic Literature* 38: 332-382.

Starmer, Chris and Robert Sugden (1989). Violations of the independence axiom in common ratio problems: an experimental test of some competing hypotheses. *Annals of Operations Research* 19: 79-102.

Sugden, Robert (1999). Alternatives to the neo-classical theory of choice. In Ian J. Bateman and Kenneth G. Willis (eds), *Valuing Environmental Preferences: Theory and Practice of the Contingent Valuation Method in the US, EU, and Developing Countries* (Oxford: Oxford University Press).

Thaler, Richard (1980). Toward a positive theory of consumer choice. *Journal of Economic Behavior and Organization* 1: 39-60.

Tversky, Amos and Daniel Kahneman (1991). Loss aversion in riskless choice: a reference-dependent model. *Quarterly Journal of Economics* 106: 1039-1061.

Tversky, Amos and Daniel Kahneman (1992). Advances in prospect theory: cumulative representation of uncertainty. *Journal of Risk and Uncertainty* 5: 297-323.