



A Microeconomic Test of Alternative Stochastic Theories of Risky Choice

GRAHAM LOOMES, PETER G. MOFFATT AND ROBERT SUGDEN*

r.sugden@uea.ac.uk

School of Economic and Social Studies, University of East Anglia, Norwich NR4 7TJ, United Kingdom

Abstract

The random preference, Fechner (or 'white noise'), and constant error (or 'tremble') models of stochastic choice under risk are compared. Various combinations of these approaches are used with expected utility and rank-dependent theory. The resulting models are estimated in a random effects framework using experimental data from two samples of 46 subjects who each faced 90 pairwise choice problems. The best fitting model uses the random preference approach with a tremble mechanism, in conjunction with rank-dependent theory. As subjects gain experience, trembles become less frequent and there is less deviation from behaviour consistent with expected utility theory.

Keywords: risk, stochastic choice, error, expected utility theory, rank dependent theory

JEL Classification: D81, C91

Over the last twenty years, the growth of experimental research into decision making under risk, and the resulting accumulation of evidence of deviations from the predictions of expected utility theory, has led to the development of many alternative theories of choice. In almost all these theories, the choices of any given individual are taken to be fully determined by fixed preferences. Recently, however, there has been a revival of interest among economic theorists in modelling the stochastic element in decision making. In this paper, by means of microeconomic analysis of experimental data, we compare the explanatory power of three of the most prominent models of stochastic choice to emerge in recent discussions.

The new interest in stochastic choice is partly a response to experimental evidence of what appears to be random variation in individuals' decisions. For example, in a number of tightly-controlled experiments in which subjects have confronted exactly the same pairwise choice problem on two occasions, separated only by a short time interval, the proportion who choose differently in the two cases has often been found to be of the order of 20 to 30 per cent.¹ Such evidence suggests that stochastic variation is an essential feature of decision-making behaviour, and not merely the outward manifestation of changes in the values of unobserved variables.

* Corresponding author.

Interest has also been generated by several studies which have attempted to find a ‘best buy’ from the set of deterministic theories of choice under risk by estimating alternative models using experimental data.² Such estimation exercises require the specification of stochastic ‘error’ mechanisms. Different researchers have assumed fundamentally different mechanisms, while comparing the performance of the same deterministic theories. But choosing between alternative stochastic specifications is not just a matter of economic convenience: whether a given pattern of behaviour is classified as a *systematic*—that is, non-random—deviation from the predictions of a given deterministic theory can depend on the stochastic specification that is used.³ We suggest that the ‘best buy’ problem is better posed as a choice between alternative *stochastic* theories. Thus, alternative stochastic specifications, just like the deterministic theories to which they are added, should be tested for their explanatory power.

Three alternative approaches to the modelling of stochastic choice have been discussed in recent papers. The ‘constant error’ or ‘tremble’ approach proposes that each individual normally behaves according to some deterministic theory of choice, but subject to some constant probability of a lapse of concentration; if such a lapse occurs, choices are random. The ‘Fechner’ or ‘white noise’ approach proposes that the individual maximizes some form of utility function which includes a stochastic disturbance term. The ‘random preference’ approach proposes that, for each individual, there is a set of alternative preference relations; facing any particular decision problem, the individual acts on one of these preference relations, selected at random. These approaches embody different explanations of stochastic variation—explaining it as the product either of error in the execution of preferences (the constant error approach), or of error in the formation of preferences (the Fechner approach), or of uncertainty about preferences (the random preference approach).

Our principal objective in this paper is to compare the explanatory power of these three approaches in conjunction both with expected utility theory and with its most prominent rival, rank-dependent theory. In principle, this analysis could be extended to include other non-expected utility theories; our main reason for considering only these two theories is to keep our investigation within manageable bounds. However, for reasons that we shall explain later in the paper, we think it unlikely that other known non-expected utility theories would perform markedly better than rank-dependent theory in explaining our data.

Our analysis also breaks new ground by investigating how far the behaviour of experimental subjects is modified in the light of experience of and familiarity with a given type of task. Critics of experimental economics sometimes suggest that experimentally-observed deviations from deterministic expected utility theory represent some combination of systematic and random errors, made by naive respondents who are unfamiliar with the decision tasks they confront. Implicit in this suggestion is the notion that, as individuals gain experience of the decision-making environment, both kinds of error will tend to become less frequent. Our analysis allows us to separate the deterministic and stochastic elements in choice, and to track how each element changes as individuals become more experienced.

1. Three models of stochastic choice

We consider decisions between pairs of risky lotteries or prospects. Let $X = \{x_0, \dots, x_n\}$ be a set of money *consequences* with $x_0 < x_1 < \dots < x_n$. A *prospect* is a probability distribution over those consequences; typical prospects will be denoted by $\mathbf{p} \equiv (p_0, \dots, p_n)$ and $\mathbf{q} \equiv (q_0, \dots, q_n)$. A *decision problem* is a pair of prospects $\{\mathbf{p}, \mathbf{q}\}$.

A *core theory* is a function $V(\cdot, \cdot)$ which assigns a real-valued index of *subjective value* to every pair (\mathbf{p}, \mathbf{z}) where \mathbf{p} is a prospect and $\mathbf{z} \in Z$ is a vector of *person-specific parameters*. The content of Z depends on the theory. For example, in expected utility (EU) theory, Z is the set of all vectors $\mathbf{z} = (u_0, \dots, u_n)$ of von Neumann–Morgenstern utility indices satisfying the monotonicity condition $u_0 < u_1 < \dots < u_n$; we then have:

$$V(\mathbf{p}, \mathbf{z}) = \sum_i p_i u_i. \quad (1)$$

Later in this paper we will also be concerned with *rank-dependent* (RD) theory (Quiggin, 1982). This can be represented as a deterministic theory in which Z is the set of all vectors $\mathbf{z} = (u_0, \dots, u_n; a_1, \dots, a_m)$ where the u_i terms are interpreted as for EU, and the a_j terms are the parameters of a *probability transformation function* $\pi: [0, 1] \rightarrow [0, 1]$ which is monotonically increasing, with $\pi(0) = 0$ and $\pi(1) = 1$. Then:

$$V(\mathbf{p}, \mathbf{z}) = \sum_i w_i u_i, \quad (2)$$

where the w_i terms (the *decision weights*) are given by

$$w_0 \equiv \pi(p_0), \quad (3a)$$

and

$$w_i \equiv \pi\left(\sum_{j=0}^i p_j\right) - \pi\left(\sum_{j=0}^{i-1} p_j\right) \quad (i = 1, \dots, n). \quad (3b)$$

Person-specific parameters are used to allow a theory to account for interpersonal variation in preferences. When a core theory is interpreted deterministically, it is assumed that for any given person, the values of these parameters are fixed, and that decisions maximize subjective value.

We now consider how a stochastic component might be grafted onto a core theory. As an expositional device, it is helpful to assume that some given core theory provides a normatively compelling account of rational choice, and to imagine the process of choosing between two prospects \mathbf{p}, \mathbf{q} being split into three stages:⁴ preference selection, calculation, and action. In the first stage, the decision-maker identifies her current preferences, represented by a particular vector \mathbf{z} . In the second stage, she calculates the values of $V(\mathbf{p}, \mathbf{z})$ and $V(\mathbf{q}, \mathbf{z})$ and resolves to choose whichever prospect has the higher subjective value. In the third stage, she implements her choice (say, by pressing a key on a keyboard). Randomness could enter the choice process at any of these stages.

The first possibility is that randomness enters at the preference selection stage: the individual is uncertain about her own preferences (while being certain that they have whatever *general* properties are required by the core theory). This can be modelled by assuming that for any given individual, there is a probability distribution over Z . The vector of person-specific parameters used for each decision problem is a random drawing $\hat{\mathbf{z}}$ from this distribution. The choice between \mathbf{p} and \mathbf{q} is then determined by the sign of $V(\mathbf{p}, \hat{\mathbf{z}}) - V(\mathbf{q}, \hat{\mathbf{z}})$.⁵ This stochastic specification is the *random preference* model. This model, with EU as the core theory, is due to Becker, DeGroot and Marschak (1963). It has recently been reconsidered and generalized by Loomes and Sugden (1995).

A second possibility is that randomness enters at the calculation stage: we might imagine that the individual has ‘true’ preferences, represented by the core theory in conjunction with some fixed \mathbf{z} , but that her calculations of subjective values are subject to random error. This error can be modelled by adding a stochastic term to the difference in subjective value, so that the choice between \mathbf{p} and \mathbf{q} is determined by the sign of $V(\mathbf{p}, \mathbf{z}) - V(\mathbf{q}, \mathbf{z}) + \varepsilon$, where ε is a continuous random variable, symmetrically distributed around zero. This stochastic specification is the *Fechner model* (named after Fechner, 1860/1966). This model, like the first, was developed by Becker, DeGroot and Marschak (1963); it has recently been used by Hey and Orme (1994) and Carbone and Hey (1994) in estimating a wide range of alternative models of choice under risk.

The final possibility is that randomness enters at the action stage. We might imagine that the individual sometimes fails properly to understand a decision problem, or suffers a lapse of concentration or *tremble*, with the result that her action is unconnected with her preferences. This kind of error can be modelled by assuming that, for all decision problems $\{\mathbf{p}, \mathbf{q}\}$, there is some fixed probability $0 < \omega < 1$ that the decision is made at random. With probability $1 - \omega$, the decision is determined by the sign of $V(\mathbf{p}, \mathbf{z}) - V(\mathbf{q}, \mathbf{z})$. This stochastic specification will be called the *constant error* model. It was developed by Harless and Camerer (1994), who use it to compare the power of various theories of choice under risk.⁶

There is no obvious reason to assume that only one of these forms of randomness is present. In principle, these models of stochastic variation can be combined with one another. Moreover, there is a practical econometric reason for considering hybrid models. The random preference model implies that, in certain classes of decision problems, choice is *not* stochastic. Take any core theory, and consider any decision problem $\{\mathbf{p}, \mathbf{q}\}$ such that, in that theory, $V(\mathbf{p}, \mathbf{z}) > V(\mathbf{q}, \mathbf{z})$ for *all* \mathbf{z} . For example, this is the case if the core theory is EU or RD and if \mathbf{p} stochastically dominates \mathbf{q} . Then the random preference model implies that \mathbf{p} is chosen with probability one. Thus, with EU or RD as the core theory, a single violation of dominance (i.e., a case in which a stochastically dominated prospect is chosen) is sufficient to refute the random preference model. In any experiment with many subjects, each of whom faces many choice problems, there are almost certain to be *some* instances of the kinds of misunderstandings, confusions and errors that the constant error model represents. If any of these trembles creates a violation of dominance, an EU- or RD-based random preference model cannot be estimated. An obvious solution to this problem is to add a tremble mechanism to the random preference model.

The Fechner model does not confront this problem so directly: for *every* decision problem $\{\mathbf{p}, \mathbf{q}\}$, the Fechner model implies that each prospect is chosen with non-zero probability. However, if experimental subjects sometimes make the kind of errors that are best modelled as trembles, adding a tremble mechanism to a Fechner model will improve its specification.

2. Existing evidence

Hey and Orme (1994), Carbone and Hey (1994) and Harless and Camerer (1994) compare the predictive power of alternative theories of choice under risk; the first two papers use the Fechner approach while the third uses the constant error approach. Although these three papers address a common problem, it is difficult to use them to compare the explanatory power of the two models of stochastic variation. Hey and Orme's and Carbone and Hey's strategy is to present parameterized forms of the various deterministic theories, to add a Fechner error term, and then to estimate each of these models, separately for each of a set of experimental subjects. In contrast, Harless and Camerer do not estimate person-specific parameters: they test predictions which are independent of those parameters, using data which are aggregated across subjects.

A small number of recent experimental studies have compared the explanatory power of different stochastic specifications. Ballinger and Wilcox (1997) test a number of implications of the Fechner and constant error models which are independent of the core theory with which those specifications are combined. Their data reject the constant error model but not the Fechner model.⁷ Carbone (1998) compares all three stochastic specifications, using EU as the core theory and estimating parameterized models subject by subject. The constant error model performs least well, but neither of the other two models emerges as obviously better than the other.

In the light of these investigations, it seems clear that the constant error model is inadequate. Given the extreme simplicity of that model, this finding is perhaps not surprising; and it does not rule out the possibility that a tremble mechanism might be *part of* a successful hybrid model of stochastic choice. Little is yet known about the relative performances of the Fechner and random preference models, particularly in conjunction with non-EU theories.

3. Data

The experiment which generated our data set is described in detail by Loomes and Sugden (1998). Here, we merely summarize its main features.

There were 92 subjects, recruited from the student population of the University of York. Each subject began the experiment by facing a set of 45 pairwise choice problems, presented in random order. After a short break, the same 45 pairwise choices were presented again, in a different random order. At the end of the experiment, one problem

was selected at random for each subject; the subject then played out the gamble described by the prospect she had chosen in that problem, and was paid accordingly.

For each subject, the 45 choices involved different probability mixes of three fixed consequences. Subjects were divided at random between two subsamples, the ‘£30 group’ and the ‘£20 group.’ For the £30 group, the consequences were $x_0 = £0$, $x_1 = £10$, $x_2 = £30$; for the £20 group, they were $x_0 = £0$, $x_1 = £10$, $x_2 = £20$. The parameters of the two sets of 45 problems are presented in Tables 1a and 1b. The final three columns of these tables show the values of the parameters $d_1 \equiv p_1 - q_1$ and $d_2 \equiv p_2 - q_2$, and the quantity $-d_1/d_2$, whose relevance will become apparent.

For each subsample, each of the *non-dominance* Problems 1–40 requires a choice between a pair of prospects **p**, **q**, where $d_1 + d_2 < 0$ and $d_2 > 0$; thus, neither prospect stochastically dominates the other, and **p** is unambiguously riskier than **q**. In EUT, a person’s preference ranking of any such pair is determined by her degree of risk aversion and by the value of $-d_1/d_2$: where the subject’s attitude to risk parameter u , to be defined in Section 5, exceeds $-d_1/d_2$, **p** is chosen; otherwise, **q** is chosen. These 40 problems can be subdivided into five sets of eight: Problems 1–8, 9–16, 17–24, 25–32, and 33–40. For a given subsample, the value of $-d_1/d_2$ is constant for all problems in any one of these sets, but increases from set to set, the lowest values of $-d_1/d_2$ being associated with the highest-numbered problems. Thus, we might expect **p** to be more attractive relative to **q** in the higher-numbered sets. Each of the *dominance* Problems 41–45 requires a choice between a pair of prospects **p**, **q** where $d_1 + d_2 \geq 0$ and $d_2 \geq 0$, with a strict inequality in at least one case. Thus, in these problems, **p** stochastically dominates **q**.

The full data set is presented in Tables 2a and 2b. In these tables, each row represents a problem, and each column represents a subject. Problems 46, . . . , 92 are identical with problems 1, . . . , 45, but represent those problems when faced for the second time. A value of 1 indicates that **p** was chosen, a value of 0 that **q** was chosen.

Casual inspection reveals that, in non-dominance problems, **p** choices become more frequent as $-d_1/d_2$ decreases, and that in all dominance problems, **q** choices are very infrequent (there are only 13 violations of dominance in 920 decisions). Another feature of the data evident from Tables 2a and 2b is the enormous variation between subjects within each subsample: some subjects never choose **p** in non-dominance problems; others choose **p** nearly every time. Clearly, subjects differ greatly in their attitudes to risk. Summing over all subjects and all non-dominance problems, the *reversal rate* (that is, the relative frequency of cases in which a subject’s first and second responses to a problem are different) is 0.183. This rate is a little lower than in some other similar experiments (see the introduction), but it still suggests a considerable degree of within-subject stochastic variation.

4. Regularities in the data

The experiment generated data for certain specific non-parametric hypothesis tests. These tests are described in Loomes and Sugden (1998). Briefly, these tests reject some implications of the constant error model that are independent of the core theory. They also

Table 1a. Parameters of problems: £20 group

Problem	p_0	p_1	p_2	q_0	q_1	q_2	d_1	d_2	$-d_1/d_2$
1	0.15	0.00	0.85	0.00	0.25	0.75	-0.25	0.10	2.5
2	0.30	0.00	0.70	0.15	0.25	0.60	-0.25	0.10	2.5
3	0.30	0.00	0.70	0.00	0.50	0.50	-0.50	0.20	2.5
4	0.15	0.25	0.60	0.00	0.50	0.50	-0.25	0.10	2.5
5	0.15	0.75	0.10	0.00	1.00	0.00	-0.25	0.10	2.5
6	0.60	0.00	0.40	0.00	1.00	0.00	-1.00	0.40	2.5
7	0.60	0.00	0.40	0.15	0.75	0.10	-0.75	0.30	2.5
8	0.90	0.00	0.10	0.75	0.25	0.00	-0.25	0.10	2.5
9	0.10	0.00	0.90	0.00	0.20	0.80	-0.20	0.10	2.0
10	0.50	0.00	0.50	0.10	0.80	0.10	-0.80	0.40	2.0
11	0.50	0.00	0.50	0.00	1.00	0.00	-1.00	0.50	2.0
12	0.10	0.80	0.10	0.00	1.00	0.00	-0.20	0.10	2.0
13	0.70	0.00	0.30	0.50	0.40	0.10	-0.40	0.20	2.0
14	0.70	0.00	0.30	0.40	0.60	0.00	-0.60	0.30	2.0
15	0.50	0.40	0.10	0.40	0.60	0.00	-0.20	0.10	2.0
16	0.90	0.00	0.10	0.80	0.20	0.00	-0.20	0.10	2.0
17	0.10	0.00	0.90	0.00	0.25	0.75	-0.25	0.15	1.67
18	0.40	0.00	0.60	0.10	0.75	0.15	-0.75	0.45	1.67
19	0.40	0.00	0.60	0.00	1.00	0.00	-1.00	0.60	1.67
20	0.10	0.75	0.15	0.00	1.00	0.00	-0.25	0.15	1.67
21	0.70	0.00	0.30	0.60	0.25	0.15	-0.25	0.15	1.67
22	0.70	0.00	0.30	0.50	0.50	0.00	-0.50	0.30	1.67
23	0.60	0.25	0.15	0.50	0.50	0.00	-0.25	0.15	1.67
24	0.85	0.00	0.15	0.75	0.25	0.00	-0.25	0.15	1.67
25	0.10	0.00	0.90	0.00	0.30	0.70	-0.30	0.20	1.5
26	0.40	0.00	0.60	0.20	0.60	0.20	-0.60	0.40	1.5
27	0.40	0.00	0.60	0.10	0.90	0.00	-0.90	0.60	1.5
28	0.20	0.60	0.20	0.10	0.90	0.00	-0.30	0.20	1.5
29	0.60	0.00	0.40	0.50	0.30	0.20	-0.30	0.20	1.5
30	0.60	0.00	0.40	0.40	0.60	0.00	-0.60	0.40	1.5
31	0.50	0.30	0.20	0.40	0.60	0.00	-0.30	0.20	1.5
32	0.80	0.00	0.20	0.70	0.30	0.00	-0.30	0.20	1.5
33	0.10	0.00	0.90	0.00	0.50	0.50	-0.50	0.40	1.25
34	0.20	0.00	0.80	0.10	0.50	0.40	-0.50	0.40	1.25
35	0.20	0.00	0.80	0.00	1.00	0.00	-1.00	0.80	1.25
36	0.10	0.50	0.40	0.00	1.00	0.00	-0.50	0.40	1.25
37	0.35	0.25	0.40	0.25	0.75	0.00	-0.50	0.40	1.25
38	0.40	0.00	0.60	0.25	0.75	0.00	-0.75	0.60	1.25
39	0.40	0.00	0.60	0.35	0.25	0.40	-0.25	0.20	1.25
40	0.60	0.00	0.40	0.50	0.50	0.00	-0.50	0.40	1.25
41	0.00	0.25	0.75	0.00	0.30	0.70	-0.05	0.05	—
42	0.55	0.20	0.25	0.65	0.15	0.20	0.05	0.05	—
43	0.80	0.00	0.20	0.85	0.00	0.15	0.00	0.05	—
44	0.10	0.75	0.15	0.15	0.75	0.10	0.00	0.05	—
45	0.70	0.30	0.00	0.75	0.25	0.00	0.05	0.00	—

Table 1b. Parameters of problems: £30 group

Subjects in the £30 group faced problems with the same probabilities as those given in Table 1a, apart from problems 33–40, which were as follows:

Problem	p_0	p_1	p_2	q_0	q_1	q_2	d_1	d_2	$-d_1/d_2$
33	0.10	0.00	0.90	0.00	0.40	0.60	-0.40	0.30	1.33
34	0.25	0.00	0.75	0.10	0.60	0.30	-0.60	0.45	1.33
35	0.25	0.00	0.75	0.00	1.00	0.00	-1.00	0.75	1.33
36	0.10	0.60	0.30	0.00	1.00	0.00	-0.40	0.30	1.33
37	0.50	0.20	0.30	0.40	0.60	0.00	-0.40	0.30	1.33
38	0.55	0.00	0.45	0.40	0.60	0.00	-0.60	0.45	1.33
39	0.55	0.00	0.45	0.50	0.20	0.30	-0.20	0.15	1.33
40	0.70	0.00	0.30	0.60	0.40	0.00	-0.40	0.30	1.33

reject some implications of the random preference and Fechner error models, given the assumption that the core theory is EU. These negative results are broadly consistent with the results of Ballinger and Wilcox (1997) and Carbone (1998), which were outlined in Section 2.

The non-parametric methods used by Loomes and Sugden do not allow any formal analysis of how these models failed. However, by using descriptive statistics, Loomes and Sugden identify two apparent regularities in the data which are inconsistent with all three stochastic models, when combined with EU.

The first regularity is the *bottom-edge effect*: for non-dominance pairs $\{\mathbf{p}, \mathbf{q}\}$, for any given value of $-d_1/d_2$, the frequency of \mathbf{p} choices is markedly greater when $q_2 = 0$ than when $q_2 > 0$. That is, \mathbf{p} choices are more frequent when \mathbf{q} is located on the bottom edge of the Marschak-Machina triangle than when it is located anywhere else.⁸ This effect has been found in other experiments. Summarizing the results of several independent experiments, Harless and Camerer (1994, p. 1285) conclude that EU predicts well when decision problems involve pairs of prospects that have the same support (i.e. when both are in the interior of the triangle), but poorly when the support is different (i.e. when at least one prospect is on an edge of the triangle). The implication is that, if we are to test the random preference and Fechner models, we need to combine them with some non-EU core theory which is consistent with the bottom-edge effect.

The second regularity is common to these data and to the data of Hey and Orme (1994) and Ballinger and Wilcox (1997).⁹ In all three experiments, subjects faced the same decision problems twice. In each case, the second set of responses are more risk averse than the first set, and this difference is statistically significant. This finding cannot be accommodated by any of the stochastic models as they have been proposed so far; it requires a model in which experience can have a systematic effect.

The present study uses parametric maximum likelihood estimation methods to investigate the unresolved issues outlined above. We compare the performance of the Fechner and random preference models, not only for the benchmark case in which the core theory is EU, but also for the case of RD. Our reasons for using RD, rather than any other non-

EU theory, will be explained in Section 5. The models we estimate allow us to separate the effects of experience from stochastic variation.

Inevitably, the use of parametric methods requires us to assume specific functional forms for the relationships within our models. The functional forms we used were selected after a good deal of experimentation with alternatives in the pursuit of goodness of fit and parsimony. Except in a few particularly significant cases, space constraints prevent us from reporting this work. We found no evidence to suggest that our main qualitative conclusions are sensitive to our choice of specifications.

5. Stochastic models based on expected utility theory

We begin by considering the implications of *deterministic* EU for the decisions of any given subject in our experiment. As explained in Section 1, we define a vector \mathbf{u} of von Neumann-Morgenstern utility indices. We adopt the normalization¹⁰ $u_0 = 0$, $u_1 = 1$, and define $u \equiv u_2$. Notice that all relevant properties of the utility function are captured by the single value u . Since we assume that utility is monotonically increasing in money, we impose the restriction $u > 1$. The value of u then represents the individual's attitude to risk, higher values of u corresponding with lower degrees of risk aversion; the individual is risk-averse, risk-neutral or risk-loving according to whether u is less than, equal to, or greater than x_2/x_1 . EU preferences are given by:

$$\mathbf{p}(\succsim) \mathbf{q} \iff d_1 + d_2 u (\succsim) 0 \quad (4)$$

We now consider various models which add a stochastic component to (4). In this Section, we consider the choices of a given individual. The problems faced by this individual will be indexed by t and the total number of problems will be denoted by T .

5.1. The random preference model

In this model, the attitude-to-risk parameter u is a random variable satisfying the restriction $u > 1$. We shall assume that $(u - 1)$ follows a lognormal distribution. Thus, letting m be the median value of u :

$$\ln(u - 1) \sim N[\ln(m - 1), \sigma_u^2] \quad (5)$$

Let $\{\mathbf{p}_t, \mathbf{q}_t\}$ be the pair of options offered in problem t , and let d_{1t} and d_{2t} be the probability differences, as defined in Section 3. Let y_t be a variable taking the value 1 if \mathbf{p}_t is chosen from $\{\mathbf{p}_t, \mathbf{q}_t\}$ and 0 if \mathbf{q}_t is chosen. For non-dominance problems, (4) and (5) imply:

$$\text{pr}(y_t = 1 | m) = \text{pr}(d_{1t} + d_{2t}u > 0 | m) = \Phi \left[\frac{\ln\left(-\frac{d_{2t}}{d_{1t} + d_{2t}}\right) + \ln(m - 1)}{\sigma_u} \right] \quad (6)$$

where $\Phi(\cdot)$ is the standard normal distribution function. Clearly, $\text{pr}(y_t = 1) = 1$ for dominance problems.

As noted in Section 1, the random preference model breaks down if any violations of dominance are observed. This problem can be avoided by adding a tremble mechanism. Let ω be a parameter representing the probability that any choice is made at random ($0 < \omega < 1$). Then, for non-dominance problems:

$$\text{pr}(y_t = 1|m) = (1 - \omega)\Phi\left[\frac{\ln\left(-\frac{d_{2t}}{d_{1t} + d_{2t}}\right) + \ln(m - 1)}{\sigma_u}\right] + \frac{\omega}{2} \quad (7)$$

For dominance problems, $\text{pr}(y_t = 1) = 1 - \omega/2$. Theoretical and practical issues relating to the estimation of the tremble parameter are discussed in detail in a more general context by Moffatt and Peters (2001).

5.2. The Fechner model

In the Fechner model, u is a fixed parameter, and stochastic variation results from an error term ε , applied additively to (4), so that \mathbf{p}_t is chosen if $d_{1t} + d_{2t}u + \varepsilon > 0$. Following Hey and Orme (1994), we assume that ε follows a normal distribution with mean zero. The particular distribution we specify is $\varepsilon \sim N[0, (\sigma_\varepsilon \ln(u))^2]$ where σ_ε is a parameter.

For a model of a single individual, it would be sufficient to represent the variance of ε by a single parameter. Our specification is designed to apply across subjects; the value of σ_ε will be held constant while u will be subject-specific. The reason why the parameter u appears in the expression for the variance is that we need to explain the behaviour of the considerable number of subjects who chose the safer option \mathbf{q} in all or almost all non-dominance problems. The only way that the responses of such a subject can be explained in the context of the Fechner model is by inferring both that the value of the subject's u parameter is close to one, and that the variance of ε is close to zero. This is made possible by our specification. In interpreting this specification, it must be remembered that the variance of ε is expressed in units of utility, and that the normalization of utility is arbitrary. Any ranking of subjects by error variance is conditional on the normalization used. Thus, it would be a mistake to interpret our model as implying that more risk-averse individuals are in any real sense less error-prone.

With this distributional assumption, we have, for all problems:

$$\text{pr}(y_t = 1|u) = \text{pr}(d_{1t} + d_{2t}u + \varepsilon > 0|u) = \Phi\left[\frac{d_{1t} + d_{2t}u}{\sigma_\varepsilon \ln(u)}\right] \quad (8)$$

5.3. The Fechner model with trembles

The model described by (8) can be estimated whether or not there are violations of dominance, and so it is not essential to add a tremble mechanism. Nevertheless, the

specification might be improved by such an addition. The Fechner model with trembles is:

$$\text{pr}(y_t = 1|u) = (1 - \omega)\Phi\left[\frac{d_{1t} + d_{2t}u}{\sigma_\varepsilon \ln(u)}\right] + \frac{\omega}{2} \quad (9)$$

with $0 < \omega < 1$.

5.4. The effect of experience

The index t identifies the $\{\mathbf{p}, \mathbf{q}\}$ pair under consideration and corresponds with the labelling of problems in Tables 1a and 1b. However, the actual order in which the problems were faced was determined randomly for each subject. For a given subject facing problem t , let τ_t be the number of problems she has faced (including problem t). Thus, if $\tau_1 = 17$, problem 1 is the 17th problem that the subject faces. The effect of experience can then be modelled by allowing $\text{pr}(y_t = 1)$ to depend on τ_t . In principle, any of the parameters m , σ_u , u , σ_ε and ω might vary with τ . In the specification search we investigated all these possibilities, but we found significant experience effects only for ω . We therefore chose to model experience through its effect on ω , as follows:

$$\omega_t = \omega_0 \exp(\omega_1 \tau_t). \quad (10)$$

The parameter ω_0 represents the tremble probability at the start of the experiment, and ω_1 indicates how this probability changes with experience. If trembles are interpreted as misunderstandings, then since we might expect misunderstandings to become less frequent with experience, we would expect a negative value for ω_1 . On the other hand, if trembles are interpreted as lapses of concentration, then it is possible that they could become *more* frequent as subjects become tired or bored, implying a positive value for ω_1 .

6. Stochastic models based on rank-dependent theory

We chose to study RD as a representative non-EU theory. In part, we made this choice in recognition of the prominence of RD in the literature: it is probably the most widely-used non-EU theory. But we were also influenced by the properties of our data.

As explained in Section 4, there is an apparently systematic bottom-edge effect in our data. If a core theory is to induce a *general* bottom-edge effect, it must have the property that, for all q_1 and \mathbf{z} , the value of $V(\mathbf{q}, \mathbf{z})$ is highly sensitive to changes in the value of q_2 when that value is equal to or close to zero. Specifically, low but non-zero values of q_2 , i.e., the probability of receiving the best consequence, must be ‘overweighted.’ The most direct way of incorporating such an effect into an EU-like theory is to allow objective probabilities to be transformed into subjective decision weights, and to make this

transformation responsive to the ranking of consequences (so that the ‘best consequence’ has a distinct status). RD has exactly these properties.¹¹

In RD, the bottom-edge effect is generated if $\pi(p) < p$ for values of p that are close to but strictly less than 1. In the RD literature, it is conventional to assume that the probability transformation function has this property, along with the mirror-image property that $\pi(p) > p$ for values of p close to but strictly greater than 0. The latter property implies that low but non-zero values of $1 - q_1 - q_2$, i.e., the probability of receiving the worst consequence, are also overweighted. This generates a *side-edge effect*: for any given value of $-d_1/d_2$, there is a tendency for the frequency of \mathbf{q} choices to be greater when $q_1 + q_2 = 1$ than when $q_1 + q_2 < 1$. However, in the general form of RD (as presented in Section 1), a bottom-edge effect can occur in the absence of a side-edge effect, and vice versa.¹²

Analogously with our treatment of EU, we begin by considering the implications of deterministic RD for a given subject. For estimation purposes, we need to specify a functional form for the probability transformation function.

Our specification search led us to the following simple functional form:

$$\pi(p) = (1 - b)p \quad (p < 1); \quad (11a)$$

$$\pi(1) = 1; \quad (11b)$$

with $0 \leq b < 1$. With $b = 0$, this model reduces to EU. Positive values of b imply the overweighting of the probability of the best outcome in each lottery (i.e., the decision weight is greater than the probability), and thus induce a bottom-edge effect.

This specification can be generalized to allow a side-edge effect by introducing a second parameter $0 \leq a < 1$ and setting $\pi(0) = 0$, $\pi(1) = 1$, and $\pi(p) = a + (1 - a - b)p$ for $0 < p < 1$. However, when we experimented with such a model, our estimates of a were not significantly different from zero.

Now consider any prospect \mathbf{p} defined for the set of consequences $\{x_0, x_1, x_2\}$. Using the notation developed in Section 1, and using the same normalization of the utility function as in Section 4, the subjective value of \mathbf{p} can be written as $V(\mathbf{p}, u, b)$. Combining (2), (3a), (3b), (11a), and (11b), $V(\mathbf{p}, u, b)$ can conveniently be expressed as:

$$V(\mathbf{p}, u, b) = (1 - b)(p_1 + p_2u) + bu \quad \text{if } p_2 > 0; \quad (12a)$$

$$V(\mathbf{p}, u, b) = (1 - b)(p_1 + p_2u) + b \quad \text{if } p_2 = 0. \quad (12b)$$

Now consider any choice $\{\mathbf{p}_t, \mathbf{q}_t\}$ in the experiment. Noting that there are no problems in which $p_{2t} = 0$ and $q_{2t} > 0$,¹³ we define a binary variable E_t which takes the value one if $p_{2t} > 0$ and $q_{2t} = 0$, and zero otherwise. Then, using (12a) and (12b), for every problem in the experiment:

$$V(\mathbf{p}_t, u, b) - V(\mathbf{q}_t, u, b) = (1 - b)(d_{1t} + d_{2t}u) + b(u - 1)E_t. \quad (13)$$

If $b > 0$, the second term on the right-hand side of (13) represents the bottom edge effect.

Equation (13) is the deterministic core of our RD model. We now consider alternative ways of combining it with a stochastic specification.

If we use the stochastic specification given by the *random preference model*, and add a tremble mechanism, we arrive at the equation:

$$\text{pr}(y_t = 1|m) = (1 - \omega)\text{pr}[(1 - b)(d_{1t} + d_{2t}u) + b(u - 1)E_t > 0] + \omega/2. \quad (14)$$

As when the core theory is EU, we treat u as a random variable, with the distribution (5). To allow for the effect of experience, ω depends on τ as in (10). To simplify the task of estimation, we treat b as deterministic. However, we allow b to vary with τ , so that our estimations can pick up any tendency for the bottom-edge effect to decay with experience. Our specification is:

$$b_t = b_0 \exp(b_1 \tau_t), \quad (15)$$

with $0 \leq b_0 < 1$, $b_1 \leq 0$.

If we use the *Fechner model without trembles*, choice probabilities are given by:

$$\text{pr}(y_t = 1|u) = \text{pr}[(1 - b_t)(d_{1t} + d_{2t}u) + b_t(u - 1)E_t + \varepsilon > 0]. \quad (16)$$

If we use the *Fechner model with trembles*, this equation becomes:

$$\text{pr}(y_t = 1|u) = (1 - \omega_t)\text{pr}[(1 - b_t)(d_{1t} + d_{2t}u) + b_t(u - 1)E_t + \varepsilon > 0] + \omega_t/2. \quad (17)$$

As in our EU-based models, we assume that ε follows the normal distribution $N[0, (\sigma_\varepsilon \ln(u))^2]$; and as in the random preference model, ω and b are assumed to depend on τ , according to (10) and (15).

We recognize that our specification of the probability transformation function by (11a) and (11b) is not conventional: it is more usual to use a functional form which makes $\pi(\cdot)$ smooth and continuous. In deference to this practice, we also report an estimation based on the following alternative specification:

$$\pi(p) = \frac{(1 - \alpha_t)p^{1-\delta_t}}{(1 - \alpha_t)p^{1-\delta_t} + (1 - p)^{1-\delta_t}} \quad (18)$$

$$\alpha_t = \alpha_0 \exp(\alpha_1 \tau_t) \quad (19)$$

$$\delta_t = \delta_0 \exp(\delta_1 \tau_t) \quad (20)$$

We shall call this the RDS model ('S' standing for 'smoothed'). The functional form (18) is discussed by Prelec (1998). It defines a family of smooth functions with the common properties that $\pi(p) > p$ at low values of p (inducing a side-edge effect), and $\pi(p) < p$

at high values (inducing a bottom-edge effect). As we shall show, the RDS model does not fit the data as well as our basic RD model.

7. The models with many subjects

In the discussion so far, we have modelled the behaviour of a single subject. It would be possible to follow the strategy of Hey and Orme (1994), Carbone and Hey (1994) and Carbone (1998), and to estimate the various models separately for each subject. However, we prefer to pool the observations from all subjects, and to be as parsimonious as possible in the use of subject-specific parameters. As with most economics experiments, our subject pool is not a systematic sample of any economically significant population. Thus, findings about the relative frequency of different kinds of behaviour in the subject pool are of limited value. Our view is that experimental findings are of most interest when they reveal systematic tendencies in the behaviour of subjects *in general*, and when those tendencies are susceptible to explanation by general theories.

As we pointed out in Section 4, our subjects appear to differ greatly in their attitudes to risk. The implication seems to be that any model which is to fit our data needs at least one subject-specific parameter to represent attitude to risk. In the Fechner models, it is natural to make the value of u subject-specific, and we do this by assuming $\ln(u - 1) \sim N(\mu, \eta^2)$. In the random preference models, the corresponding move is to make m (i.e., the median value of u) subject-specific, and we do this by assuming $\ln(m - 1) \sim N(\mu, \eta^2)$. In the interests of parsimony, the RD parameters b_0 and b_1 , the RDS parameters $\alpha_0, \alpha_1, \delta_0$ and δ_1 , and the tremble parameters ω_0 and ω_1 , are required to take the same value for all subjects.

We estimate each model in a random effects framework, by assuming that the subject-specific parameters vary randomly across the population according to a distribution whose parameters we set out to estimate. An alternative would be to adopt a fixed effects approach and to obtain an estimate of the subject-specific parameter separately for each subject. We prefer the random effects approach for two reasons. First, there are many less parameters to estimate in the random effects approach. Second, with the fixed effects approach, we encounter difficulties with those subjects who chose \mathbf{q} in almost every non-dominance problem and with the smaller number who almost always chose \mathbf{p} . The problem is that it is impossible to estimate such a subject's attitude to risk, and so these subjects need to be discarded in order to estimate the model. No such problems are encountered when random effects are used, so all subjects can be included in the estimation. Our estimation methods are described in the Appendix.

8. Results

The stochastic models which we have presented can be classified by their core theories (EU, RD, and RDS) and by their stochastic specifications: random preference with trembles (RP-T), Fechner without trembles (F); and Fechner with trembles (F-T). Hence-

forth, we shall use the core theory acronyms as subscripts to the stochastic specification acronyms. So, for example, $RP-T_{EU}$ represents the random preference model with trembles applied to expected utility theory.

We have estimated each of the six possible combinations of the core theories EU and RD and the three stochastic specifications. To show why we chose our particular parameterization of RD, we also report an estimation of $RP-T_{RDS}$.

Each model was estimated separately for the two subsamples of subjects. Each subsample contains 46 subjects, so the sample size is always $46 \times 90 = 4140$. The results are presented in Tables 3a and 3b. All entries are maximum likelihood estimates of parameters, with asymptotic standard error estimates in parentheses. Each column refers to a different model.

We begin by addressing the question of which of these models is best at explaining the observed data. We can go some way towards answering this question by adopting the straightforward criterion of the maximized log-likelihood. Where one model is nested within another, we can compute a likelihood ratio test statistic (LR) as twice the difference of the maximized log-likelihoods between the two models, whose distribution is $\chi^2(q)$ under the truth of the nested model, where q is the number of restrictions under test. We have conducted several such nested LR tests, and the results are displayed in Figure 1(a) and (b), in which single-pointed arrows represent movement from a nested model to a nesting model.

As it happens, all of the tests we have conducted are tests of two restrictions. This means the null distribution is always $\chi^2(2)$. If the LR test statistic exceeds 6.0, there is evidence in favour of the nesting model; if the statistic exceeds 13.8, the evidence is overwhelming. Figure 1(a) and (b) show that all of the LR tests conducted result in overwhelming rejections of the restrictions under test, and therefore overwhelming evidence in favour of the nesting model in each case.¹⁴ These results have two distinct implications. In relation to core theories, these results establish that the RD model has significantly greater explanatory power than the EU model. In relation to stochastic specifications, they establish that the addition of a tremble mechanism to the Fechner model leads to a significant increase in explanatory power. Thus, in comparing the performance of the Fechner and random preference approaches, we should compare the models $F-T_{RD}$ and $RP-T_{RD}$.

A formal comparison of these two models is less straightforward than the comparisons discussed in the last paragraph, because these two models are non-nested. For the purpose of this comparison, we make use of Vuong's (1989) non-nested likelihood ratio test, which is described in the Appendix. The results of using this test are shown in Figures 1a and 1b, where double-pointed arrows indicate non-nested tests, with positive values indicating evidence in favour of random preference models relative to Fechner models, and in favour of RD models relative to RDS ones. The tests comparing $F-T_{RD}$ and $RP-T_{RD}$ result in a Z-statistic of +2.57 for the £20 group, and of +6.19 for the £30 group. In each case, this is strong evidence that $RP-T_{RD}$ is closer to the true data generating process than $F-T_{RD}$. In other words, when combined with the most successful core theory, the most successful model of stochastic variation is the random preference model with trembles.

Table 3a. Maximum likelihood estimates: £20 group

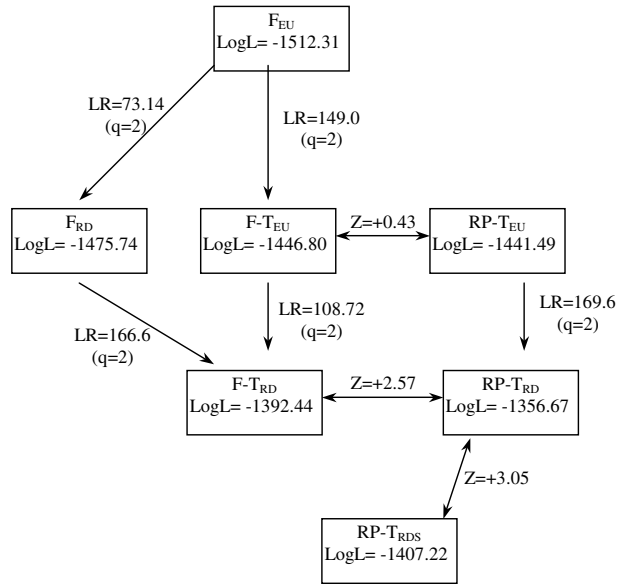
	F_{EU}	$F-T_{EU}$	$RP-T_{EU}$	F_{RD}	$F-T_{RD}$	$RP-T_{RD}$	$F-T_{RDS}$	$RP-T_{RDS}$
μ	-1.33(0.10)	-1.34(0.10)	-1.46(0.14)	-1.53(0.10)	-1.55(0.095)	-1.60(0.126)	-1.96(0.154)	-1.96(0.154)
η	0.67(0.08)	0.66(0.09)	0.91(0.12)	0.64(0.08)	0.595(0.077)	0.77(0.103)	0.97(0.125)	0.97(0.125)
σ_e	0.31(0.01)	0.23(0.01)		0.30(0.01)	0.209(0.013)			
σ_u			0.72(0.04)			0.62(0.031)	0.80(0.04)	0.80(0.04)
ω_0		0.07(0.02)	0.07(0.02)		0.050(0.018)	0.043(0.016)	0.053(0.02)	0.053(0.02)
ω_1		-0.016(0.007)	-0.010(0.006)		-0.010(0.008)	-0.007(0.007)	-0.008(0.007)	-0.008(0.007)
b_0				0.167(0.021)	0.171(0.017)	0.150(0.018)		
b_1				-0.009(0.003)	-0.006(0.002)	-0.006(0.002)		
α_0							0.51(0.036)	0.51(0.036)
α_1							-0.003(0.001)	-0.003(0.001)
δ_0							0.145(0.057)	0.145(0.057)
δ_1							-0.024(0.015)	-0.024(0.015)
n	46	46	46	46	46	46	46	46
T	90	90	90	90	90	90	90	90
Log L	-1512.31	-1446.80	-1441.49	-1475.74	-1392.44	-1356.67	-1407.22	-1407.22

Note: Asymptotic standard errors in parentheses.


Table 3b. Maximum likelihood estimates: £30 group

	F_{EU}	$F-T_{EU}$	$RP-T_{EU}$	F_{RD}	$F-T_{RD}$	$RP-T_{RD}$	$RP-T_{RDS}$
μ	-0.68(0.10)	-0.72(0.11)	-0.72(0.13)	-0.93(0.09)	-1.00(0.11)	-1.06(0.12)	-0.88(0.15)
η	0.63(0.07)	0.72(0.09)	0.87(0.11)	0.61(0.07)	0.678(0.085)	0.79(0.097)	0.93(0.11)
σ_ε	0.29(0.01)	0.19(0.01)		0.27(0.01)	0.144(0.010)		
σ_u			0.88(0.049)			0.67(0.039)	0.88(0.05)
ω_0		0.19(0.04)	0.107(0.046)		0.176(0.032)	0.114(0.034)	0.08(0.04)
ω_1		-0.016(0.005)	-0.035(0.013)		-0.010(0.004)	-0.023(0.008)	-0.025(0.01)
b_0				0.195(0.017)	0.183(0.016)	0.202(0.018)	
b_1				-0.007(0.002)	-0.005(0.001)	-0.006(0.002)	
α_0							0.37(0.06)
α_1							-0.013(0.005)
δ_0							0.15(0.03)
δ_1							0.0003(0.004)
n	46	46	46	46	46	46	46
T	90	90	90	90	90	90	90
Log L -2026.54		-1959.93	-1870.35	-1911.23	-1795.00	-1705.35	-1799.37

Note: Asymptotic standard errors in parentheses.



(a)



(b)

Figure 1. Nested and non-nested hypothesis tests. (a) £20 group. (b) £30 group.

When the Vuong test is used to compare RP-T_{RD} and RP-T_{RDS}, the Z-statistics are +3.05 (£20 group) and +5.65 (£30 group). This is the evidence for our claim that our specification of the probability transformation function in (11a) and (11b) has greater explanatory power than the more conventional specification (18).

Using our estimated RP-T_{RD} model, we turn to the relationship between experience and stochastic variation. Our estimates of σ_u , the parameter which represents within-subject variation, are significantly positive. At the specification search stage, we found no evidence that the value of this parameter is affected by experience. The implication is that stochastic variation includes a constant component, which RP theory interprets as the individual's uncertainty about her own 'true' preferences. However, there is a second component, represented by the tremble parameters ω_0 and ω_1 . The fact that our estimates of ω_0 are significantly positive for both subsamples shows that a model without a tremble mechanism would be inadequate. The effect of experience on trembles is picked up by ω_1 . Our negative estimates of this parameter (significantly negative for the £30 group) imply that the kind of errors that are modelled by the tremble mechanism become less frequent as subjects gain experience. Substituting our estimates for the £30 group into (10), the predicted relationship between the tremble probability and experience is:

$$\hat{\omega}_t = 0.114 \exp(-0.023\tau_t) \quad (21)$$

This formula implies that the value of ω is 0.111 at the beginning of the experiment ($\tau = 1$) but falls to 0.014 by the end ($\tau = 90$). This effect is equivalent to a fall in the reversal rate from 0.10 to 0.01.

In our model, experience also impacts on the bottom-edge effect through the parameters b_0 and b_1 . Both our estimates of b_0 are significantly positive, indicating a bottom-edge effect that is inconsistent with EU but can be explained by RD. However, both our estimates of b_1 are significantly negative, indicating that this effect declines with experience. Substituting our estimates for the £30 group into (15), the relationship between the parameter b and experience is:

$$\hat{b}_t = 0.202 \exp(-0.006\tau_t) \quad (22)$$

Recall that b is a measure of the degree to which the probability of the best outcome (when close to zero) is overweighted; in EU, $b = 0$. The formula (22) implies that the value of b falls from 0.201 to 0.118 during the course of the experiment; however it remains significantly different from zero even after 90 decisions have been made.¹⁵ In our specification search, we found that the attitude-to-risk parameter μ did not show any similar dependence on experience. Thus, the apparent increase in risk-aversion during the course of the experiment, mentioned in Section 4, seems to be due principally to the decay of the bottom-edge effect.

Extrapolating from our estimated model and taking the limit as $\tau \rightarrow \infty$, we can identify a simple model, towards which choice behaviour is apparently tending. The core theory of this model is EU, and the stochastic specification is RP. For any given individual, this model is fully described by m and σ_u^2 , the parameters of the probability




Figure 2. Von Neumann–Morgenstern utility functions.

distribution of the quantity u for that individual. For the population as a whole, it is described by μ and η , i.e., the parameters of the population-wide probability distribution of m , and by σ_u^2 . From our estimates of μ and η , it is straightforward to deduce, using an inversion of (A1), the quartiles of m over the population. For the £20 group, the quartiles of m are $(Q1, Q2, Q3) = (1.119, 1.202, 1.340)$. For the £30 group, the quartiles are $(1.203, 1.346, 1.590)$. In Figure 2, the solid piecewise line is the estimated median utility function for the median individual in the population, and the two dotted piecewise lines are the median utility functions for the first and third quartile individuals. The quartile utility functions we have estimated turn out to be remarkably close to power functions of the form $u(x) = (x/10)^\beta$, where x is the subject's winnings (measured in pounds) from the experiment,¹⁶ with the value of β approximately 0.16 for the first quartile, 0.27 for the median, and 0.42 for the third quartile.¹⁷ The dotted and dashed curve in Figure 2 represents the power function for the median individual.

9. Dominance

One major difference between the random preference and Fechner models of stochastic variation is in their implications for choices in dominance problems. For the moment,

we consider each of these models in the absence of a tremble mechanism. When combined with a core theory in which preferences respect stochastic dominance, the random preference model implies that dominating prospects are always chosen. In contrast, the Fechner model permits violations of dominance. In the latter model, the frequency of such violations depends on the difference between the subjective values of the two prospects and on the variance of the error term ε . Thus, *for dominance problems in general*, dominance violations cannot be infrequent unless the variance of ε is small; and if that variance is small, reversal rates on all problems will be low. In other words: the Fechner model cannot easily accommodate *both* high reversal rates on non-dominance problems *and* low rates of dominance violation. Recall that in our data, the reversal rate on non-dominance problems (0.183) is very much higher than the rate of dominance violation (0.014). It is therefore natural to ask how far the superiority of the random preference specification in our tests is attributable to the dominance problems.

To investigate this question, we estimated the F_{RD} , $F-T_{RD}$, RP_{RD} , and $RP-T_{RD}$ models using only the data from the 40 non-dominance problems. Notice that the elimination of the dominance problems makes it possible for us to estimate a random preference model without a tremble mechanism (i.e., RP_{RD}). The results are shown in Figure 3(a) and (b), which use the same notation as Figure 1(a) and (b). The Fechner and random preference specifications are now more evenly matched. If tremble mechanisms are not added, the random preference model (RP_{RD}) still performs somewhat better than the Fechner model (F_{RD}) for the £30 group, but there is almost no difference between the models for the £20 group. If tremble mechanisms are added, the ranking is reversed:

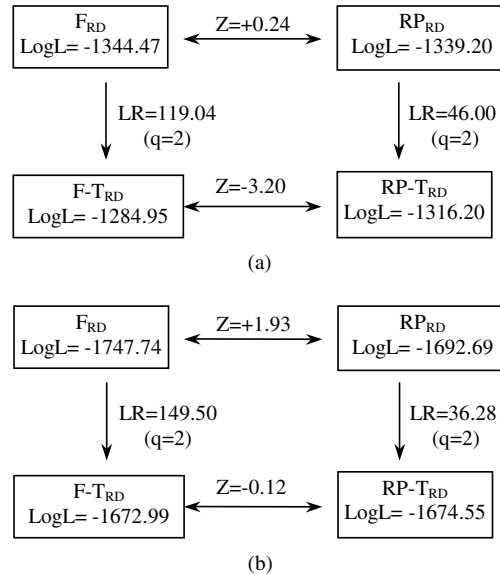


Figure 3. Model comparisons in the absence of dominance problems. (a) £20 group. (b) £30 group.

the Fechner model ($F-T_{RD}$) performs significantly better for the £20 group, while there is almost no difference for the £30 group.

It seems, then, that the superiority of the random preference specification in our principal tests owes something to the dominance problems; for non-dominance problems, there is little to choose between random preference and Fechner specifications. Nevertheless, dominance problems surely fall within the scope of any theory of stochastic decision making. The low rate of dominance violations in our data must count as evidence against the Fechner approach to explaining stochastic variation in choice.

10. Conclusions

In relation to the set of problems presented in this experiment, our subjects' decision-making behaviour can be modelled as converging towards a random preference model of stochastic choice with expected utility as the core theory. However, we are inclined to be cautious about interpreting this as evidence in support of the *general* validity of a stochastic form of expected utility theory. To draw this conclusion would be to claim that the utility functions we have estimated represent subjects' 'true' preferences, and can be used to predict the decisions that those subjects would make in relation to *other* choice problems (after sufficient experience of the relevant choice environment). Our estimation used data from pairwise choices involving a very limited set of payoffs. The experimental data sets used by other researchers who have estimated competing theories of risky choice are similar to ours in this respect. Before concluding in favour of stochastic expected utility theory, we would need to be sure that the parameters of our estimated models are stable across a wider range of payoffs and a richer set of decision tasks.

However, as far as the analysis of the current data set is concerned, we can give some clear answers to the questions which prompted our research. Our first objective was to compare the explanatory power of alternative models of the stochastic component in individual choice under risk. Although the constant error or tremble model, as used by Harless and Camerer (1994), is inadequate as a *general* theory of stochastic choice, our results show that the explanatory power of other stochastic models can be significantly increased by the addition of a tremble term. One significant implication of this result is that stochastic variation in decision making does not have a single cause. We focused on two other approaches to modelling stochastic choice: the Fechner model (as used by Hey and Orme, 1994) and the random preference model (as proposed by Loomes and Sugden, 1995). We found that the random preference model performed significantly better than the Fechner model in explaining our data as a whole; when choices between stochastically dominating and stochastically dominated prospects were excluded from the data set, neither model was clearly superior to the other.

Our second objective was to investigate the relationship between stochastic variation and decision-making experience. We found that the frequency of Harless-Camerer trembles decayed rapidly as subjects gained experience: after subjects had faced the 90 decision problems of our experiment, trembles had almost disappeared. This suggests that the tremble component of stochastic variation may be a form of error, which indi-

viduals can learn to avoid. However, that part of the variation that can be explained by the random preference model did *not* decay with experience. The latter variation might be interpreted as the result of imprecision in people's preferences. We speculate that such imprecision may be an inherent and stable property of preferences rather than a transitory phenomenon. We suggest that future theoretical and empirical work on decision making under risk should give serious attention to the formulation, testing and refinement of decision theories in stochastic form.

Appendix

A.1. Estimation

The random effects models we construct are similar to the random effects probit model of Avery et al. (1983). We let n denote the number of subjects in the sample, and we index them by i . In this Appendix, we index the binary variable y by both i and t , so that $y_{it} = 1$ if subject i chose \mathbf{p} in problem t . In the random preference models, we assume that the median of u varies across the population according to a lognormal distribution, so that:

$$\ln(m - 1) \sim N(\mu, \eta^2) \quad (\text{A1})$$

We may then construct the log-likelihood function:

$$\text{Log } L = \sum_{i=1}^n \ln \left[\int_1^{\infty} \prod_{t=1}^T [\text{pr}(y_{it} = 1 | \tilde{m})^{y_{it}} \text{pr}(y_{it} = 0 | \tilde{m})^{1-y_{it}}] f_m(\tilde{m}; \mu, \eta) d\tilde{m} \right] \quad (\text{A2})$$

where $f_m(\tilde{m}; \mu, \eta)$ is the log-normal density function for the random variable m , evaluated at the value \tilde{m} . The conditional probability term $\text{pr}(y_{it} = 1 | \tilde{m})$ is defined in either (7) or (14), depending on which of the random preference models is under consideration.

In the Fechner models, it is u itself which is assumed to vary across the population. We therefore specify:

$$\ln(u - 1) \sim N(\mu, \eta^2) \quad (\text{A3})$$

and the log-likelihood function is:

$$\text{Log } L = \sum_{i=1}^n \ln \left[\int_1^{\infty} \prod_{t=1}^T [\text{pr}(y_{it} = 1 | \tilde{u})^{y_{it}} \text{pr}(y_{it} = 0 | \tilde{u})^{1-y_{it}}] f_u(\tilde{u}; \mu, \eta) d\tilde{u} \right] \quad (\text{A4})$$

where $f_u(\tilde{u}; \mu, \eta)$ is the log-normal density function for the random variable u , evaluated at \tilde{u} , and $\text{pr}(y_{it} = 1 | \tilde{u})$ is defined in either (8), (9), (16) or (17), depending on which particular model is under consideration.

We use the MAXLIK routine in GAUSS to maximise the log-likelihood functions (A2) and (A4). We always use the BHHH algorithm (Berndt et al., 1974) with analytic first derivatives. The GAUSS quadrature routine INTQUAD1 is used to evaluate the integral appearing in the formula for the log-likelihood. Convergence to the MLE is usually achieved in around 12 iterations. The asymptotic standard error estimates reported in Tables 3a and 3b are obtained from an estimated covariance matrix which is returned by the MAXLIK routine.

A.2. Vuong's non-nested likelihood ratio test

Consider any two non-nested models 1 and 2. Let \hat{f}_i be the estimated probability of observing the T actual choices made by subject i , on the assumption that model 1 is the true model. Let \hat{g}_i be the estimated probability of observing the same T choices on the assumption that model 2 is the true model. The Vuong test is based on the quantity D , defined by:

$$D = n^{-1/2} \sum_{i=1}^n \log \left(\frac{\hat{f}_i}{\hat{g}_i} \right) \quad (\text{A5})$$

D is similar to the log-likelihood ratio of the two models, but since the models are non-nested, it can be of either sign. To implement the test, we need to estimate the variance of D . An appropriate variance estimator is:

$$\hat{V} = \frac{1}{n} \sum_{i=1}^n \left[\log \left(\frac{\hat{f}_i}{\hat{g}_i} \right)^2 - \left[\frac{1}{n} \sum_{i=1}^n \log \left(\frac{\hat{f}_i}{\hat{g}_i} \right) \right]^2 \right] \quad (\text{A6})$$

The Vuong test statistic is then:

$$Z = \frac{D}{\sqrt{\hat{V}}} \quad (\text{A7})$$

As proved by Vuong (1989), the statistic Z defined in (A7) has a limiting standard normal distribution under the hypothesis that the two models are equivalent. A significantly positive value of Z indicates that model 1 is closer to the true data generating process than model 2, while a significantly negative value of Z indicates the converse.

Acknowledgments

This research was supported by the Economic and Social Research Council (award numbers L 122 251 024 and L 211 252 053). Robert Sugden's work was also supported by the Leverhulme Trust. We are grateful to Enrica Carbone, Andrew Chesher, John Hey, Chris Orme, Chris Starmer, the Editor and an anonymous referee for comments, and to Norman Spivey who programmed the experiment.

Notes

1. See Starmer and Sugden (1989), Camerer (1989), Hey and Orme (1994, p. 1296), and Ballinger and Wilcox (1997).
2. The pioneering studies are those of Harless and Camerer (1994), Hey and Orme (1994), and Carbone and Hey (1994).
3. For more on this, see Loomes and Sugden (1995).
4. We do not assume that decision-makers actually go through these three stages when they deliberate.
5. We assume that the probability density function over Z has continuity properties such that, for all \mathbf{p}, \mathbf{q} : $\text{pr}[V(\mathbf{p}, \hat{\mathbf{z}}) - V(\mathbf{q}, \hat{\mathbf{z}}) = 0] = 0$.
6. This model is discussed in more detail by Ballinger and Wilcox (1997).
7. Ballinger and Wilcox test and reject three parametric distributions of ε , but the hypothesis of strong stochastic transitivity (which is an implication of all Fechner models, irrespective of the distribution of ε) survives their tests.
8. For given values of x_0, x_1, x_2 satisfying $x_0 < x_1 < x_2$, the Marschak-Machina triangle represents all possible prospects \mathbf{p} in (p_0, p_2) space.
9. This regularity is not reported in Hey and Orme's paper, but was found in subsequent analysis. We are grateful to John Hey for telling us of this effect.
10. We choose this normalization because $x_0 = \text{£}0$ and $x_1 = \text{£}10$ for both subsamples, while x_2 refers to a different consequence in the two cases.
11. The cumulative form of *prospect theory* (Starmer and Sugden, 1989; Tversky and Kahneman, 1992) has the same properties. If (as in our data set) there are no negative consequences, RD and cumulative prospect theory are formally equivalent.
12. *Prospective reference theory* (Viscusi, 1989) offers an alternative way of modelling the bottom-edge effect. This theory effectively 'overweights' all small non-zero probabilities. Thus, the bottom-edge effect cannot occur in the absence of a corresponding side-edge effect. For this reason, prospective reference theory is less well adapted than RD to explaining data like ours, which show the former effect but not the latter.
13. This is a consequence of the fact that in non-dominance problems, \mathbf{p}_i is riskier than \mathbf{q}_i , while in dominance problems, \mathbf{p}_i dominates \mathbf{q}_i .
14. An issue here is that when the null hypothesis of no tremble is being tested, the LR test does not have the $\chi^2(2)$ null distribution because the parameter values under the null are on the boundary of the parameter space. However, simulations reported by Moffatt and Peters (2001) reveal that the true distribution of this statistic is not far from $\chi^2(2)$, and certainly not far enough from it to alter our conclusion that a tremble is present in every case that we consider.
15. The estimate of the bottom-edge parameter at the end of the experiment is 0.118, and the standard error of this estimate (obtained using the delta method) is 0.012.
16. Notice that this function has constant relative risk aversion with respect to *changes in* wealth. If utility is interpreted as a function of *levels of* wealth (the standard interpretation in EU), this type of constant risk aversion would occur only by an astonishing coincidence. (The coincidence is that, given the utility-of-wealth function, the subject's initial wealth should happen to be at the level that induces constant relative risk aversion with respect to changes in wealth.) Thus, this function's success in fitting the data suggests that preferences should be defined relative to a person's status quo position, as in prospect theory.
17. After observing this regularity, we estimated the RP-T_{RD} model using pooled data from the £20 and £30 groups, with the power function specification $u(x) = (x/10)^\beta$. Our estimate of the median value of β for the median individual was 0.22. Space constraints prevent us from describing this estimation in detail.

References

- Avery, Robert B., Lars Peter Hansen, and V. Joseph Hatz. (1983). "Multiperiod Probit Models and Orthogonality Condition Estimation," *International Economic Review* 24, 21–35.

- Ballinger, T. Parker and Nathaniel Wilcox. (1997). "Decisions, Error and Heterogeneity," *Economic Journal* 107, 1090–1105.
- Becker, Gordon, Morris DeGroot, and Jacob Marschak. (1963). "Stochastic Models of Choice Behavior," *Behavioral Science* 8, 41–55.
- Berndt, Ernst et al. (1974). "Estimation and Inference in Nonlinear Structural Models," *Annals of Economic and Social Measurement* 3, 653–666.
- Camerer, Colin. (1989). "An Experimental Test of Several Generalized Expected Utility Theories," *Journal of Risk and Uncertainty* 2, 61–104.
- Carbone Enrica. (1998). "Investigation of Stochastic Preference Theory Using Experimental Data," *Economics Letters* 57, 305–312.
- Carbone, Enrica and John Hey. (1994). "Estimation of Expected Utility and Non-Expected Utility Preference Functionals Using Complete Ranking Data." In Bertrand Munier and Mark Machina (eds.), *Models and Experiments on Risk and Rationality*. Dordrecht: Kluwer.
- Fechner, Gustav. (1860/ 1966). *Elements of Psychophysics*, Vol. 1. New York: Holt, Rinehart and Winston.
- Harless, David and Colin Camerer. (1994). "The Predictive Utility of Generalized Expected Utility Theories," *Econometrica* 62, 1251–1289.
- Hey, John and Chris Orme. (1994). "Investigating Generalizations of Expected Utility Theory Using Experimental Data," *Econometrica* 62, 1291–1326.
- Loomes, Graham and Robert Sugden. (1995). "Incorporating a Stochastic Element Into Decision Theories," *European Economic Review* 39, 641–648.
- Loomes, Graham and Robert Sugden. (1998). "Testing Alternative Stochastic Specifications for Risky Choice," *Economica* 65, 581–598.
- Moffatt, Peter and Simon Peters. (2001). "Testing for the Presence of a Tremble in Economics Experiments," *Experimental Economics* 4, 221–228.
- Prelec, Drazen. (1998). "The Probability Weighting Function," *Econometrica* 66, 497–527.
- Quiggin, John. (1982). "A Theory of Anticipated Utility," *Journal of Economic Behavior and Organisation* 3, 323–343.
- 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.
- Tversky, Amos and Daniel Kahneman. (1992). "Advances in Prospect Theory: Cumulative Representation of Uncertainty," *Journal of Risk and Uncertainty* 5, 297–323.
- Viscusi, W. Kip. (1989). "Prospective Reference Theory: Toward an Explanation of the Paradoxes," *Journal of Risk and Uncertainty* 2, 235–264.
- Vuong, Quang (1989). "Likelihood Ratio Tests for Model Selection and Non-Nested Hypotheses," *Econometrica* 57, 307–333.