Sheffield Economic Research Paper Series

SERP Number: 2005017

Paul Dolan, Oyeyemi Oluboyede and Jennifer Roberts*

Testing models of stochastic choice in health state valuation data

November 2005

* Corresponding author:
Department of Economics
University of Sheffield
9 Mappin Street
Sheffield
S1 4DT
United Kingdom
www.shef.ac.uk/economics

Testing models of stochastic choice in health state valuation data

Paul Dolan[1], Oyeyemi Oluboyede[1], Jennifer Roberts[2]

[1]Centre for Well-being in Public Policy (CWiPP), The

[2] Department of Economics, School for Health and Related Research and CWiPP, University of Sheffield

University of Sheffield, Western Bank, Sheffield, S10 2TN, UK

Abstract

Expected Utility (EU) theory is the standard economic model of individual preferences under uncertainty. However, observed violations of the axioms of EU have generated interest in the incorporation of a stochastic element into deterministic models of decision-making. Previous empirical investigation of the theories of stochastic choice has involved monetary gambles in risky conditions using convenience samples of students. The aim of this study is to test generalisations of these models in the context of eliciting the preferences of the general public over health states under conditions of certainty. Our findings lend support to the 'white noise' stochastic specification of Hey and Orme (1994) which indicates that the stronger the preferences of an individual, the less likely they are to make a mistake and attach a lower value to their truly preferred alternative.

JEL Classification: D0

Key words: Stochastic preferences, utility assessment, expected utility theory

1. Introduction

Expected utility (EU) is the standard economic model of decision making under risk. Violations of the axioms of EU (as stated by von Neumann and Morgenstern, 1944) have led to large numbers of generalisations, most of which are deterministic in the sense that they preclude any error on the part of individuals when they are expressing their preferences (see Harless and Camerer, 1994). There have been recent attempts to incorporate a stochastic element into deterministic models. These fall into three broad categories: 1) ‘tremble’ models, which assume that individuals have true underlying preferences that are subject to error, the rate of which is independent of the strength of preference (Harless and Camerer, 1994); 2) ‘white noise’ models, which also assume that individuals have underlying preferences that are subject to error, but which assume that the error rate is a negative function of the strength of preference (Hey and Orme, 1994); and 3) ‘imprecise preference’ models, which do not assume that individuals have true underlying preferences, but instead assume that an individual’s preferences are ‘inherently variable or imprecise’ (Loomes and Sugden, 1995).

There has been some empirical investigation of these models using the preferences of student samples over gambles with monetary pay-offs. These experimental studies have generally shown that no single model of stochastic choice can adequately explain the findings; rather, combinations of elements of different theories seem to better represent the data (Carbone and Hey, 2000; Loomes and Sugden, 1998). However, the generalisability of these findings is somewhat limited. First, the general elements of each of these models have not been tested under conditions of certainty. Even where individuals are not faced with the task of incorporating risk into their decisions, there is still scope for them to make a mistake when expressing their preferences or for their preferences themselves to be imprecise. Second, the findings of previous empirical research are based on experiments using convenience samples of students and such respondents may be representative of the wider population. Finally, empirical investigation of the various stochastic specifications has been restricted to the domain of monetary payoffs.

In this paper, we test the three types of stochastic theory under conditions of certainty using data relating to the general public's preferences over different states of health. In the next section, we define the three stochastic error specifications and Section 3 provides a brief review of the existing evidence that has compared the models. Section 4 describes the data reported in this paper, Section 5 describes the methods used to test the three models and Section 6 presents the results. Finally, Section 7 discusses the findings and their implications for future research.

2. The three stochastic models

We begin by assuming that any given subject possesses (or acts as if possessing) a 'von Neuman – Morgenstern utility function' $U(.)$ or 'von Neumann – Morgenstern utility index' $\{U_i\}$ defined over some set of outcomes. Applied to our setting, the model proceeds by specifying a set of health states (or objects) of choice. Presuming that an individual possesses a preference ordering over these health states, they may be represented by a real 'preference function' $V(.)$, in the sense that the first health state is preferred to the second if and only if it is assigned a higher value by this preference function. This deterministic model can then include a random stochastic component, which may differ according to assumptions that are made about the structure of the stochastic component or about the stability of the preference function. The three alternative specifications of the stochastic component of individuals' preference functions tested here are those described by Harless and Camerer, 1994 (HC), Hey and Orme, 1994 (HO) and Loomes and Sugden, 1995 (LS).

2.1 The HC stochastic specification

In the HC specification, it is assumed that there is a constant probability *e* that an individual will make a 'mistake' (i.e. indicate that they value health state *y* higher than state *x*, when in fact they truly value *x* higher than *y*). HC allow for the possibility that individuals make erroneous deviations from their underlying deterministic preferences, describing this as a ‘tremble’. They assume that the probability of deviation, *e,* is independent of the nature of the pairwise choice itself and that it is the same value for all individuals ($e_i = e$); a parameter within the range $0.5 \le e \ge 0$.

Therefore, in the case of health states *x* and *y*:

$$\pi(x\{x,y\}) = 1 - e \qquad \text{if } x>y \qquad (1a)$$
$$\pi(x\{x,y\}) = 0.5 \qquad \text{if } x\sim y \qquad (1b)$$
$$\pi(x\{x,y\}) = e \qquad \text{if } x<y \qquad (1c)$$

where π is the probability of an individual giving a higher value to *x* than to *y*. In (1a) this is the case and π is synonymous with the probability of the individual not making an error; in (1b) the individual is indifferent between the two health states and thus π corresponds to the 50:50 chance of them making an error; and in (1c) the individual prefers *y* to *x* and thus π is the probability of the individual making an error.

2.2 The HO stochastic specification

HO assume that a decision maker’s preference between *x* and *y* is determined by *V(x,y)*. To accommodate individuals making errors when expressing their preferences, HO suggest that actual decisions are taken on the basis of whether $V(x,y) + \varepsilon \ge 0$. Here, ε is a “genuine” error, which HO describe as the prospect of individuals’ mistakes, carelessness, inattentiveness etc (p. 1301). To make the model operational, they assume that ε is normally distributed with a mean of zero (assuming no bias in the errors that any given individual makes). The model therefore implies from the example above:

$$\pi(x,\{x,y\}) = [V(x,y) + \varepsilon \ge 0] \qquad (4)$$

The magnitude of the error variance ϑ^2 measures the error spread: the larger is ϑ the greater will be the measurement error. The model differs from the HC model because $\pi(x,\{x,y\})$ (the probability of an individual ranking *x* higher than *y* from a binary choice set (*x,y*)) is an increasing function of $V(x,y)$ i.e ($w(x)$-$w(y)$) implies that the stronger an individual's true preference, the less likely they are to make a mistake.

2.3 The RP stochastic specification

Both HC and HO assume that individuals have underlying deterministic preference functions, whereas LS do not. For LS, an individual's choices are a reflection of "the individually small and collectively unsystematic impact on preferences of many unobserved factors" (p. 643). The 'random preference' approach proposes that, for each individual, there is a set of alternative preference relations facing any particular decision problem, and an individual acts on one of these preference relations selected at random. The stochastic element thus enters the choice process at the preference selection stage: the individual is uncertain about his/her own preferences, while being certain that they have the general properties required by the deterministic theory, such as EU (although the model does allow for generalisations of EU).

Following LS, let **A** represent the set of conceivable health states, and the choice task that an individual faces is to choose one and only one of its elements from a binary choice set. Further, we let **R** be the set of all possible preference orderings on **A**. The fundamental assumption of the random preference model is the existence of an additive probability measure ϕ on **R**, so that for every $\mathbf{R`} \subseteq \mathbf{R}$ there is a probability $\phi(\mathbf{R`})$. For example, **R`** might be a set of preference relations with the property $x>y$: in this case, $\phi(\mathbf{R`})$ can be written as $\phi(x>y)$. The random preference stochastic function can thus be written as:

$$\pi(x\{x, y\}) = \phi(x > y) + \phi(x \sim y)/2 \qquad (4)$$

where the probability of an individual indicating that they prefer health state *x* over health state *y* is represented by π .

3. Existing evidence

The evidence in relation to the different models comes from choices involving monetary gambles. The results from HC, which support their model, come from 23 datasets from a range of studies which consisted of approximately 8,000 choices and 2,000 choice patterns and aggregated results (Harless and Camerer, 1994). The results from HO, which support their model, come from an experiment with 80 students at the University of York. The data relevant to testing the HO model consisted of a two part experiment (separated by a period of 7-10 days) in which individuals were asked to choose between two sets of one hundred pairwise gambles consisting of two risky prospects. Individuals could indicate whether they had a preference for either of the choices available to them, or if they were indifferent between them. To motivate truthful responses, the choice from one pairwise gamble would be chosen at random and played out for real at the end of the experiment (Hey and Orme, 1994). A direct comparison of the two sets of results is not possible however because of 'the way the data was fitted: HC fitted the data across all subjects whilst HO fitted the data subject by subject' (Carbone and Hey, 2000).

In order to allow for a direct comparison of the findings of the two empirical studies discussed above Carbone and Hey (2000) fit both error specifications to the Hey and Orme (1994) data set. They carried out a maximum likelihood estimation to test the HO model and a maximum score estimation procedure to fit the HC model. This allowed them to fit both models subject-by-subject instead of across all subjects like in the Harless and Camerer (1994) study. The authors found that for some subjects the HC error specification best explained their observed behavior best; whilst for others the HO error specification suited their behavior most adequately. They concluded that 'trying to get one error story as 'best' for all subjects would appear to be seriously misleading' (Carbone and Hey, 2000).

Loomes and Sugden (1998) compare the HC, HO and RP specifications using data from 92 students at the University of York. Respondents were asked to answer two identical sets of 45 monetary pairwise gambles with a short break in between each session. These questions were made up of 40 pairwise choices between ‘safe’ and ‘risky’ alternatives, and 5 choices where one alternative dominated the other. The incentive mechanism was that same as that adopted by HO. The authors find that there is no concrete evidence suggesting that any one of the three error specifications is fully superior to the others. Overall, the results are inconclusive about which model performs best.

4. Data

All of the existing empirical work has involved monetary gambles using experimental data from samples of students. It is informative to test the performance of the models in other contexts, in particular with samples that are more representative of the general population. The data used here is from the York Measurement and Valuation of Health (MVH) study, which sought to estimate the relative valuations attached to different health states by members of the general public (see Dolan et al, 1996 for a detailed account).

The health states in question are defined by the EQ-5D health state descriptive system (Figure 1). States are defined according to five dimensions: mobility, self care, usual activities, pain/discomfort and anxiety/depression. Each of these dimensions has three levels of severity, which essentially equate to ‘no problems’, ‘some problems’ and ‘severe problems’. The descriptive system thus generates $3^5 = 243$ theoretically possible health states. Each state has a five-digit identifier that signifies the level of severity on each dimension; for example, 11111 is a state with ‘no problems’ on all five dimensions, whereas state 33333 has ‘severe problems’ on all dimensions. Some of these states have a logical ordering while others do not; for example 11111 dominates all other states, but without preference information, we cannot know whether 11211 is better or worse than 11121, because we do not know if people ‘prefer’ to have ‘some problems’ on the usual activities or pain dimension.

Each respondent was asked to rank a set of 13 EQ-5D health states (including 11111, or ‘full health’), which were randomly drawn from a larger subset of 42 states, plus ‘unconscious’ and ‘immediate death’, which are important health outcomes not defined by the EQ-5D. Each state was then valued using the time trade-off (TTO) method, which has been widely used to elicit preferences over health states. A TTO question presents the respondent with a choice between t years (where t=10 in this study) in one of the dysfunctional EQ-5D health states or x years (where x < t) in full health. Assuming no discounting, the value of the dysfunctional state is calculated as x/t (For details of the estimates with discounting, see Dolan and Jones-Lee, 1997). In total, 3395 respondents took part in the study.

Additionally, in order to test the reliability of TTO valuations, a sub-sample of 221 respondents (who were representative of the full sample) were taken through an identical retest interview by the same interviewer around 10 weeks after the first interview (Dolan et al. 1996). It is the test-retest data that we use here. The TTO valuations can be interpreted as 13*12/2 = 78 pairwise choices and, after five respondents with unusable data and other missing values are excluded, there are 16524 useable responses.

5. Methods

The analysis is based on comparing the set of pairwise choices inferred from the initial set of TTO valuations at time t_1 (i.e. if $w(x) > w(y)$, then $x > y$) with the choices inferred from the retest set of TTO valuations at time t_2. If a different choice is made at t_1 compared to t_2, this is termed a preference 'reversal'. The HC model assumes that the observed difference is due to a 'trembling hand'. The HO explanation is that the likelihood of an observed difference is a decreasing function of difference in preference ordering. The LS specification suggests that the difference in preference ordering is the result of the randomness inherent in preference functions. The three models were tested in ways that are consistent with using EU as the deterministic core model.

The HC stochastic specification

The HC stochastic specification in conjunction with deterministic EU theory has implications that are unique. If we assume that decisions are stochastically independent (in other words, the random component incorporated in one pair-wise choice decision is independent of any other pair-wise choice decision that an individual makes), the stochastic specification of an individual who acts as if maximizing expected utility can be represented as in (2).

$$r(a\{a,b\}) \equiv r(b\{a,b\}) = 2\pi(a\{a,b\})\pi(b\{a,b\}) \qquad (2)$$ [1]

where the probability of an individual reversing their preferences between the two time periods is equal to the probability of the individual ranking health state *x* higher than health state *y* in time period one and then ranking health state *y* higher than health state *x* in time period two or vice versa.

We can calculate *e* if we assume that the probability of an individual being indifferent between two health states ($x \sim y$) is zero, then for any individual a pairwise comparison between any of the health states, following from equations (1a) and (1c):

$$r(x,y)=2e(1-e) \qquad (3)$$

where *r* represents the reversal probability and *e* is the error.

In order to test the HC model, we need to examine whether the probability of an error, *e*, is the same for all individuals and across all pairwise choice problems. Calculation of *e* from (3) requires the average reversal frequency, *r*, which is the ratio of total reversals to the total number of observed responses. Because each respondent valued a different set of health states, we can only test whether the reversal rate is constant across individuals and not whether it is constant across questions.

The HO stochastic specification

---

[1] $2\pi(a\{a,b\})\pi(b\{a,b\}) = \pi(a\{a,b\})\pi(b\{a,b\}) + \pi(b\{a,b\})\pi(a\{a,b\})$

The HO stochastic specification combined with EUT as the core deterministic theory has the implication, that for any given binary choice problem (*a*,*b*), an individual will prefer *a* to *b* if and only if $V(a,b) \geq 0$ (where $V(a,b)=w(a)-w(b)$ and $w(.)$ represents the observed rankings that an individual assign to a given health state). The probability of an individual making erroneous mistake in the process of calculating the subjective values they attach to each of the alternative prospects, is subject to the strength of the individual's true preferences. For our measure of strength of preference, we use the distance between the implied rankings of all health states at t_2. This is in line with the with the idea of Learning direction theory proposed by Reinhard Selten (Selten and Stoecker, 1986) based on the belief that after some experience, people think about what might have been a better decision last time, and then adjust their behavior in that direction (Selten and Buchta, 1999). However, regardless of whether t_1 or t_2 is used as the base case, assuming that individuals have well constructed underlying preferences, this condition will hold true in both time periods and thus there should be little difference between them. If the HO model holds true, we would be able to see a relationship between the strength of individual's preferences and the occurrence of a preference reversals for the given pair-wise comparisons. The model predicts that, the greater the distance (or the strength of underlying preferences) between health states in a given pair-wise comparison, the less likely it is that an individual will reverse their preferences when valuing the same health states over the two time periods. A probit regression was used to estimate the relationship between strength of preference and the probability of preference reversal. If the HO model were correct, we would expect a negative relationship.

The RP stochastic specification

Testing the LS stochastic specification in conjunction with EU as the core deterministic theory involves an analysis of the pairwise comparisons where a logical ordering exists. The model is founded upon the idea that individuals do not have distinctive underlying preferences, so for pairwise comparisons where there is no dominant alternative, it is reasonable for an individual to express different preferences when facing the same choice at different time periods. However, assuming all individuals in the data set have the properties required by EU theory, dominated options should never be chosen over dominant ones. Thus we investigate the number of reversals that occur when dominant ordering exists, and compare this to the number of reversals when there is no dominant ordering.

6. Results

The average reversal frequency, r = 4175/16524 = 0.253, so from (2) e = 0.148. This enables us to calculate (using a binomial distribution) the expected number of reversals according to the HC model. Table 1 compares the expected and observed reversal frequencies. The expected results suggest the most common number of reversals is 11 (made by 27 individuals). However, in the observed data the frequency of reversals is more widely distributed. The most common numbers of reversals are 18 and 12, both given by 15 respondents. Two individuals exhibit no reversals in their preferences, and four individuals exhibit more than 40 reversals. These results are plotted in Figure 2. A chi-squared test reveals a significant difference in these frequencies, suggesting that the HC hypothesis should be rejected. According to the HC model, we would expect many more respondents to have fewer reversals, with only four respondents expected to have in excess of 20 reversals, whereas in our data almost half of the respondents exhibit more than 20 reversals.

The HO model predicts that the number of reversals is inversely rated to how far apart the states are in the ranking implied by the TTO values, and this is borne out by the data. The scatter plot in Figure 3 shows the number of reversals against the absolute difference in ranks at t_2. The probit model confirms a significant negative relationship between the difference in ranks and the probability of reversal (β =-0.19 s.e = 0.004)[2]. A regression model shows that 78% of the observed results were correctly predicted.

---

[2] The full results are available from the authors on request.

The LS model requires that dominance should not be violated. Evaluation of the retest data (we chose to focus on the valuations elicited in t_2 for the same reasons as earlier stated) shows that there are 7415 observed responses to pairwise comparisons where a dominant order exists, and 9109 where there is no dominance between the health states. Of the former, 344 choices (5%) resulted in a violation of dominance. This number is significantly different from zero (p<0.05). However, preference reversals are less likely when a dominant ordering exists than when it does not. Taking the data from both time periods we were able to calculate the reversal frequency for logically ordered pairs as 18% compared to 31% per cent for non-ordered pairs; these proportions are significantly different (p<0.05).

7. Discussion

There have been recent attempts to incorporate a stochastic component into deterministic models of decision-making. Tests of these models have involved monetary gambles using samples of students. In this paper, we compare three models of stochastic choice in the context of preferences over health states elicited from a representative sample of the general public. These results are similar to those reported in Loomes and Sugden (1998): our estimate of the probability of making an error, *e*, is 0.148, which compares favorably to their estimate of 0.102. In both cases, the results reject the 'tremble' model of HC. For monetary gambles and for health state valuations, it appears that likelihood of a reversal falls as the difference in implied ranks increases, which lends some support to the HO model.

The Random Preference model of LS fails to account for the small proportion of preference reversals that occur when a dominant ordering exists. There is, however, much less chance of observing reversals when a logical ordering does exist than when it does not, and a possible avenue for future theoretical development might be, as Loomes and Sugden (2002) suggest, to add a HC 'tremble' to the Random Preference specification. However, as Roberts and Dolan (2004) have shown using the MVH data, there is an enormous amount of heterogeneity in pairwise comparisons of health states and people do not always prefer logically dominant states. Hence the methods utilized by Loomes and Sugden might not yield the same results when applied to health valuation data

Our findings regarding the HO model rests on the fundamental assumption that the valuations elicited from individuals in the retest data are an approximation of their true preferences. There are, however, good reasons to believe that preferences, especially over health states, are not stable. Whilst a considerable amount of research effort has been devoted to quantifying individual preferences, most of this work has been concerned with valuations for conventional goods and services, where the market provides a forum in which, by trial and error, the true values can be quickly recognised. The valuation of a health state is an unfamiliar task for members of the general public and thus it may be that more than two valuations are required for these preferences to be revealed. Further more, it is argued by Chu and Chu (1990) that the prevalence of preference reversals decline in a market-like setting where subjects are able to revise their choices in response to learning and feedback effects.

Having said that, studies carried out by Shiell et al (2000) and San Miguel et al. (2002) provide evidence of stability in preferences. In addition, although something may have happened to individuals in the data set to warrant them altering the values of the states between test and retest, it is likely that it would have to be a pretty dramatic event to substantively change their underlying preferences. Given the short interval between test and retest, very few respondents are likely to have experienced such an event.

Stochastic models of preferences have so far been applied to money gambles and so there is plenty of scope to test the models in other settings. We have shown how the models perform in a health state valuation context and some differences from the results reported for money are found. Future research might consider the performance of the stochastic models using non-EU deterministic theories as the core model. In the meantime, we hope our results have provided further evidence that individual preferences can be modeled with a stochastic component that accounts for the strength of preference between any two alternatives.

Acknowledgements

We would like to thank Michael Dietrich, Martin Jones and Jolian McCardy for their helpful comments.

References

Carbone E, Hey J. Which error story is best? *Journal of Risk and Uncertainty* 2000; 20:2; 161 - 176

Chu Y, Chu R. The subsidence of preference reversals in simplified and market-like experimental settings: A note. *American Economic Review* 1990; 80; 902-11.

Dolan P, Cudex C, Kind P, Williams A.. The time trade – off method: results from a general population study. Health Economics 1996; 5; 141 - 154

Dolan P, Jones-Lee M. The time trade-off: A note on the effect of lifetime reallocation of consumption and discounting. *Journal of Health Economics* 1997; 16:6; 731 - 739

Harless D, Camerer, C. The predictive utility of generalized expected utility theories. *Econometrica* 1994; 62; 1251 - 289

Health State Valuations from the British General Public, 1993 [data file]. Colchester, Essex: The Data Archive [distributor], 31 Oct. 1995. SN:34444.

Hey J, Orme C. Investigating generalizations of expected utility theory using experimental data. *Econometrica* 1994; 62; 1291 - 326

Loomes G, Sugden R. Incorporating a stochastic element into decision theories. *European Economic Review* 1995; 39; 641 - 648

Loomes G, Sugden R. Testing different stochastic specifications of risky choice. *Economica* 1998; 65; 581 - 98

Loomes G, Moffatt P, Sugden R. A microeconomic test of alternative stochastic theories of risky choice. *Journal of Risk and Uncertainty* 2002; 24:2; 103 – 130

Quiggin J. Risk perception and risk aversion among Australian Farmers. *Australian Journal of Agricultural Economics* 1981; 25; 160-169

Roberts J, Dolan P. To what extent do people prefer health states with higher values? A note on evidence from the EQ – 5D valuation set. *Health Economics* 2004; 13; 733 -737

San Miguel F, Ryan M, Scott A.. Are preferences stable? The case of health care. *Journal of Economic Behavior and Organization* 2002; 48; 1 – 14

Selten R, Stoecker R. End Behaviour in Sequences of Finite Prisoner‘s Dilemma Supergames – A Learning Theory Approach. *Journal of Economic Behaviour and organisation* 1986; 7:1; 47–70.

Selten R, Buchta J 1999. *Experimental Sealed Bid First Price Auctions with Directly Observed Bid Functions*. In: Bodescu D, Erev I, Zwick R (Eds.) *Games and Human Behaviour,* Mahwah (NJ), Lawrence Erlbaum Associates Inc., 79–102.

Schmeidler D. Subjective probability and Expected Utility without additively. *Econometrica* 1989; 57; 571-587

Shiell A., Seymour J, Hawe P, Cameron S. Are preferences over health states complete? *Health Economics* 2000; 9; 47 – 55

Figure 1: The EQ5D Descriptive System

Mobility

1. No problems walking about
2. Some problems walking about
3. Confined to bed

Self – Care

1. No problems with self - care
2. Some problems with self - care
3. Unable to wash or dress - self

Usual Activities

1. No problems with performing usual activities (e.g. work, study, housework, family or leisure activities)
2. Some problems with performing usual activities
3. Unable to perform usual activities

Pain/Discomfort

1. No pain or discomfort
2. Moderate pain or discomfort
3. Extreme pain or discomfort

Anxiety/Depression

1. Not anxious or depressed
2. Moderately anxious or depressed
3. Extremely anxious or depressed

A five digit code number relating to each of the relevant level of each dimension is assigned to each health state under study. For example for health state AP, assigned dimensions 11111, means:

1 No problems walking about
1 No problems with self – care
1 No problems with performing usual activities
1 No pain or discomfort
1 Not anxious or depressed

Table 1: Observed and Expected Reversal Frequencies

No. of reversals (out of 78[3]) made by each individual	**Expected**	**Observed**
0	0.0008	2
1	0.0107	0
2	0.0716	0
3	0.3159	0
4	1.0318	1
5	2.6596	2
6	5.6359	0
7	10.0964	3
8	15.6066	5
9	21.1416	8
10	25.4074	10
11	27.3558	6
12	26.6022	15
13	23.5229	8
14	19.0217	7
15	14.1355	14
16	9.6940	10
17	6.1577	7
18	3.6346	15
19	1.9990	10
20	1.0271	12
21	0.4941	9
22	0.2230	9
23	0.0945	3
24	0.0377	8
25	0.0142	2
26	0.0050	8
27	0.0017	6
28	0.0005	1
29	0.0002	8
30	4.5822E-05	3
31	1.2357E-05	5
32	3.1612E-06	3
33	7.6748E-07	0
34	1.7692E-07	4
35	3.8738E-08	4
36	8.0588E-09	2
37	1.5933E-09	2
38	2.9941E-10	0
39	5.3485E-11	0
40	9.0827E-12	0
41 - 50	1.7292E-12	3
51 - 60	1.1282E-21	1
61-70	2.5907E-33	0
71-78	2.7155E-48	0
Total No. of Individuals	216	216

[3] true for all but 15 individuals

Figure 2: Observed and expected reversal frequencies

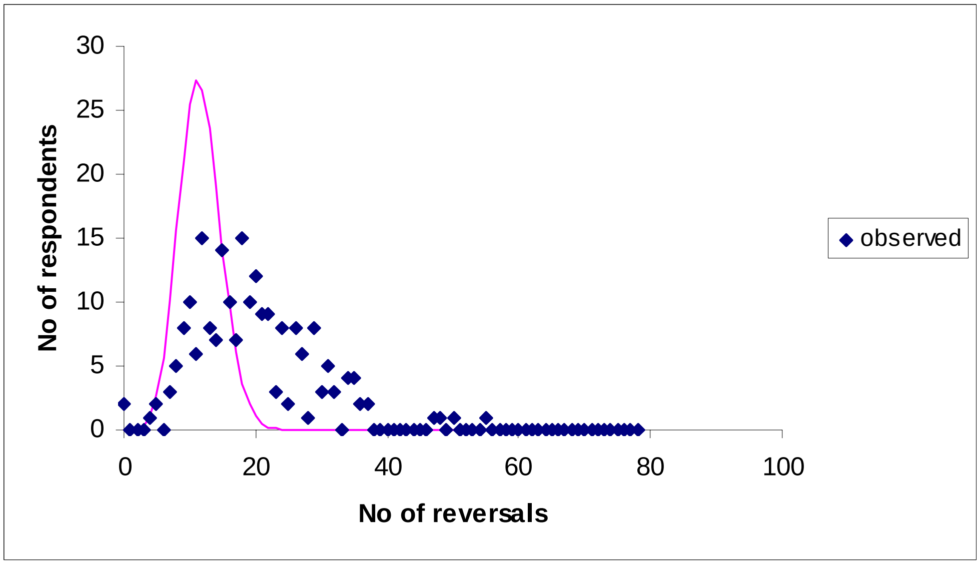

Figure 3: Reversals by distance between ranks at t_2.