On the Conjunction Fallacy in Probability Judgment: New Experimental Evidence Regarding Linda

Gary Charness, University of California, Santa Barbara
Edi Karni, Johns Hopkins University
Dan Levin, The Ohio State University

August 19, 2009

Abstract: This paper reports the results of a series of experiments designed to test whether and to what extent individuals succumb to the conjunction fallacy. Using an experimental design of Kahneman and Tversky (1983), it finds that given mild incentives, the proportion of individuals who violate the conjunction principle is significantly lower than that reported by Kahneman and Tversky. Moreover, when subjects are allowed to consult with other subjects, these proportions fall dramatically, particularly when the size of the group rises from two to three. These findings cast serious doubts about the importance and robustness of such violations for the understanding of real-life economic decisions.

Keywords: Conjunction fallacy, representativeness bias, group consultation, incentives

1. Introduction

The conjunction rule—namely, the fact that the probability of the intersection of events (that is, their conjunction) cannot exceed the probabilities of the constituent events—is one of the simplest and most basic rules of probability. Abiding by this rule is therefore a basic tenet of any theory of rational choice in the face of uncertainty. Moreover, unlike other principles of choice, the simplicity of the conjunction rule makes it compelling, particularly in situations in which its applicability is transparent. To test whether decision-makers abide by the conjunction rule, Kahneman and Tversky (1983) asked subjects to rank the likelihoods of certain conclusions that can be drawn from hypothetical personality sketches of fictitious individuals. In one version of the experiment (the "transparent" test), subjects were given the following personality sketch and asked to identify which of the two alternatives below the sketch was more probable (Kahneman and Tversky, 1983, p. 297):

> Linda is 31 years old, single, outspoken, and very bright. She majored in philosophy. As a student, she was deeply concerned with issues of discrimination and social justice, and also participated in anti-nuclear demonstrations.
> > (a) Linda is a bank teller.
> > (b) Linda is a bank teller and is active in the feminist movement.

Kahneman and Tversky (1983) report that in this version of the experiment, 85 percent of respondents indicated that (b) is more likely than (a), thereby violating the conjunction rule. Moreover, they report similar rates of violations in several variations of the question, including injecting more alternatives and using a different fictitious character. They conclude that the conjunction fallacy (assigning higher probability to the conjunction than its constituents) is prevalent in situations in which likelihood judgments are mediated by intuitive heuristics such as *representativeness* and *availability*. In the example above, the conjunction fallacy may be

1

accounted for by the impression that the conjunction is more representative of the personality described than the constituent proposition "Linda is a bank teller." In such situations, representative bias may lead subjects to reverse the likelihood ranking of the events.

Camerer (1995, p. 598) offers an alternative explanation of the high rate of violations of the conjunction principle in terms of linguistic conventions. According to him:

> … some apparent biases might occur because the specific words used, or linguistic convention subjects assume the experimenter is following, convey more information than the experimenter intends. In other words, subjects may read between the lines. The potential linguistic problem is this: in the statement "Linda is a feminist bank teller," subjects might think that this statement "Linda is a bank teller" tacitly excludes feminists; they might think it actually means "Linda is a bank teller (and not feminist)." If subjects interpret the wording this way none of the statements are conjunctions of others and no probability rankings are wrong.

Even if these findings do indeed reflect the use of underlying intuitive heuristics, as suggested by Kahneman and Tversky, there is still an issue concerning their importance for understanding real-life decisions with significant consequences when individuals can discuss their options before deciding which course of action best serves their interests. The findings of Kahneman and Tversky indicate systematic errors committed by individuals acting in isolation and without monetary incentives, which is seldom the case in real life.

However, our evidence indicates that these errors are substantially reduced by both consultation and monetary incentives. Thus, while we welcome the additional insights, realism, and enrichment coming from psychological studies, we also must guard against "jumping the gun" and changing the way economists model the world. By carefully replicating and retesting those findings we may find that these biases have substantially less significance for understanding real-life economic behavior. The result that consultation sharply reduces errors is

in a sense related to the List (2003) finding that market experience substantially reduces the endowment effect.

Charness, Karni, and Levin (2007) finds that while a substantial number of subjects choose first-order stochastically dominated alternatives when making decisions in isolation, when social interaction is allowed, the number of such violations tends to decrease dramatically with the size of the group.[1] Our paper applies the same approach to test whether the violations of the conjunction rule persist in an environment in which mild monetary incentives are offered and consultation with others permitted. To find out the extent to which these errors persist when individuals are allowed to consult with others, we tested the hypothesis that merely deliberating the alternative courses of action with other subjects helps improve the understanding of the decision problem.

To establish a baseline against which the social-interaction effects may be evaluated, we report our attempt to replicate the results of Kahneman and Tversky (1983), using the transparent test described above as the main experiment. We varied whether each participant made an unassisted or whether he or she consulted with either one or two other participants. Each of these cases was conducted both with and without small monetary incentives.[2]

While our paper considers the specific issue of the conjunction fallacy, it also contributes to the nascent literature in economics on social interaction with individual decision-making. Charness, Rigotti, and Rustichini (2007) shows that group membership affects individual choices in strategic games (the Battle of the Sexes and the Prisoner's Dilemma), but only when this group membership is made salient by factors such as shared group payoffs or the presence of an

---

[1] In that study the observed choices were group decisions; in this study the observations are individual choices.
[2] Studies such as Holt and Laury (2002) find that payoff incentives affect behavior, in that people were substantially less risk averse when choices have no financial impact. However, Ariely, Bracha, Gneezy, and Mazar (2009) point out that providing very large financial incentives may actually backfire. See Camerer and Hogarth (1999) for a survey of the effect of financial incentives on behavior.

audience.   Sutter (2008) extends these findings on group membership to a non-strategic investment experiment, where individual decisions with salient group membership are essentially the same as team decisions.   Chen and Li (2009) investigate the minimal-group paradigm and find a substantial increase in charity concerns and social-welfare-maximizing actions (see Charness and Rabin [2002]) when participants are matched with ingroup members.[3]

The next section provides a brief discussion of some of the other relevant literature, while Section 3 describes the experiment and the results.   We present an analysis of the combined data set in Section 4, and additional remarks and discussion appear in section 5.   Section 6 presents the main conclusions.

2. Related Literature

The evidence on groups versus individuals' performance on decision tasks in which there is a correct decision is inconclusive, and it seems that psychologists and economists often reach opposite conclusions.   Generally speaking, the conclusion that emerges from the psychology literature dealing with teams versus individuals is that teams typically fall well short of the "truth-wins" norm (see [Davis [1992] and references therein).[4,5]   On the other hand, Cooper and Kagel (2005) and Blinder and Morgan (2005), find that groups consistently play more strategically than individuals and generate positive synergies in more difficult games.

---

[3] The minimal-group paradigm involves groups that only have in common a factor such as a preference for one artist over another (Kandinsky or Klee). Tajfel, Billig, Bundy, and Flament (1971) were the first to demonstrate that even minimal groups were sufficient to affect behavior. However, note that the presence of minimal groups *per se* in Charness, Rigotti, and Rustichini (2007) was not enough to affect behavior in strategic environments.

[4] Lorge and Solomon (1955) developed this norm, which provides a metric for synergy. In a "eureka-type" problem where a single insight leads to the then-obvious solution, if p is the likelihood that an individual solves a problem, then a randomly-formed group with n members should solve the problem with probability $1 - (1-p)^n$.

[5] The Heath and Gonzalez (1995) study is also closer to this type of problem, where decision makers interact with others before making a final decision alone. Although they find less improvement than the truth-wins norm would predict, they conjecture that interaction forces people to explain their choices to others, thereby increasing their confidence.

Critiques of the use of heuristics to explain cognitive illusions include Gigerenzer (1991, 1997, 2005, and in press), who claims that the heuristics or biases that arise in synthetic experimental environments often do not represent how people make decisions in natural environments with natural stimuli. Cosmides and Tooby (1996) and Catrin et al. (1999) suggest that people may be reasonably good intuitive statisticians when problems are presented in terms of frequencies rather than probabilities. Inhelder and Piaget (1964) report an experiment in which 82 percent of eight-year-olds gave responses consistent with set inclusion. Fiedler (1988) finds that while more than 80 percent of undergraduate subjects commit the conjunction fallacy when asked for the probability of single events, only 20 percent do so when asked for relative frequencies.

While it appears that the conjunction fallacy is less prevalent when the representation takes a more natural form, there is still an issue concerning its robustness in the presence of incentives and the ability to consult others. Most important economic decisions—including decisions whose consequences affect individual decision units, such as buying a home or choosing a health insurance plan—are taken after some consultations with others. Yet, not much is known about the effect of social interaction on the quality of such decision-making. [6]

The literature distinguishes between situations in which decisions involve judgment (that is, there is no "correct decision") and situations involving finding the correct solution to a decision problem. The tasks in the experiments reported here are of the second type. Consequently, the measure of success is the ability of decision-makers to arrive at the correct

---

[6] A large body of literature in experimental psychology deals with group versus individual decision-making (see, for example, Heath and Gonzalez [1995] and references therein) and, in economics an entire field is devoted to the study of collective decision-making. This literature should not be confused with the social interaction studied here. In particular, we are not concerned with the issue of aggregation of individual preferences that is at the heart of social choice theory. Our focus is on individual decisions and how their quality may be affected following consultation with others.

response. When assessing the improvement, if any, afforded by consultation with other members of a group, it should be kept in mind that some improvement is due to a pure order-statistic effect. Intuitively, a team should be at least as likely to solve a problem as its most able members acting alone. Thus, by comparing the performance of freely interacting teams with this norm, it is possible to identify the presence and nature of synergies resulting from social interaction.[7]

3. The Experiment and the Findings

3.1 The experiment

Our study focuses on the version of Kahneman and Tversky's (1983) experiment referred to as the "transparent" test. Accordingly, in our experiment sheets of paper were distributed with the Linda scenario found in the introduction to their paper. To neutralize possible order effects, the statement "Linda is a bank teller" was placed first in some scenarios and second in other scenarios. The instructions can be found in the Appendix.[8]

Our experiments were conducted at the University of California, Santa Barbara (UCSB). The 361 participants were recruited from the general student population that had expressed interest in participating in such experiments. To reduce the chance that psychology majors who may have seen this experiment before contaminate the results, no psychology majors were

---

[7] Suppose each individual has a 50 percent chance of seeing the solution to a problem, which is obvious ex post, and that this likelihood is not correlated across individuals. The expected error rate is 50 percent for individuals, 25 percent for pairs, and 12.5 percent for trios. This constitutes the "truth wins" norm against which actual performance is evaluated.

[8] We also conducted sessions in which the word "single" was deleted from the description of Linda. This sharply reduced the error rate, and there was a considerably smaller effect of incentives on behavior. Once again, the decline in error rates is monotonic with respect to the number of interacting subjects: For individuals, the error rate with and without incentives was 28.1% and 35.8%, respectively; for pairs these respective rates were 20.8% and 27.0%, and for trios these respective rates were 6.2% and 12.3%. Thus, the overall error rate without incentives was reduced to about half when the word "single" was omitted, but was relatively unchanged when incentives were provided. See Charness, Karni, and Levin (2008) for a more detailed discussion.

included in the sample and no subject was allowed to participate in the experiment more than once.[9]

As this experiment is short, it was added as an exercise at the conclusion of other experiments, so that participants could mark their responses while payoffs were being calculated for the other experiments. The other experiments (a public goods experiment, an experiment on a generalized winner's curse, and an experiment on hidden information and communication) were quite different from this experiment. There seem to be no reason to suspect that the earlier experiment affected the decisions made, as there was no difference in behavior across these different experiments (and in two cases, there was no earlier experiment, because of software crashes).

Experiments were conducted either with or without incentives. In both cases, the sheets of paper were distributed, the scenario was read aloud, and no questions were taken. In the treatment with incentives, participants were informed that there was a correct answer and that anyone who chose this correct answer would receive $4. In the treatment without incentives, participants were told that they would receive $2 for filling out the questionnaire. When no other experiment was conducted, participants also received a $5 show-up fee.

These experiments were conducted with individuals, pairs, and trios of participants. In the second and third treatments, participants formed groups of two or three based on group numbers written on the (shuffled) decision sheets. They were told that they could (quietly) discuss the scenario with their partner or partners before making their choices; it was not required that responses be the same for all members of the group. After this period of consultation, each subject made his or her own choice. Subjects (or groups) were seated apart,

---

[9]After the instructions were read, participants were also asked if they had seen this question before. In the rare event that they had, they were excluded from the experiment.

so that they did not observe one another's answers. In the case of pairs or trios, the subjects were also asked to speak softly, so that members of other groups could not hear their discussion. No time constraint was imposed.

3.2 The findings

Several interesting findings emerged from this experiment (figure 1 and table 1). First, we were unable to replicate the findings of Kahneman and Tversky (1983), who found a violation rate of the conjunction rule of more than 85 percent. Our results show a much lower rate of such violations. In this study 58 percent of participants committed the conjunction fallacy, a highly significant difference, with the test of the difference in proportions (Glasnapp and Poggio 1986) yielding $Z = 4.58$, $p = 0.000$. Second, the mere addition of small incentives reduced this rate to 33 percent. This is less than 40 percent of the rate Kahneman and Tversky report.

Table 1. Violations of the Conjunction Rule

Study	Details	Incorrect answers/total sample	Error rate (percent)
Individuals			
T&K, 1983	UBC undergrads, no incentives	121/142	85.2
CKL, 2008	UCSB students, singles, no incentives	50/86	58.1
CKL, 2008	UCSB students, singles, incentives	31/94	33.0
CKL, 2008	UCSB students, total singles	**81/180**	**45.0**
Pairs			
CKL, 2008	UCSB students, in pairs, no incentives	27/56	48.2
CKL, 2008	UCSB students, in pairs, incentives	5/38	13.2
CKL, 2008	UCSB students, total in pairs	**32/94**	**34.0**
Trios			
CKL, 2008	UCSB students, in trios, no incentives	10/39	25.6
CKL, 2008	UCSB students, in trios, incentives	5/48	10.4
CKL, 2008	UCSB students, total in trios	**15/87**	**17.2**

Third, and most important, the violation rate falls significantly when communication between—and particularly among—participants is allowed. Without incentives a majority of participants acting alone—that is, without communicating with others—commit the conjunction fallacy. When participants interact with one other, the decline in the violation rate is significant at the 5 percent level ($Z = 1.75$, $p = 0.040$, one-tailed test). When participants interact with two other subjects, the error rate declines dramatically, falling to less than half that observed for individuals acting alone without incentives and to less than a third of that observed when incentives are introduced. When moving from pairs to trios, the decline is even more pronounced than when moving from individuals to pairs. This decline is significant at the 1 percent level ($Z = 2.58$, $p = 0.005$, one-tailed test).

Fourth, the violation rate of the conjunction rule declines across all group sizes as a result of incentives, suggesting that subjects put more effort into the task and, as a result, made fewer errors, when coming up with the right answer was rewarded. Using pair wise one-tailed tests, we find that the incentive effect is statistically significant for individuals ($Z = 2.58$, $p = 0.005$), pairs ($Z = 3.52$, $p = 0.000$), and trios ($Z = 1.87$, $p = 0.031$).

Finally, as would be predicted, consultation leads to a substantially higher correlation within groups than is present between groups. One way to see this is to compare the number of times that the group decision is mixed with the predicted number of times for random groups, given the individual error rate in the population.[10] We observe significant differences. For pairs with and without incentives, the test of the difference of proportions gives $Z = 1.56$ and $Z = 2.54$, respectively, with corresponding p-values (one-tailed tests) of 0.059 and 0.006; for trios with and without incentives, the test of the difference of proportions gives $Z = 1.63$ and $Z = 2.74$,

---

[10] For example, if the individual error rate is 50%, we would expect that in 25% of the cases both people in a pair make an error, in 50% of the cases exactly one person in a pair would make an error, and in 25% of the cases neither person in a pair would make an error.

respectively, with corresponding *p*-values (one-tailed tests) of 0.052 and 0.003. The levels of statistical significance do not differ greatly across pairs and trios, but they differ according to whether or not incentives are present.

4. Discussion

Our findings suggest that the rate of violation of the conjunction rule by individuals acting in isolation is much lower than reported in Tversky and Kahneman (1983). Furthermore, when subjects are given the opportunity to consult with others before choosing their responses, the error rate declines dramatically. Not surprisingly, financial incentives for providing the correct answer are effective in inducing individuals to make efforts to find the correct answer. The weight of this evidence suggests that the conjunction fallacy is a phenomenon of no serious concern for economics.

A subsidiary finding (also found in Charness, Karni, and Levin [2007]) is that there is a substantially larger drop in the error rate when the group size is increased from two to three than when it is increased from one to two: "three heads are better than one." This result is perhaps even more surprising in the Linda problem than in our earlier study, where the error reflected a failure to recognize first-order stochastic dominance. To the extent that these problems are eureka problems, one would expect the truth-wins norm to be applicable (assuming that the insight regarding the conjunction principle is i.i.d across the subject population).

However, the group error rate falls short of the truth-wins norm.[11] Although teams rarely meet this norm in the studies reported by psychologists (see Davis [1992]), this finding contrasts

---

[11] If we accept as representative the 58.1% error rate for single without incentives, this norm predicts error rates of 33.8% for pairs and 19.6% for trios (compared to the observed rates of 48.2% and 25.6%); similarly, the 33.0% error rate for singles with incentives predicts error rates of 10.9% for pairs and 3.6% for trios (compared to the observed rates of 13.2% and 10.4%).

with the results in Cooper and Kagel (2005), in which teams meet or beat the truth-wins norm in two of the three treatments. What might explain the difference? First, our experiment considers a simple decision-making task like the typical case in the psychology literature, rather than a more complex strategic game where the additional layer of complexity may lead to a greater degree of synergy. Second, Cooper and Kagel (2005) have participants repeat the same task 24-32 times, while we consider a one-shot choice. Their data suggest that experience is a key to realizing the synergies that they observe.

Finally, although the Linda choice is a clear-cut eureka problem, the patterns observed for decisions within groups do not meet the truth-wins standard. This standard require that if any person in a group has the (seemingly) clear insight and gets the right answer, all people in the group should also see the correct answer; otherwise no one in the group should get the correct answer. Thus, we should observe perfect correlation among the answers within the groups. However, while consultation does in fact lead to significantly more correlation of decisions within groups than between groups, we do not observe completely homogenous decisions among the decisions of group members. Instead, in the experiment with pairs, the number of pairs in which one subject gave the correct answer and one subject gave the incorrect answer is six out of 47 pairs. In the experiments with groups of three, the number of groups in which one subject gave the right answer and the other two subjects belonging to the same group gave the wrong answer is three (out of the 29 such groups). There is no group of three subjects in which two subjects gave the right answer and the third subject gave the wrong answer.[12]

---

[12] Without incentives, 12 pairs made all correct choices, five pairs made one correct choice, and 11 pairs made no correct choices; nine trios made all correct choices, two trios made one correct choice, and two trios made no correct choices. With incentives, 16 pairs made all correct choices, one pair made one correct choice, and two pairs made no correct choices; 14 trios made all correct choices, one trio made one correct choice, and one trio made no correct choices.

In any case, the particular way in which the truth-wins norm fails suggests the presence of relatively more synergy with trios than with pairs. The presence of a third person in a group may create a "cascading effect," in which the person who finds the correct answer has a better chance of persuading one other member of the group and the weight of the two may be sufficient to make the third member accept their conclusion even if not persuaded. In any event, the effects of group interaction may well not be monotonic with respect to group size; our result parallels that of Sutter (2005), who finds only a marginal difference between the choices of individuals and two-person groups, but a significant difference between the choices of two-person and four-person groups in an experimental guessing game.

5. Conclusions

Individuals rarely make important decisions in a vacuum: before changing jobs, buying homes, or investing in a mutual fund, they consult others (friends, colleagues, parents, spouses, experts). Moreover, in facing important decisions they recognize that the choice they make has consequences. Experimental environments in which subjects are not allowed to consult with others and not rewarded for making the right decision therefore seem far removed from those in which actual economic agents operate.

The results presented here suggest that the phenomenon described by Kahneman and Tversky (1983) as the conjunction fallacy in individual decision-making is much less prevalent than their findings indicate. More important, to the extent it exists, the conjunction fallacy is a psychological phenomenon of little consequence for economic analysis.

Some critics cite experimental evidence to fault economics for its reliance on the assumption that economic agents act rationally. Whether or not rational behavior is a useful

premise, the results presented here suggest that observing individuals making decisions in isolation and without financial incentives is not always a good guide for understanding the conduct of economic agents in real-world environments.

References

Ariely, Dan, Anat Bracha, Uri Gneezy, and Nina Mazar (2009), "Large Stakes and Big Mistakes," *Review of Economic Studies*, 76, 451-469.

Blinder, Alan and John Morgan. (2005). "Are Two Heads Better than One? An Experimental Analysis of Group versus Individual Decision-making," *Journal of Money, Credit, and Banking*, 37(5): 789-811.

Camerer, Colin, (1995) "Individual Decision-making," *The Handbook of Experimental Economics,* Eds. John Kagel and Alvin Roth. Princeton: Princeton University Press.

Camerer, Colin and Robin Hogarth (1999), "The Effects of Financial Incentives in Economics Experiments: A Review and Capital-labor-production Framework," *Journal of Risk and Uncertainty*, 19(1), 7-42.

Catrin, Rodea, Leda Cosmides, Wolfgang Hell and John Tooby. (1999). "When and Why do people avoid unknown probabilities in decisions under uncertainty? Testing some predictions from optimal foraging theory," *Cognition* 72, 269-304.

Charness, Gary, Edi Karni and Dan Levin (2007). "Individual and Group Decision Making under risk: An Experimental Study of Bayesian Updating and Violations of First-Order Stochastic Dominance," *Journal of Risk and Uncertainty*, 35, 129-48.

Charness, Gary, Edi Karni and Dan Levin (2008). "On the Conjunction Fallacy in Probability Judgment: New Experimental Evidence," mimeo.

Charness, Gary and Matthew Rabin (2002), "Understanding Social Preferences with Simple Tests," *Quarterly Journal of Economics*, 117, 817-869

Charness, Gary, Luca Rigotti, and Aldo Rustichini (2007), "Individual behavior and group membership," *American Economic Review*, 97(4), 1340-1352.

Chen, Yan and Sherry Li (2009), "Group Identity and Social Preferences," *American Economic Review*, 99(1), 431-457.

Cooper, David and John Kagel (2005). "Are Two Heads Better Than One? Team versus Individual Play in Signaling Games," *American Economic Review*, 95(2), 477-509.

Cosmides, Leda and John Tooby, 1996, "Are humans good intuitive statisticians after all? Rethinking some conclusions from the literature on judgment under uncertainty," *Cognition* 58, 1-73.

Davis, James H. (1992) "Some Compelling Intuitions About Group Consensus Decisions, Theoretical and Empirical Research, and Interpersonal Aggregation Phenomena: Selected Examples, 1950-1990." *Organizational Behavior and Human Decision Processes*, 52(1), 3-38.

Fiedler, Klaus (1988) "The dependence of the Conjunction Fallacy on subtle linguistic factors," *Psychological Research*, 50 123-129.

Gigerenzer, Gerd (1991) "How to make Cognitive Illusions Disappear: Beyond Heuristics and Biases," *European Review of Social Psychology*, 2, 83-115.

Gigerenzer, Gerd (1997) "Bounded Rationality: Models of Fast and Frugal
        Inferences," *Swiss Journal of Economics and Statistics*, 133, 201-218.
Gigerenzer, Gerd (2005) "I think therefore I Err," *Social Research*, 72, 1-24.
Gigerenzer, Gerd (in press) "What is bounded Rationality," *Supreme Court Economic Review*.
Glasnapp, Douglas and John Poggio (1986), *Essentials of Statistical Analysis for the Behavioral
Sciences*, Columbus: Merrill.
Heath, Chip and Rich Gonzalez (1995) "Interaction with Others Increases Decisions
        Confidence but Not Decision Quality: Evidence Against Information Collection Views of
        Interactive Decision-making," *Organizational Behavior and Human Decision Processes*,
        61(3), 305-326.
Holt, Charles and Susan Laury, "Risk Aversion and Incentive Effects, *American Economic
        Review*, 92(5), 1644-1655.
Inhelder, Barbel and Jean Piaget (1964) *The early growth of logic in the child: Classification and
        Seriation*, London: Routledge and Kegan Paul.
Kahneman, Daniel, and Amos Tversky (1972) "Subjective Probability: A Judgment of
        Representativeness," *Cognitive Psychology*, 3, 430-454.
List, John (2003), "Does Market Experience Eliminate Market Anomalies?," *Quarterly Journal
        of Economics*, 118(1), 41-71.
Lorge, Irving and Herbert Solomon (1955), "Two Models of Group Behavior in the Solution of
        Eureka-Type Problems," *Psychometrika*, 20(2), 139-148.
Sutter, Matthias (2005), "Are Four Heads Better Than Two? An Experimental Beauty-contest
        Game with Teams of Different Size," *Economics Letters*, 88, 41-46.
Sutter, Matthias (2008), "Individual behavior and group membership: Comment," forthcoming in
        American Economic Review.
Tajfel, Henri, Michael Billig, R. P. Bundy, and Claude Flament (1971), "Social categorization
        and intergroup behavior,'' *European Journal of Social Psychology*, **1**, 149-178.
Tversky, Amos and Daniel Kahneman (1983) "Extensional versus Intuitive Reasoning:
        The Conjunction Fallacy in Probability Judgment," *Psychological Review*, 90, 293-315.

Appendix – Instructions

[The following instructions were handed out, with the alternatives listed either in this order or the reverse order. No questions were taken]

Linda is 31 years old, single, outspoken and very bright. She majored in philosophy. As a student, she was deeply concerned with issues of discrimination and social justice, and also participated in anti-nuclear demonstrations.

Which of these two alternatives do you think is more probable?

Linda is a bank teller   _____

Linda is a bank teller and is active in the feminist movement   _____